CHALMERS
UNIVERSITY OF TECHNOLOGY

DockStream: a docking wrapper to enhance de novo molecular design

Downloaded from: <https://research.chalmers.se>, 2023-04-26 09:29 UTC

Citation for the original published paper (version of record):

Guo, J., Janet, J., Bauer, M. et al (2021). DockStream: a docking wrapper to enhance de novo molecular design. Journal of Cheminformatics, 13(1). <http://dx.doi.org/10.1186/s13321-021-00563-7>

N.B. When citing this work, cite the original published paper.

RESEARCH ARTICLE

Open Access

DockStream: a docking wrapper to enhance de novo molecular design

Jeff Guo¹, Jon Paul Janet², Matthias R. Bauer³, Eva Nittinger⁴, Kathryn A. Giblin⁵, Kostas Papadopoulos¹, Alexey Voronov¹, Atanas Patronov¹, Ola Engkvist^{1,6} and Christian Margreitter^{1*}

Abstract

Recently, we have released the de novo design platform REINVENT in version 2.0. This improved and extended iteration supports far more features and scoring function components, which allows bespoke and tailor-made protocols to maximize impact in small molecule drug discovery projects. A major obstacle of generative models is producing active compounds, in which predictive (QSAR) models have been applied to enrich target activity. However, QSAR models are inherently limited by their applicability domains. To overcome these limitations, we introduce a structure-based scoring component for REINVENT. DockStream is a flexible, stand-alone molecular docking wrapper that provides access to a collection of ligand embedders and docking backends. Using the benchmarking and analysis workflow provided in DockStream, execution and subsequent analysis of a variety of docking configurations can be automated. Docking algorithms vary greatly in performance depending on the target and the benchmarking and analysis workflow provides a streamlined solution to identifying productive docking configurations. We show that an informative docking configuration can inform the REINVENT agent to optimize towards improving docking scores using public data. With docking activated, REINVENT is able to retain key interactions in the binding site, discard molecules which do not fit the binding cavity, harness unused (sub-)pockets, and improve overall performance in the scaffold-hopping scenario. The code is freely available at <https://github.com/MolecularAI/DockStream>.

Keywords: De novo design, Generative Models, Reinforcement Learning (RL), Molecular docking, Structure-based drug discovery (SBDD)

Introduction

Machine learning has emerged as a versatile tool with potential to accelerate drug discovery. One of the quintessential problems is de novo drug design which involves finding promising candidate molecules that satisfy a multi-parameter optimization (MPO) objective. [1, 2] The major obstacle is the sheer number of possible molecules, estimated to be on the order of 10^{10^60} [23–60], effectively preventing a brute-force search of chemical space. [3] Recently, generative models have been

proposed to sample chemical space beyond what is covered by established datasets by conferring the ability to sample novel compounds. Neural network architectures including recurrent neural networks (RNNs), variational autoencoders (VAEs), generative adversarial networks (GANs), and graph neural networks (GNNs) have demonstrated success in using input data as SMILES or molecular graphs to generate promising chemical ideas [2, 4–7]. Moreover, reinforcement learning (RL) has been applied in conjunction with generative models to apply an iterative design process in which an agent (a model) learns to generate compounds achieving increasing scores.

over time [4, 8, 9] RL encourages the agent to make decisions in order to maximize a reward function which

*Correspondence: christian.margreitter@astrazeneca.com

¹ Molecular AI, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden

Full list of author information is available at the end of the article

© The Author(s) 2021. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

can be tailored to optimize drug-like properties. The synergistic application of generative models and RL has demonstrated potential for de novo drug design by providing a solution to MPO and notably mitigating the computational burden of a brute-force search of chemical space. This approach is also very versatile in terms of components that can be optimized. Notably, quantitative structure–activity relationship (QSAR) models have been applied to great effect to enrich target activity [10]. However, such models are limited in their generalizability and thus often restricted to relatively small domains of applicability which hinders the ability to sample truly novel compounds [11, 12].

On the other hand, existing physics-based methods such as molecular docking continues to be an invaluable tool to identify molecules that are promising drug candidates [13–16]. The advent of high-performance computing (HPC) has enabled in silico virtual screening (VS) to consider increasingly larger datasets. This enhanced capability has seen success in identifying more hits with diverse chemotypes and scaffolds, often desirable in structure-based drug discovery (SBDD) [16]. VS typically screens molecular libraries whose compounds are readily available, offering the potential to expedite experimentation. However, chemical ideas are inherently restricted to the pre-defined chemical space of these collections which may not satisfy the rigorous criteria in a project. Instead, VS is often used to steer human creativity to a relevant chemical sub-space where iterative design discovers the final drug compound. While strategies such as screening privileged scaffolds can narrow the search space, computational costs of in silico docking rapidly becomes prohibitive and resources may be sub-optimally allocated towards exploring unproductive chemical space [17]. Thus, efficient traversal of chemical space remains a non-trivial endeavour.

More recently, molecular docking has been incorporated into RL generative model paradigms, offering a proposed solution that integrates structural information, steering molecular design by rewarding compounds that exhibit good docking scores and circumventing some limitations of QSAR models. [18–21] However, it is often challenging to ascertain what exactly constitutes a ‘good’ docking score. Docking algorithms are inherently sensitive to the three-dimensional (3D) representation of the protein and ligands [22, 23]. Moreover, different docking configurations may be better suited for particular targets, and performance is intricately dependent on ligand embeddings to sample sufficiently many conformations to access a binding pose deemed favourable. Consequently, while docking scores are a proxy for binding free energies, an accurate prediction is beyond the capability of docking algorithms, caused to a large extent by a lack

of sampling of the receptor dynamics and insufficient treatment of entropic effects [24, 25]. Thus, in RL scenarios where docking is a component of the reward function, a poor choice of the docking algorithm or the docking score reward can accentuate the limitations of docking and misinform the agent.

Herein, we introduce DockStream, a molecular docking wrapper providing a unified interface to access a collection of ligand embedders and docking backends. DockStream provides an automated and streamlined platform to run docking experiments, supporting a large variety of possible configurations. Automated analysis of docking results expedites search for a docking configuration that performs best for the specific target and set of ligands—i.e. one that displays good correlation with experimental binding affinity or potency. The use of DockStream as a stand-alone tool for SBDD is demonstrated by benchmarking the DEKOIS 2.0 dataset which curates 81 targets with provided sets of active and decoy ligands [26]. The flexibility of DockStream to specify different docking configurations facilitates tailored protocols for different end applications as docking performance necessarily varies depending on the target system. Subsequently, DockStream is integrated with REINVENT 2.0, the recently published de novo design platform [4]. Docking provides structural information to the REINVENT agent, facilitating docking score optimization via RL and steering exploration to relevant chemical space. Using DockStream as a scoring function component enables REINVENT to design compounds that retain key interactions in the binding cavity while simultaneously discovering new ones.

Application overview

DockStream

DockStream is a molecular docking wrapper providing access to a collection of ligand embedders: Corina, LigPrep, OMEGA, and RDKit, and docking backends: AutoDock Vina, Glide, GOLD, Hybrid, and rDock (Fig. 1) [27–41]. DockStream streamlines molecular docking by handling all necessary steps under a unified platform: target preparation, ligand embedding, and docking. Target preparation usually starts with refining a protein crystal structure, often involving adding missing hydrogen atoms, defining side chain ionization and tautomeric states, and minimizing the conformational energy [42]. The final orientation of amino acid residues around a reference ligand defines the binding cavity and the docking search space. Ligand embedding refers to the generation of 3D molecular configurations for the ligands that are to be docked, thus defining an initial ligand conformational state that can affect docking search space traversal [23]. By enumerating tautomeric states and stereoisomers,

this stage can also lead to an expansion of the ligand set. Finally, docking itself adheres to specified parameters that control the rigor of a conformational search to generate so-called binding poses and scores that are typically used to inform SBDD. All of these preparatory steps must be considered in tandem to identify a useful docking configuration. DockStream provides numerous solutions to each component required for molecular docking via the supported backends whose parameters can be controlled by a single input JSON. The flexibility of the parameter definitions lends itself to applicability across diverse docking problems. Beyond providing access to a variety of backends, notable features of DockStream include:

1. Ligand tautomer/stereoisomer handling providing a thorough enumeration of atom spatial arrangements and states.
2. Different write-out modes to output poses and scores corresponding to either the best per ligand (and all its enumerations), best per enumeration (each tautomer/stereoisomer generated), or all ligand enumerations.
3. Parallelization across cores to speed-up execution (see Additional file 1: Fig. S1).

An example LigPrep configuration in an input JSON is shown below (see Additional file 1 for an example full input JSON that includes docking) [30]:

```
"ligand_preparation": {
  "embedding_pools": [
    {
      "pool_id": "Ligprep",
      "type": "Ligprep",
      "parameters": {
        "prefix_execution": "module load schrodinger/2019-4",
        "parallelization": {
          "number_cores": 4,
          "max_compounds_per_subjob": 1
        },
        "use_epik": {
          "target_ph": 7.0,
          "ph_tolerance": 2.0
        },
        "force_field": "OPLS3e"
      }
    },
    {
      "input": {
        "standardize_smiles": false,
        "input_path": "<path to SMILES file>",
        "type": "smi"
      },
      "output": {
        "conformer_path": "<output directory for conformers>",
        "format": "sdf"
      }
    }
  ]
}
```

Notable settings include:

1. The “parallelization” block specifies to parallelize LigPrep over 4 CPU cores, distributing a maximum of 1 compound per core at a time.
2. The “use_epik” block provides access to LigPrep capabilities including specifying a desired pH and a pH tolerance.
3. The “force_field” block allows one to specify the desired force field which may be particularly relevant when reproducing previous docking experiments.
4. The “input” block specifies the input path for the SMILES to be embedded by LigPrep.
5. The “output” block is optional and specifies the desired output directory for the embedded ligands.

Once an input JSON is constructed, DockStream can be executed via the command line:

```
python docker.py -conf <path to input JSON>
```

Docking evaluation

The suitability of an *in silico* docking configuration is usually assessed by its sampling or scoring power [23]. The former refers to the ability to reproduce binding poses of reference ligand-receptor co-crystal structures, often measured by the root-mean-square deviation (RMSD) of the atomic coordinates. The latter refers to the ability to rank binding affinities based on the docking score. In the ideal scenario, the docking configuration displays excellent sampling and scoring power, but this is typically not observed. Rather, docking algorithms may be better suited for either task where choice of a docking configuration will be dependent on the desired use case [39, 43]. In practice, finding an informative docking configuration is challenging owing to the inherent sensitivity of scoring functions to the 3D representation of proteins and ligands. DockStream expedites this search by automating the execution and analysis of molecular docking experiments via the benchmarking and analysis scripts, respectively (Fig. 2).

Benchmarking

Using DockStream, a single input JSON file defines all parameters required for a molecular docking experiment. Modifying parameters in the JSON defines a new docking configuration and offers complete flexibility in executing all supported backends. The benchmarking script in turn takes as input an arbitrary number of configuration JSONs, and executes all defined docking experiments sequentially, returning the docked poses and scores (Fig. 2). Analysis of resulting docking scores can be automated by the analysis script which supports three

modes (see Additional file 1 for details regarding each mode):

1. Enrichment analysis to evaluate whether a docking configuration scores active ligands better than decoys (on average).
2. Correlation analysis to evaluate how well, if at all, the docking scores correlate with experimental results.
3. Thresholds analysis to probe the distribution of docking scores and experimental results given a set of defined thresholds that serve as a hard boundary between classification as active or inactive. This functionality is well suited to investigate general separation of data points in cases where poor correlation between docking scores and experimental results is observed.

Similar to docking, a single input JSON defines all parameters required for the analysis script. Moreover, all analysis modes take as input an arbitrary number of docking outputs. Therefore, the DockStream benchmarking and analysis scripts workflow is pertinent to expedite search for a docking configuration well suited for various SBDD applications.

Enrichment analysis

The ability of a docking configuration to distinguish between active and decoy ligands is useful for hit discovery where the goal is to identify a sufficiently active ligand against a target. The logarithmic receiver operating characteristic (pROC) area under the curve (AUC) is calculated and provides a measure of enrichment. The logarithmic transform of the classical ROC curve biases the contributions to the AUC towards early actives recovered and is therefore particularly advantageous for evaluating early enrichment [44]. The pROC AUC is given by Eq. 1:

$$pROCAUC = \frac{1}{n} \sum_i^n \log_{10} \left(\frac{1}{\beta_i} \right) \quad (1)$$

where n is the total number of active and decoy ligands and β_i is the false positive rate (FPR) for when the i th active is recovered in the ordered (by docking score) list. A greater pROC AUC value demonstrates enrichment and is attributed to the ability of a docking configuration to distinguish between actives and decoys. The pROC AUC expected for random selection is 0.434 (in contrast to 0.5 for the classical ROC AUC). Importantly, unlike the ROC AUC that enforces an upper bound of 1.0, the pROC AUC is formally unbounded. In practice, however, it is constrained by the relative frequency of recovering actives amongst decoys based on HTS [44].

Correlation analysis

In lead optimization, one may be interested in increasing potency. Docking scores are often used as a crude proxy for binding free energies and thus, it would be straightforward to simply consider compounds exhibiting good scores. However, accurate prediction of binding free energies is often beyond the scope of these scoring functions and there is no guarantee the aforementioned approach will yield novel leads with enhanced potency [24, 25]. Instead, identifying a docking configuration that displays good correlation between docking scores and experimental binding affinities for a calibration set can guide iterative lead compound selection and bolster confidence in the results. The Spearman (ρ) $\in [-1, 1]$ and Kendall Tau-b (τ_B) $\in [-1, 1]$ rank correlation coefficients provide a quantitative measure between ordinal variables [45, 46]. The Spearman correlation is given in Eq. 2:

$$\rho = 1 - \frac{6 \sum_{i=1}^n d_i^2}{n(n^2 - 1)} \quad (2)$$

where n is the total number of ligands and d_i is the difference between the ranks of the i th docking score and experimental binding affinity. If the docking scores and experimental binding affinities are perfectly monotonic with respect to each other, then $\rho = 1$ (see Additional file 1 for details). The Kendall Tau-b correlation is given in Eq. 3:

$$\tau_B = \frac{C - D}{\sqrt{(C + D + T_{dock}) * (C + D + T_{exp})}} \quad (3)$$

where C and D are concordant and discordant pairs, respectively, and T_{dock} and T_{exp} are the number of ties in the docking scores and experimental binding affinities data, respectively. Similar to Spearman, if the data is perfectly monotonic with respect to each other, then $\tau_B = 1$. An important distinction between Spearman and Kendall is that the latter accounts for ties in the data [46]. This is particularly relevant in binding assays where compound activities may exceed the limit of quantification (LOQ). Consequently, these compounds are assigned the value at the limit, resulting in ties in the data. In general, Kendall is a more informative measure than Spearman in these cases.

REINVENT overview

REINVENT is a de novo design platform that uses generative models to sample compounds in SMILES format. [4, 47] SMILES generation is formulated as a natural language processing (NLP) problem whereby atoms are tokenized into machine-readable vocabulary. Compound SMILES are then generated token by token based on sequential conditional probabilities, analogous to a typical NLP problem. The underlying architecture of REINVENT uses RNNs and is based on work by Arús-Pous et al. [48]. The specific architecture used in this work is a RNN with embedding size 256, three hidden layers of 512 gated recurrent unit (GRU) cells, and a linear layer with softmax activation [49]. Model training followed two steps: first, a prior generative model was trained on the ChEMBL dataset [50]. The agent policy was initialized based on the prior before diverging into individual

experiments. The RL was conducted for 1000 epochs, thus iteratively calibrating the agent's likelihood of sampling desirable compounds (see Additional file 1 for details).

Scoring function

The scoring function, $S(x) \in [0, 1]$ offers complete flexibility in the defined properties to optimize, including topological polar surface area (TPSA), molecular weight (MW), number of hydrogen bond donors (HBD), molecular docking score, and custom QSAR models. In this work, $S(x)$ is formulated as a weighted geometric mean given in Eq. 4:

$$S(x) = \left(\prod_{i=1}^n P_i(x)^{w_i} \right)^{1/\sum_{i=1}^n w_i} \quad (4)$$

where n is the number of properties, $P_i(x)^{w_i}$ is the score calculated for the i th property for the sampled compound, and w_i is the weighting for the i th property. Properties assigned a greater weighting contribute more significantly to $S(x)$ and the weighted geometric mean necessitates every property to be reasonably satisfied. Otherwise, the total score diminishes eliciting negligible change during agent feedback. In this work, all properties were assigned a weighting of 1, denoting equal importance. Greater weightings such as 5 or 10 may be assigned to target properties to increase their relative importance in the scoring function. For instance, given a lead optimization task to enhance predicted potency and solubility, one could assign a greater weighting to docking and the corresponding QSAR model.

Diversity filters (DF)

A consequence of the iterative design cycle is that the agent may become stuck in local minima and sample increasingly similar compounds so as to exploit a single solution found. A diversity filter (DF) is applied to penalize sampled compounds based on the notion of bucket saturation (details as in the REINVENT 2.0 paper) [4, 47]. The topological DF is used in this work which takes a compound scaffold and converts all atoms into sp³ carbons. Compounds with similar scaffolds that satisfy $S(x)$ are stored in buckets with limited size (25 in this work). Buckets are filled when such compounds are repeatedly sampled. Once full, the next compound with scaffold similarity above a certain threshold will penalize the agent to disincentivize further sampling of the scaffold and encourage exploration to other chemical space. This prevents mode collapse and ensures a diverse library of generated compounds [51].

REINVENT-DockStream integration

DockStream can be specified as a component to REINVENT (see Additional file 1: Fig. S15 for more details), conferring the ability to incorporate molecular docking into $S(x)$. While it is straightforward to impose docking score optimization in REINVENT, defining a useful docking configuration remains challenging. Ideally, the docking score should be biased to either exhibit high correlation with experimental activity or at least be able to distinguish active and decoy ligands. Similar to using DockStream as a stand-alone docking tool, the benchmarking and analysis workflow (Fig. 2) can be used to automate search for such a docking configuration. There are two use cases:

1. Enrichment analysis to assess the ability of the docking configuration to distinguish between active and decoy ligands. Here, the range of docking scores for active ligands is important to define an appropriate transformation function, $P \in [0, 1]$, that normalizes the raw score such that it can be integrated into $S(x)$. Ultimately, the agent learns to propose compounds that score well using the given docking configuration.
2. Correlation analysis to assess the ability of the docking configuration to rank compounds relative to their experimental activity. Here, higher values for the correlation metrics are desired with attention given to experimental assays that artificially restrict the range of values due to the LOQ. Compounds that score well should have a greater likelihood of being potent binders.

Identifying an informative docking configuration enhances the quality of the compounds proposed by REINVENT. In scaffold hopping scenarios with docking, sampled compounds are able to retain key interactions while simultaneously exploring new ones. As long as the binding pose is similar to that of known binders, this approach can guide the agent to produce promising candidate molecules. Importantly, this can dissuade entropically driven binding which is notorious for off-target activity and toxicity, provided the generated ligand maintains similar lipophilic ligand efficiency [52].

Results

DockStream DEKOIS 2.0 evaluation

The use of DockStream as a stand-alone tool was demonstrated by benchmarking the DEKOIS 2.0 dataset which curates 81 targets each with 40 actives and 1200 decoys. The dataset was curated to represent a reasonable fair assessment of docking performance by addressing diverse properties, including removal of Pan Assay Interference Compounds (PAINS) and covalent binders in

bioactivity data, ensuring similar physico-chemical property matching between actives and decoys, and ensuring diverse chemotypes to avoid scaffold bias [26, 53]. The overall docking results reproduce and extend the observations of the original DEKOIS 2.0 benchmarking. The 5 supported docking backends: AutoDock Vina, GOLD, Glide, Hybrid, and rDock were paired with 3 ligand embedders: Corina with TautEnum, LigPrep, and RDKit with TautEnum to generate 15 distinct docking configurations [27–30, 33–38, 40, 41, 54, 55]. It is important to note that each docking configuration has many tunable parameters and screening every permutation is infeasible. Therefore, the default settings were kept with a few exceptions (see [Methods](#) for details). The benchmarking script was used to automate docking using each docking configuration for all 81 targets, resulting in 1215 docking runs and over 1,500,000 ligands docked. The analysis script was used to automate enrichment analysis and the pROC AUC values are shown in Fig. 3. The pROC AUC for a docking configuration that is no better than random selection in distinguishing actives and decoys is 0.434. Glide led to the most observed enrichment followed by Hybrid and GOLD, while the open-source docking backends, AutoDock Vina and rDock yielded less overall enrichment. In some cases, AutoDock Vina and GOLD failed to dock the entire set of ligands (see “[Methods](#)” section for details). In general, docking backends are not very sensitive to the choice of ligand embedder for the DEKOIS 2.0 targets. However, there are cases where the ligand embedder does have a significant effect on observed enrichment. For instance, AutoDock Vina with Corina and TautEnum for BCL2, Hybrid with LigPrep for RXRa, and rDock with RDKit and TautEnum for COX2, perform best within their ligand embedder series. Therefore, no single ligand embedder always performs best and while screening through n different ligand embedders necessitate n docking runs, increases in enrichment may identify a valuable docking configuration that otherwise would have been overlooked. Furthermore, no single docking backend performed best for all targets (compare COX2 Glide and AutoDock Vina ACHE for example). Alternatively, the goal may be to identify a reasonably good docking configuration with the highest throughput. Hybrid is particularly well suited for this owing to its speed and it may be the case that it also performs best or close to best performing as observed in the case of KIF11 and AR, respectively. Naturally, docking varies in performance for different target systems and ligand sets. The flexibility and ease of specifying different docking configurations and automating their execution and analysis demonstrates the value of a tool like DockStream.

REINVENT-DockStream experiments

REINVENT experiments were performed with DockStream as a specified component to demonstrate the agent can learn to optimize docking scores and hence implicitly incorporate 3D structural information. To enforce “drug-likeness”, the MPO problem was defined as optimizing the following properties with resemblance to Lipinski’s Rule of 5 [56]:

- (1) Docking score (docking backend and target specific, see Additional file 1: Figs. S17–S33)
- (2) Quantitative Estimate of “Druglikeness” (QED) Score $\in [0, 1]$ [57]
- (3) Number of hydrogen bond donors (HBD) $\in [0, 7]$ (see Additional file 1: Fig. S16a)
- (4) Molecular weight (MW) $\in [200, 575]$ (see Additional file 1: Fig. S16b)

In our experience, enforcing QED optimization made the molecular generation process more challenging but was necessary for the agent to propose “drug-like” compounds [58]. In the absence of QED, the agent learns to exploit the weak spots of docking scoring functions, proposing unreasonable molecules with artificially good docking scores. Furthermore, a well known phenomenon in RL is catastrophic forgetting in which the agent forgets previous learnings over the course of training [59]. To mitigate this, inception is used which is based on experience replay whereby top scoring compounds are randomly replayed to the agent during training and is implemented as described by Blaschke et al. [4].

REINVENT-DockStream experiments were performed for all 15 backend and ligand embedder combinations used in the DEKOIS 2.0 benchmarking (Fig. 3). 15 targets were selected (1 for each backend and ligand embedder combination) based on exhibiting a pROC AUC > 1 , indicating actives scoring better than decoys (on average) (see Additional file 1: Table S1 for details). The rationale was that given the agent learns to optimize the docking scores, selecting an appropriate docking score transformation will steer the REINVENT agent to propose compounds that are more likely to be true binders. Sampled training plots (1 for each docking backend) for the REINVENT-DockStream experiments are shown in Fig. 4 (see Additional file 1: Figs. S16–33 for all plots). The agent learned to propose compounds with increasingly favourable docking scores for every docking backend. In the case of AutoDock Vina, GOLD, and rDock (Fig. 4a, c, and e), the docking scores learning curves fluctuate and sometimes exhibit sharp spikes, in contrast to the smooth optimizations observed for Hybrid and Glide (Fig. 4b and d). The difference in agent performance is attributed to the stochasticity of AutoDock

Fig. 3 DockStream DEKOIS 2.0 benchmarking.²⁶ Docking configurations used represent every combination possible between the docking backends: AutoDock Vina (ADV), GOLD, Glide, Hybrid, and rDock, and the ligand embedders: Corina, LigPrep, and RDKit. CTE is Corina with TautEnum and RDKit TE is RDKit with TautEnum. See [Methods](#) for parameters used. The pROC AUC of a docking configuration that is no better than random selection in distinguishing active and decoy ligands is 0.434. A pROC AUC greater than 0.434 denotes enrichment

Vina, GOLD, and rDock where compounds proposed at select epochs can be notably worse while docking scores continue to improve, on average [33, 38, 41]. The sharp spikes correspond to the agent exploring chemical space as enforced by the DF and in which generated compounds fail to dock or dock poorly. While changing the docking parameters such as removing constraints or increasing sampling time can make docking more robust to these events, the momentary drop in performance did not hinder overall optimization across all docking backends. At the end of the 1000 epochs, the docking scores for all experiments converged to a point where the average score for the batch of compounds proposed by the agent are in the same range as the active ligands provided in the DEKOIS 2.0 dataset. Moreover, the top compounds (based on total score) display even more favourable docking scores, especially when $S(x)$ uses a weighted geometric mean, necessitating every property to be reasonably optimized. Thus, the top proposed compounds at the very least exhibit docking scores that are similar, if not better, than verified active ligands. The QED scores were also improved on average for every docking backend, ensuring proposed compounds are “drug-like” (Fig. 4). Similar to docking, momentary drops in QED optimization can be attributed to the stochasticity of agent exploration. Other properties that were optimized include the number of hydrogen bond donors and molecular weight (see Additional file 1: Figs. S16–33). The transformed scores for these properties across all docking backends and across all epochs ranged between 0.77 and 0.99, as expected since the prior was trained on ChEMBL which curates “drug-like” compounds (see Additional file 1: Figs. S16–S3). While most compounds in ChEMBL are Lipinski compliant, the rules are simply a guideline and there are many drug molecules that violate it [56]. The rationale in enforcing the number of hydrogen bond donors (increased to 7) and molecular weight (200–575 Da) was that keeping proposed compounds within this expanded guideline should increase the likelihood of obtaining a successful first candidate molecule when domain knowledge is typically not well established. It follows that docking scores are generally more meaningful amongst Lipinski compliant compounds, preventing molecules from being completely decorated with

-OH groups or possessing excessive hydrophobicity that exploits the docking algorithm to achieve an artificially high docking score. Furthermore, the number of unique SMILES found exhibit almost a linear relationship for all docking backends (Fig. 4) indicating agent exploration. The exception was rDock which generated less favourable and unique compounds at the start of the REINVENT experiment (Fig. 4e) which is attributed to the agent exploring similar and unfavourable chemical space at the beginning of training. Importantly, agent training is demonstrated as all properties including docking were still optimized over time. The percentage of valid SMILES for all experiments and across all epochs was between 90 and 99% which was also expected as the prior was trained on ChEMBL (see Additional file 1: Figs. S17–33). The outcome of the REINVENT-DockStream experiments are libraries of proposed compounds that are drug-like and exhibit docking scores similar or better than verified active ligands.

Agent exploration and exploitation

Compound diversity is exemplified in Fig. 5 which displays 3 out of the top 10 compounds (based on total score) of each REINVENT-DockStream experiment. In general, the top compounds were generated at later epochs and possess diverse scaffolds from sampling numerous local minima, effectively circumventing mode collapse [51]. Agent exploration is demonstrated by both the training plots which display gradual optimization of docking scores (Fig. 4) and the top sampled compounds (Fig. 5) which feature diverse scaffolds. Furthermore, agent perceived desirable scaffolds can be generated multiple times, as illustrated by GOLD and rDock (Fig. 5c and e). Interestingly, the top GOLD and rDock compounds shown in Fig. 5c and e share subtle differences and vary only by 1 or 2 atoms. The scores achieved by these compounds were nearly identical and suggests the agent implicitly learns 3D structural information when docking is incorporated into the scoring function. This is further exemplified by the Glide docking experiments performed against cyclooxygenase 2 (COX2) which is a target for anti-inflammation (Fig. 5b) [60, 61]. The proposed sulfonamide moiety (also observed ubiquitously in replicate experiments) is present in the approved

(See figure on next page.)

Fig. 4 REINVENT-DockStream agent reinforcement learning training progress for selected experiments: **a** AutoDock Vina with RDKit and TautEnum against MK2 (PDB ID: 3K33), **b** Glide with LigPrep against COX2 (PDB ID: 1CX2), **c** GOLD with RDKit and TautEnum against HDAC2 (PDB ID: 3MAX), **d** Hybrid with LigPrep against RXR α (PDB ID: 2P1RT), and **e** rDock with RDKit and TautEnum against TS (PDB ID: 1I00) (see Additional file 1: Figs. S16–33 for training plots of all experiments). ‘Known Actives’ docking scores (fitness scores for GOLD) from the DEKOIS 2.0 dataset are shown.²⁶ Docking and QED score optimization and the number of SMILES found are shown. Each epoch proposes batch size (128) number of compounds. Lower docking scores for AutoDock Vina, Glide, Hybrid, and rDock, and higher fitness scores for GOLD are considered better. The direction of the docking score optimizations reflect this difference. ‘SMILES found’ refers to the cumulative number of unique compounds proposed that pass a total score threshold. If every epoch generates only unique, valid, and favourable compounds, the plot is linear

a. AutoDock Vina RDKit and TautEnum Against MK2**'Known Actives' Docking Scores: Range: (-5.40) – (-11.20) Mean: -8.74 Median: -8.85** **b. Glide LigPrep Against COX2****'Known Actives' Docking Scores: Range: (-6.78) – (-11.71) Mean: -10.21 Median: -10.61** **c. GOLD RDKit and TautEnum Against HDAC2****'Known Actives' Fitness Scores: Range: 59.12 – 122.18 Mean: 80.62 Median: 81.85** **d. Hybrid LigPrep Against RXRα****'Known Actives' Docking Scores: Range: (-5.18) – (-19.25) Mean: -12.61 Median: -12.65** **e. rDock RDKit and TautEnum Against TS****'Known Actives' Docking Scores: Range: (-19.90) – (-42.25) Mean: -32.22 Median: -33.43** **Fig. 4** (See legend on previous page.)

a. AutoDock Vina RDKit and TautEnum Against MK2	b. Glide LigPrep Against COX2	c. GOLD RDKit TautEnum Against HDAC2	d. Hybrid LigPrep Against RXRα	e. rDock RDKit and TautEnum Against TS
'Known Actives' Docking Scores: Range: (-5.40) – (-11.20) Mean: -8.74 Median: -8.85	'Known Actives' Docking Scores: Range: (-6.78) – (-11.71) Mean: -10.21 Median: -10.61	'Known Actives' Fitness Scores: Range: 59.12 – 122.18 Mean: 80.62 Median: 81.85	'Known Actives' Docking Scores: Range: (-5.18) – (-19.25) Mean: -12.61 Median: -12.65	'Known Actives' Docking Scores: Range: (-19.90) – (-40.25) Mean: -32.22 Median: -33.43
Epoch 987 Docking Score: -12.00	Epoch 911 Docking Score: -11.65	Epoch 997 Fitness Score: 106.64	Epoch 995 Docking Score: -18.00	Epoch 780 Docking Score: -60.10
Epoch 966 Docking Score: -11.80	Epoch 677 Docking Score: -11.49	Epoch 968 Fitness Score: 106.98	Epoch 976 Docking Score: -18.60	Epoch 764 Docking Score: -62.00
Epoch 895 Docking Score: -11.90	Epoch 311 Docking Score: -11.43	Epoch 710 Fitness Score: 110.51	Epoch 931 Docking Score: -17.77	Epoch 747 Docking Score: -60.19

Fig. 5 Selected top compounds generated from the experiments shown in Fig. 4. Epoch refers to the cycle in which they were observed, and corresponding docking scores (fitness scores for GOLD) are shown. **a** AutoDock Vina with RDKit and TautEnum against MK2 (PDB ID: 3KC3), **b** Glide with LigPrep against COX2 (PDB ID: 1CX2), **c** GOLD with RDKit and TautEnum against HDAC2 (PDB ID: 3MAX), **d** Hybrid with LigPrep against RXRα (PDB ID: 2P1T), and **e**. rDock with RDKit and TautEnum against TS (PDB ID: 1I00). 'Known Actives' docking scores (fitness scores for GOLD) from the DEKOIS 2.0 dataset are shown. [26]

drugs, Celecoxib and Valdecoxib, although the latter was discontinued due to cardiovascular toxicity [61]. The agent was initialized based on a random sampling of the ChEMBL dataset and therefore did not possess any preconceived structural bias. The generation of sulfonamide containing compounds suggests agent structural awareness capable of exploiting the shape and electrostatics of the binding cavity. Overall, the top compounds (based on total score) across all experiments contained compounds that optimized every component specified in the scoring function. The general diversity observed for the top proposed compounds provide multiple solutions to the MPO problem and demonstrates the synergistic application of agent exploration with a DF enforced.

Compounds generated at similar epochs can share the same scaffold, as exemplified by the GOLD (epochs 997 and 968) and rDock (epochs 780 and 747) top selected compounds (Fig. 5c and e).

Intuitively, agent generated compounds from neighbouring epochs are expected to share a greater similarity than epochs further apart, as the agent is sampling from an area “close-by” in chemical space. Also, one would expect this trend to be more pronounced the more focused the agent gets over the course of the training. In order to quantify this behaviour, the average linkage similarity (Tanimoto) between every 5 epochs has been calculated (see “Methods” section for details) for every REINVENT-DockStream experiment (See Additional

file 1: Figs. S35–51) [62] Fig. 6b shows the resulting Tanimoto matrix for the GOLD experiment displayed in Figs. 4c and 5c. Firstly, the main diagonal of the Tanimoto matrix.

space. Furthermore, even once exploitation begins, the Tanimoto similarities do not converge to 1. The agent continues to explore chemical space as necessitated by the DF applied. In addition, one can observe that the Tanimoto similarities corresponding to surrounding epochs is greater (on average), further supporting gradual and iterative policy update. The results show that REINVENT achieves both chemical space exploration and exploitation and enforces sampling of numerous local minima.

Comparison of generated compounds to known actives and decoys

As an initial investigation into the quality of generated compounds, REINVENT-DockStream using Glide with LigPrep against COX2 was performed in triplicate

and used as the model experiment to assess the similarity of agent generated compounds to known actives and decoys. The top 150 compounds (based on total score) from each replicate were pooled and duplicates removed, resulting in 298/450 unique compounds, which amounts to about 33% overlap between the individual replicates. The Tanimoto similarities were calculated for each top generated compound compared to each known active (40 total) or decoy (1200 total), as provided in the DEKOIS 2.0 dataset, respectively. [26] The highest Tanimoto similarities were kept, resulting in distributions describing the maximum resemblance of a top generated compound (based on total score) to a known active or known decoy (Fig. 7). The prior was trained on the ChEMBL dataset and contains 33/40 and 1/1200 of the known actives and known decoys, respectively. While no actives were recovered, it is evident that the agent generated compounds with notable structural similarity to known actives as measured by a Tanimoto similarity > 0.7

Fig. 6 Average linkage similarity between epochs (Tanimoto) for every 5 epochs for REINVENT-DockStream using GOLD with RDKit and TautEnum against HDAC2 (PDB ID: 3MAX). **a** GOLD fitness score training plot (same as Fig. 4c). The vertical black line at around epoch 275 indicates the start of convergence whereby the GOLD fitness score begins to plateau. **b** GOLD Tanimoto matrix illustrating the Tanimoto similarities between batches of generated compounds across the entire 1000 epochs REINVENT-DockStream experiment (x-axis is on the scale of 5 epochs, e.g. epoch 10 and 200 correspond to epoch 50 and 1000, respectively). The main diagonal is darker shaded, indicating notable intra-batch compound similarity. The overall matrix transitions from lighter (top left) to darker (bottom right) shaded areas. Cross-referencing with subplot **a**, neighbouring epochs display notably greater Tanimoto similarity, coinciding with GOLD fitness Score convergence. The results suggest the agent begins exploitation once a state of productivity is achieved (as measured by fitness score convergence). The overall transition of the matrix demonstrates agent exploration and exploitation is darker (indicating higher similarity) relative to surrounding epochs and gradually becomes even darker, which indicates increased intra-batch similarity as the agent increasingly focuses on regions in chemical space. Moreover, the transition between the lighter shaded top left corner to the darker shaded bottom right corner exemplifies balance between agent exploration and exploitation. By cross-referencing the REINVENT-DockStream training plot for GOLD docking (Fig. 6a), one can identify that the GOLD docking score begins to converge at around epoch 275. At around epoch 55 in Fig. 6b (corresponds to epoch 275), the Tanimoto matrix gradually becomes darker shaded, indicating increased Tanimoto similarity within the same batch and neighbouring epochs batches (Fig. 6b). The results demonstrate policy update, reaching a state of productivity and enforcing the agent to begin exploitation of chemical.

(Fig. 7a). It follows from Fig. 6b that many of the top compounds contain the sulfonamide moiety which is present in the approved drugs, Celecoxib and Valdecoxib [60, 61]. The maximally similar generated compound displays a Tanimoto similarity of 0.937 and shares both the sulfonamide moiety and the *cis*-stilbene scaffold with its corresponding known active (Fig. 7a). In contrast, the highest Tanimoto similarity observed for the most similar decoy to a top generated compound is 0.631 (Fig. 7b). While the scaffold is shared, the vital sulfonamide moiety is not present. Moreover, the decoy ligands set (1200 total) is 30 × larger than the active ligands set (40 total). If the agent policy was not meaningfully updated during the generative process, the likelihood of observing a high Tanimoto similarity (>0.7) in the decoys set should be increased. Thus, the integration of DockStream with REINVENT facilitates agent convergence to relevant chemical space and the absence of actives recovered does not preclude the generated compounds from being true actives and potentially even more potent. Overall, the absence of Tanimoto similarities above 0.631 amongst the decoys and the presence of Tanimoto similarities > 0.7 amongst the actives suggests the agent is biased towards generating compounds that are more likely to be true actives.

3D structural awareness

While generated compounds can satisfy the scoring function almost perfectly, their usefulness strongly depends

on the plausibility of their binding poses, as commonly assessed in virtual screening campaigns. To further investigate agent 3D structural awareness, the predicted binding poses of one ligand from the top 10 compounds (based on total score) produced by the experiments shown in Fig. 4 were analyzed (Fig. 8). Docking against MK2, a serine and threonine kinase target for anti-inflammation, was performed using AutoDock Vina with RDKit and TautEnum. Structure–activity relationship (SAR) studies demonstrated interactions between the co-crystallized ligand and residues Lys 93 and Thr 206 are vital for potency [63]. The generated ligand retains these interactions and is predicted to exploit an additional hydrogen bonding interaction with Asn 191 (Fig. 8a). The overlap between the reference and generated ligand suggests the latter as a plausible binding pose and represents a solution obtained using only free software. Similarly, docking against COX2, a target for anti-inflammation, was performed using Glide with LigPrep [60, 61]. SAR studies demonstrated the *cis*-stilbene motif and the phenylsulfonamide moiety are paramount for potency and selectivity over COX1, respectively. The latter is attributed to the smaller residue Val 523 in COX2 in place of Ile 523 in COX1 facilitating easier access to the binding cavity [61]. The generated ligand retains the *cis*-stilbene backbone and the sulfonamide group, which is predicted to form hydrogen bond interactions with Leu 352, Ser 353, His 90, and Arg 513 (Fig. 8b). Notably, approved drugs against COX2 include Celecoxib

and Valdecoxib which also contain the aforementioned structural features [61]. Crucially, the predicted binding pose for the de novo compound overlaps significantly with the co-crystallized ligand which is another selective COX2 inhibitor (SC-558), yielding greater confidence in its plausibility [60, 61]. The same analysis was conducted for the GOLD, Hybrid, and rDock experiments shown in Fig. 4. Docking against HDAC2, a histone deacetylase, was performed using GOLD with RDKit and TautEnum. SAR studies demonstrated the interaction with zinc and large bulky groups that extend deep into the ‘foot pocket’ are important for potency [64]. The generated ligands retain interactions with zinc and Gly 154 and overlap significantly with the co-crystallized ligand (see Additional file 1: Fig. S34i). In another experiment, we used Hybrid with LigPrep as a docking component to design compounds against retinoid X receptor alpha (RXR α), a vital component in nuclear receptors. The co-crystallized ligand, CD3254 is an agonist and binds by modulating the H12 helix into its ‘active’ conformation [65]. While the generated ligands do not retain these interactions, their predicted binding poses are not in conflict with the required conformation of the H12 helix. On the other hand, a common interaction with Arg 316 which is formed by endogenous ligands such as 9-*cis* retinoic

acid is retained by the generated ligands, although there is some precedent that its role to facilitate binding is not essential (see Additional file 1: Fig. S34k) [66]. Finally, docking against thymidylate synthase (TS), responsible for converting deoxyuridine monophosphate (dUMP) into deoxythymidine monophosphate (dTMP) was performed using rDock with RDKit and TautEnum. The co-crystallized ligand is Tomudex, an approved drug which interacts with dUMP via stacking interactions and the quinazoline ring forms the only hydrogen bond interactions with Asp 218 and Gly 222 (see Additional file 1: Fig. S34o) which are retained by the generated ligands [67]. Overall, all REINVENT-DockStream experiments generated compounds that retain vital interactions and display excellent agreement with reference ligand binding poses, suggesting the agent learns implicitly 3D structural information (see Additional file 1: Table S1 for all REINVENT-DockStream experiments and Fig. S34 for all selected binding poses).

Steering chemical space exploration to diverse local minima

The diversity and convergence of REINVENT-DockStream was investigated on a larger scale. Uniform Manifold Approximation and Projection (UMAP) was

used as a dimensionality reduction technique to visualize the span of chemical space occupied by generated compounds compared to random sampling by the prior which has not been focused on any task (Fig. 9) [68]. REINVENT using Glide with LigPrep against COX2 was performed in triplicate and visualized by reducing the Morgan fingerprints (radius 3, 1024 bits) of the top 1000 and 3000 proposed compounds to 2D [69]. It is evident that the random sampling of molecules from the ChEMBL dataset occupies a single cluster (Fig. 9). In contrast, agent perceived desirable compounds occupy a much more diverse chemical space, forming numerous clusters representing different local minima. Importantly, this necessitates the compounds to be diverse, further supporting the application of a DF in molecular generation and demonstrating balance between agent exploration and exploitation. Interestingly, the chemical space spanned by the top compounds from the Glide REINVENT triplicate experiments overlap significantly. The agent is initialized at a random chemical space and is iteratively updated to satisfy the pre-defined scoring function. The UMAP results suggest that irrespective of the chemical space starting point and the stochasticity associated with Glide and the conditional probabilities

of token sampling, iterative optimization can lead to convergence. These results demonstrate steering of chemical space exploration via RL and suggest that while replicate experiments yield a more complete coverage of chemical space (deemed favourable by the agent), single experiments may not drastically compromise agent exploration.

Discussion

In this work, we present DockStream, a molecular docking wrapper providing facile access to a collection of ligand embedders and docking backends. The capabilities of the platform are extended by the benchmarking and analysis workflow which automates molecular docking and post hoc analysis. The use of DockStream as a stand-alone docking tool was demonstrated by reproducing and extending the DEKOIS 2.0 dataset with additional docking backends (Hybrid and rDock) and ligand embedders (Corina and RDKit). DockStream facilitates large-scale automation of molecular docking and supports various computational chemistry software suites. The supported analysis modes can enhance VS endeavours by expediting the identification of a suitable docking configuration. It was unsurprising that the DEKOIS 2.0 benchmarking demonstrated that docking can be particularly sensitive

Fig. 9 Uniform Manifold Approximation and Projection (UMAP) dimensionality reduction to illustrate chemical space coverage of triplicate Glide with LigPrep REINVENT-DockStream experiments against COX2 (PDB ID: 1CX2). Morgan fingerprints (radius 3, 1024 bits) for the top 1000 and 3000 compounds (based on total score) for each replicate were extracted to illustrate topology features. Chemical space coverage is stochastic in nature and thus, multiple replicates are used in practice to optimize output diversity. “Random Sampling” refers to sampling compounds from the prior generative model whose policy has not been focused on any task. **a** Top 1000 (based on total score) compounds generated by the agent in each triplicate experiment **b** Top 3000 (based on total score) compounds generated by the agent in each triplicate experiment. Some overlap between replicate experiments is observed, demonstrating steering of chemical space exploration to similar local minima

to the receptor and docking parameters, whereby small modifications in the protocol can have an enormous impact on enrichment. DockStream also outputs the binding poses, providing a convenient intermediate for downstream computational experiments such as rescoring or more accurate methods to assess binding free energy. To this end, the extensive capabilities of DockStream ligand enumeration provides excellent compatibility with the aforementioned techniques which often require numerous conformational representations [70].

The integration of DockStream with the de novo design platform, REINVENT, enhances its generative capabilities beyond standard QSAR models which often suffer from limited domain applicability [11, 12]. In particular, one is often interested in focusing on either agent chemical space exploration or exploitation, which are both supported in REINVENT. While QSAR models can perform well in the latter scenario provided the predictive model was trained on similar compounds (often belonging to the same series), extrapolating to new chemical space will likely be inaccurate. Molecular docking is chemical space agnostic and provides a more generalized solution to incorporating structural information during the generative process. However, docking will greatly increase the computational costs relative to trained QSAR models, warranting consideration over the quality of compounds generated. To assess the performance of REINVENT-DockStream, 15 different targets from the DEKOIS 2.0 dataset along with their corresponding docking configurations were selected [26]. We emphasize the choice of docking configurations that display enrichment and offer distinction between ‘good’ docking scores, providing an unambiguous endpoint for docking score optimization and increases the likelihood of generated compounds being true actives. We show that REINVENT-DockStream optimizes docking scores across all docking backends while maintaining “drug-likeness” as enforced by the QED score [57]. Furthermore, molecular weight was incorporated in the scoring function which helped circumvent artificially high docking scores displayed in relatively large molecules, owing to the sheer number of interactions possible. Large molecules may also be more prone to entropically driven binding, causing off-target effects [52]. Thus, generating compounds that conform to a reasonable size increases the overall quality of the results. Finally, the use of a diversity filter maintains agent exploration as demonstrated by the number of unique SMILES proposed [4].

In order to elucidate the balance between agent exploration and exploitation in REINVENT-DockStream experiments, the average linkage similarity between epochs (Tanimoto) was calculated [62]. The results demonstrate increased intra-batch similarity relative

to surrounding epochs, indicating agent sampling from similar chemical space within a given epoch. Moreover, as agent training proceeds, there is a clear transition to greater similarity values at later epochs, paralleling the transition from agent exploration to exploitation. To further assess generated compounds from agent exploitation, the generated top compounds (based on total score) were compared to known actives and known decoys as provided in the DEKOIS 2.0 dataset [26]. While the known actives were not recovered, the top compounds show a clear enrichment in the Tanimoto similarities over the known decoys. The results demonstrate the agent is steered to chemical space similar to that of experimentally validated active compounds. Furthermore, the quality of generated compounds was analyzed by comparing their binding poses relative to known binders. We show that generated compounds retain vital interactions in the binding cavity and exploit new ones. Crucially, the binding poses overlap with known binders, increasing confidence in the prediction. Moreover, we show that 3D structural information is implicitly incorporated during the generative process by exploiting structural motifs and moieties that conform to subtle changes in the binding cavity. Finally, the stochasticity of REINVENT-DockStream was investigated by visualizing the chemical space spanned by the top generated molecules using Uniform Manifold Approximation and Projection (UMAP) as a dimensionality reduction technique [68]. Compared to random sampling, agent generated compounds form clusters over a diverse area and generally converge over replicate runs.

Further work will focus on expanding the capabilities of DockStream to address limitations in the reliability of docking scores. Currently, the output binding poses can only be triaged manually but provide an invaluable database for further structure-based methods. In particular, negative image-based rescoring (R-NiB) has been shown to enhance enrichment in terms of activity [71]. Other possible enhancements include an explicit handling of ligand internal strain energy which is typically not well captured by docking algorithms. It has been shown that proper consideration can enhance docking enrichment and limit the number of false positives [72]. DockStream will receive continued support to ensure all supported backends are up to date, such as the integration of the recently reported AutoDock Vina 1.2.0 functionalities [73]. The DockStream codebase is provided at <https://github.com/MolecularAI/DockStream>. An additional repository with executable tutorials on supported features and workflows is provided at <https://github.com/MolecularAI/DockStreamCommunity>.

Conclusions

Generative models have been successfully applied to de novo design, generating compounds which satisfy a wide variety of properties. The integration of QSAR models aim to incorporate structural information in the generative process. However, these models can suffer from limited applicability domains. Consequently, extrapolating out of sample can be unreliable and potentially misinform the generative agent especially when chemical space exploration is desired. We built upon recent work that incorporates molecular docking in place of QSAR models for better generalization. DockStream is presented as a molecular docking wrapper which provides access to a collection of ligand embedders and docking backends. The use of DockStream as a stand-alone docking tool was demonstrated by large-scale docking efforts for a wide variety of targets and showcasing the automation of docking execution and post hoc analysis via the benchmarking and analysis workflow. The results show that it is beneficial to integrate several ligand embedders and docking backends in DockStream, so as to find a productive docking configuration for diverse end applications. The integration of DockStream with the recently published de novo design platform, REINVENT, demonstrates docking score optimization across all supported docking backends. The generated compounds retain vital interactions in the binding cavity and exploit new interactions. Extensive agreement with binding poses of known binders increases the confidence of the generated compounds and demonstrates the value of using docking to implicitly inform the generative process of 3D structural information. DockStream has the potential to be especially impactful in REINVENT experiments aiming at exploring novel chemotypes by providing a chemical space agnostic component that helps the agent learn to exploit the shape and intermolecular interactions of binding cavities.

Methods

Target preparation

Receptor crystal structures were obtained from the Protein Data Bank (PDB) with corresponding PDB IDs as specified in the DEKOIS 2.0 dataset (except SIRT2 which used PDB ID 5Y0Z so as to provide a reference ligand) [26, 74]. The PDB structures were processed by removing redundant chains possessing duplicate ligands, cofactors, and ions. All water molecules were also removed. The configurations for further processing depended on the docking backend to be used. For docking with AutoDock Vina, Glide, or rDock, the Protein Preparation Wizard in Maestro (release 2019–4) was used [75, 76]. The PDB structures were pre-processed with default parameters, followed by PROPKA hydrogen bond network

optimization at pH 7.4, and minimized using the OPLS3e force-field [77]. Alternatively, for docking with Hybrid or GOLD, PDBFixer (Conda package version 1.7) was used to remove heterogens and add missing heavy atoms and hydrogens [78]. All processed receptor structures were saved as PDB files and used as is for grid generation.

Receptor grid generation

Receptor grids were generated using the processed PDB structures and followed different docking specific configurations. For AutoDock Vina, Open Babel (Conda package version 3.1.1) was used to convert the PDB file into PDBQT format and coordinates for the binding cavity were defined using the reference ligand [79]. The grid box size was 15 X 15 X 15 Å. For Glide, the reference ligand was used to define the binding cavity in Receptor Grid Generation in Maestro [80]. Default parameters were used, specifying a grid box size of 20 X 20 X 20 Å to generate the corresponding ZIP grid file. For GOLD, the BindingSiteFromLigand class in the Docking API (release 2020.0.1 CSD) using the reference ligand with the distance parameter set to 10 Å was used to generate the grid [38, 81]. For Hybrid, the OEMakeReceptor method in the OEDocking API (release 3.0.8) using the reference ligand was used to generate the OEB grid [39]. For rDock, the rbcavity method (version 2013.1) using the reference ligand (in cases where an atom parsing error occurred, the reference ligand was first processed in the Protein Preparation Wizard) with radius set to 5 Å centered on each atom was used to generate the updated PRM grid file. The output from each of these configurations was used as is for docking.

Ligand preparation

The ligands SMILES for the actives and decoys sets were downloaded from the DEKOIS 2.0 web server [26]. The ligands were processed following different configurations depending on the ligand embedder used. For Corina with TautEnum (tool to enumerate tautomers and protonation states, version 2.0.0), default parameters were used [27–29, 54, 55]. For LigPrep, default parameters were used except for Epik, for which the target pH was set to 7.0 with a tolerated range of ± 2.0 [30]. The ligand energies were minimized using the OPLS3e force-field [77]. For RDKit (version 2020.09.3) with TautEnum, default parameters were used except for the number of maximum iterations, which was set to 600 (for Universal Force-Field, UFF) to ensure convergence for all compounds [82]. All prepared ligands were saved as SDF files and also stored internally in DockStream as RDKit molecules for subsequent docking.

Docking

The entire DEKOIS 2.0 dataset was docked using AutoDock Vina, Glide, GOLD, Hybrid, and rDock using ligands embedded by Corina with TautEnum, LigPrep, and RDKit with TautEnum, resulting in 15 docking runs per protein target and a total of 1215 docking experiments amounting to over 1,506,600 ligands docked [27–30, 33–38, 40, 41, 54, 55]. Default docking parameters were used unless specified as follows: AutoDock Vina and rDock were configured to return 2 and 12 poses, respectively. GOLD used the ChemPLP scoring function and was allowed up to 10 docking attempts with autoscale set to 0.25. Generally, the active and decoy ligands sets were comprised of 40 and 1200 ligands, respectively (39 actives and 1170 decoys for SARS-HCoV and 1199 decoys for TS). Only the best tautomer or ionization state (if applicable) per ligand as assessed by its docking score was kept. In some cases, ligands failed to be embedded or docked, which are attributed to ligand embedder failure or the docking backends' inability to find a good pose. In addition, AutoDock Vina and GOLD failed to dock all ligands for 23/81 targets (ACE, ACE2, ADAM17, COX1, COX2, CYP2A6, EPHB4, FXA, HDAC2, HDAC8, HIV1PR, HIV1RT, KIF11, MMP2, NA, PDE4B, PDE5, PNP, QPCT, SARS-HCoV, SIRT2, Thrombin, and TK) and 4/81 targets (11betaHSD1, ALR2, HIV1PR, and INHA), respectively. In the original DEKOIS 2.0 work, this was resolved via manual expansion of the binding cavity which was not explored in this work [26].

Average linkage similarity (Tanimoto)

The average linkage similarity between epochs (Tanimoto) was calculated for each REINVENT-DockStream experiment for every 5 epochs instead of every epoch to reduce the computational time while maintaining good interpretability. The average linkage similarity, L , is given by Eq. 5 [62]:

$$L(e_1, e_2) = \frac{1}{N_1 N_2} \sum_i^{N_1} \sum_j^{N_2} T(x_{e_1 i}, x_{e_2 j}) \quad (5)$$

where e_1 and e_2 are epochs 1 and 2, respectively, N_1 and N_2 are the number of molecules in epochs 1 and 2, respectively, and T is the Tanimoto similarity between molecules, denoted x . The result of Eq. 5 is a 200×200 heatmap (as the average linkage similarity was calculated for every 5 epochs) for every REINVENT-DockStream experiment (see Additional file 1: Figs. S35–51).

Abbreviations

MPO: Multi-parameter optimization; RNN: Recurrent neural network; VAE: Variational autoencoder; GAN: Generative adversarial network; GNN: Graph neural network; RL: Reinforcement learning; QSAR: Quantitative structure–activity

relationship; VS: Virtual screening; SBDD: Structure-based drug discovery; 3D: 3-Dimensional; RMSD: Root-mean-square deviation; pROC: Logarithmic receiver operating characteristic; AUC: Area under the curve; LOQ: Limit of quantification; NLP: Natural language processing; GRU: Gated recurrent unit; TPSA: Topological polar surface area; MW: Molecular weight; HBD: Number of hydrogen bond donors; DF: Diversity filter; PAINS: Pan-assay interference compounds; QED: Quantitative Estimate of Druglikeness; UMAP: Uniform Manifold Approximation and Projection; R-NiB: Negative image-based rescoring.

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s13321-021-00563-7>.

Additional file 1. Zipped archive of the JSON files used to generate the study.

Acknowledgements

We would like to thank Esben Jannik Bjerrum for discussions around docking and Graeme Robb and Martin Packer for providing invaluable help with integrating bespoke features. We are also grateful to numerous people from the REINVENT community and the Molecular AI department at AstraZeneca for fruitful feedback and comments.

Authors' contributions

JG and CM developed the code and tutorials. JG wrote the main manuscript and all other authors revised it. JPJ and MB contributed to the evaluation. CM, AP and OE supervised the work. All authors read and approved the final manuscript.

Declarations

Competing interests

The authors declare no conflicts of interest.

Author details

¹Molecular AI, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden. ²Medicinal Chemistry, Research and Early Development, Cardiovascular, Renal and Metabolism (CVRM), BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden. ³Structure & Biophysics, Discovery Sciences, R&D, AstraZeneca, Cambridge, UK. ⁴Medicinal Chemistry, Research and Early Development, Respiratory and Immunology (R&I), BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden. ⁵Medicinal Chemistry, Research and Early Development, Oncology R&D, AstraZeneca, Cambridge, UK. ⁶Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden.

Received: 5 August 2021 Accepted: 29 October 2021

Published online: 17 November 2021

References

- Schneider P, Walters WP, Plowright AT, Sieroka N, Listgarten J, Goodnow RA, Fisher J, Jansen JM, Duca JS, Rush TS, Zentgraf M, Hill JE, Krutsholow E, Kohler M, Blaney J, Funatsu K, Luebkemann C, Schneider G (2020) Rethinking Drug design in the artificial intelligence era. *Nat Rev Drug Discov* 19(5):353–364. <https://doi.org/10.1038/s41573-019-0050-3>
- Jiménez-Luna J, Grisoni F, Weskamp N, Schneider G (2021) Artificial intelligence in drug discovery: recent advances and future perspectives. *Expert Opin Drug Discov*. <https://doi.org/10.1080/17460441.2021.1909567>
- Polishchuk PG, Madzhidov TI, Varnek A (2013) Estimation of the size of drug-like chemical space based on GDB-17 data. *J Comput Aided Mol Des* 27(8):675–679. <https://doi.org/10.1007/s10822-013-9672-4>
- REINVENT 2.0: An AI tool for de novo drug design. *J Chem Inf Model*. <https://doi.org/10.1021/acs.jcim.0c00915>. Accessed 14 Jun 2021.

5. Lim J, Ryu S, Kim JW, Kim WY (2018) Molecular generative model based on conditional variational autoencoder for de novo molecular design. *J Cheminformatics* 10(1):31. <https://doi.org/10.1186/s13321-018-0286-7>
6. Maziarka Ł, Pocha A, Kaczmarczyk J, Rataj K, Danel T, Warchol M (2020) Mol-CycleGAN: a generative model for molecular optimization. *J Cheminformatics* 12(1):2. <https://doi.org/10.1186/s13321-019-0404-1>
7. Mercado R, Rastemo T, Lindelöf E, Klambauer G, Engkvist O, Chen H, Jan-nik Bjerrum E (2021) Graph networks for molecular design. *Mach Learn Sci Technol* 2(2):025023. <https://doi.org/10.1088/2632-2153/abc9f1>
8. Popova M, Isayev O, Tropsha A (2018) Deep reinforcement learning for de novo drug design. *Sci Adv* 4(7):eaap7885. <https://doi.org/10.1126/sciadv.aap7885>
9. Pereira T, Abbasi M, Ribeiro B, Arrais JP (2021) Diversity oriented deep reinforcement learning for targeted molecule generation. *J Cheminformatics* 13(1):21. <https://doi.org/10.1186/s13321-021-00498-z>
10. Kell DB, Samanta S, Swainston N (2020) Deep learning and generative methods in cheminformatics and chemical biology: navigating small molecule space intelligently. *Biochem J* 477(23):4559–4580. <https://doi.org/10.1042/BCJ20200781>
11. Kausar S, Falcão AO (2018) An automated framework for QSAR model building. *J Cheminformatics* 10(1):1. <https://doi.org/10.1186/s13321-017-0256-5>
12. Zhao L, Wang W, Sedykh A, Zhu H (2017) Experimental errors in QSAR modeling sets: what we can do and what we cannot do. *ACS Omega* 2(6):2805–2812. <https://doi.org/10.1021/acsomega.7b00274>
13. Pinzi L, Rastelli G (2019) Molecular docking: shifting paradigms in drug discovery. *Int J Mol Sci* 20(18):4331. <https://doi.org/10.3390/ijms20184331>
14. El Kerdawy AM, Osman AA, Zaater MA (2019) Receptor-based pharmacophore modeling, virtual screening, and molecular docking studies for the discovery of novel GSK-3 β inhibitors. *J Mol Model* 25(6):171. <https://doi.org/10.1007/s00894-019-4032-5>
15. Zhao W, Xiong M, Yuan X, Li M, Sun H, Xu Y (2020) In silico screening-based discovery of novel inhibitors of human cyclic GMP–AMP synthase: a cross-validation study of molecular docking and experimental testing. *J Chem Inf Model* 60(6):3265–3276. <https://doi.org/10.1021/acs.jcim.0c00171>
16. Lyu J, Wang S, Balius TE, Singh I, Levit A, Moroz YS, O'Meara MJ, Che T, Algaa E, Tolmachova K, Tolmachev AA, Shoichet BK, Roth BL, Irwin JJ (2019) Ultra-large library docking for discovering new chemotypes. *Nature* 566(7743):224–229. <https://doi.org/10.1038/s41586-019-0917-9>
17. Welsch ME, Snyder SA, Stockwell BR (2010) Privileged scaffolds for library design and drug discovery. *Curr Opin Chem Biol* 14(3):347–361. <https://doi.org/10.1016/j.cbpa.2010.02.018>
18. Thomas M, Smith RT, O'Boyle NM, de Graaf C, Bender A (2021) Comparison of structure- and ligand-based scoring functions for deep generative models: A GPCR case study. *J Cheminformatics* 13(1):39. <https://doi.org/10.1186/s13321-021-00516-0>
19. Ma B, Terayama K, Matsumoto S, Isaka Y, Sasakura Y, Iwata H, Araki M, Okuno Y (2021) Structure-based de novo molecular generator combined with artificial intelligence and docking simulations. *J Chem Inf Model*. <https://doi.org/10.26434/chemrxiv.14371967.v1>
20. Bai Q, Tan S, Xu T, Liu H, Huang J, Yao X (2021) MolAICal: a soft tool for 3D drug design of protein targets by artificial intelligence and classical algorithm. *Brief Bioinform*. <https://doi.org/10.1093/bib/bbaa161>
21. Jeon W, Kim D (2020) Autonomous molecule generation using reinforcement learning and docking to develop potential novel inhibitors. *Sci Rep* 10(1):22104. <https://doi.org/10.1038/s41598-020-78537-2>
22. Li X, Li Y, Cheng T, Liu Z, Wang R (2010) Evaluation of the performance of four molecular docking programs on a diverse set of protein-ligand complexes. *J Comput Chem* 31(11):2109–2125. <https://doi.org/10.1002/jcc.21498>
23. Wang Z, Sun H, Yao X, Li D, Xu L, Li Y, Tian S, Hou T (2016) Comprehensive evaluation of ten docking programs on a diverse set of protein-ligand complexes: the prediction accuracy of sampling power and scoring power. *Phys Chem Chem Phys* 18(18):12964–12975. <https://doi.org/10.1039/C6CP01555G>
24. Du X, Li Y, Xia Y-L, Ai S-M, Liang J, Sang P, Ji X-L, Liu S-Q (2016) Insights into protein-ligand interactions: mechanisms, models, and methods. *Int J Mol Sci*. <https://doi.org/10.3390/ijms17020144>
25. Chodera JD, Mobley DL (2013) Entropy-enthalpy compensation: role and ramifications in biomolecular ligand recognition and design. *Annu Rev Biophys* 42:121–142. <https://doi.org/10.1146/annurev-biophys-083012-130318>
26. Bauer MR, Ibrahim TM, Vogel SM, Boeckler FM (2013) Evaluation and optimization of virtual screening workflows with DEKOIS 2.0—a public library of challenging docking benchmark sets. *J Chem Inf Model* 53(6):1447–1462. <https://doi.org/10.1021/ci400115b>
27. Sadowski J, Gasteiger J, Klebe G (1994) Comparison of automatic three-dimensional model builders using 639 X-Ray structures. *J Chem Inf Comput Sci* 34(4):1000–1008. <https://doi.org/10.1021/ci00020a039>
28. Schwab CH (2010) Conformations and 3D pharmacophore searching. *Drug Discov Today Technol* 7(4):e245–e253. <https://doi.org/10.1016/j.ddtec.2010.10.003>
29. 3D Structure Generator CORINA Classic, Molecular Networks GmbH, Nuremberg, Germany, www.Mn-Am.Com
30. Schrödinger Release 2019-4: LigPrep, Schrödinger, LLC, New York, NY, 2019.
31. Hawkins PCD, Skillman AG, Warren GL, Ellingson BA, Stahl MT (2010) Conformer generation with OMEGA: algorithm and validation using high quality structures from the protein databank and cambridge structural database. *J Chem Inf Model* 50(4):572–584. <https://doi.org/10.1021/ci100031x>
32. OMEGA 4.1.0.2: OpenEye Scientific Software, Santa Fe, NM. <http://www.eyesopen.com>
33. Trott O, Olson AJ (2010) AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization and multithreading. *J Comput Chem* 31(2):455–461. <https://doi.org/10.1002/jcc.21334>
34. Friesner RA, Banks JL, Murphy RB, Halgren TA, Klicic JJ, Mainz DT, Repasky MP, Knoll EH, Shelley M, Perry JK, Shaw DE, Francis P, Shenkin PS (2004) Glide: a new approach for rapid, accurate docking and scoring. 1. Method and assessment of docking accuracy. *J Med Chem* 47(7):1739–1749. <https://doi.org/10.1021/jm0306430>
35. Halgren TA, Murphy RB, Friesner RA, Beard HS, Frye LL, Pollard WT, Banks JL (2004) Glide: a new approach for rapid, accurate docking and scoring. 2. Enrichment factors in database screening. *J Med Chem* 47(7):1750–1759. <https://doi.org/10.1021/jm030644s>
36. Friesner RA, Murphy RB, Repasky MP, Frye LL, Greenwood JR, Halgren TA, Sanschagrin PC, Mainz DT (2006) Extra precision glide: docking and scoring incorporating a model of hydrophobic enclosure for protein-ligand complexes. *J Med Chem* 49(21):6177–6196. <https://doi.org/10.1021/jm051256o>
37. Schrödinger Release 2019-4: Glide, Schrödinger, LLC, New York, NY, 2019.
38. Jones G, Willett P, Glen RC, Leach AR, Taylor R (1997) Development and validation of a genetic algorithm for flexible docking 11 edited by F. E. Cohen. *J Mol Biol* 267(3):727–748. <https://doi.org/10.1006/jmbi.1996.0897>
39. OEDOCKING 4.0.0.2: OpenEye Scientific Software, Santa Fe, NM. <http://www.eyesopen.com>
40. McGann M (2012) FRED and HYBRID docking performance on standardized datasets. *J Comput Aided Mol Des* 26(8):897–906. <https://doi.org/10.1007/s10822-012-9584-8>
41. Ruiz-Carmona S, Alvarez-Garcia D, Foloppe N, Garmendia-Doval AB, Juhos S, Schmidtke P, Barril X, Hubbard RE, Morley SD (2014) RDock: a fast, versatile and open source program for docking ligands to proteins and nucleic acids. *PLOS Comput Biol* 10(4):e1003571. <https://doi.org/10.1371/journal.pcbi.1003571>
42. Anderson AC (2012) Structure-based functional design of drugs: from target to lead compound. *Methods Mol Biol Clifton NJ* 823:359–366. https://doi.org/10.1007/978-1-60327-216-2_23
43. Kelley BP, Brown SP, Warren GL, Muchmore SW (2015) POSIT: flexible shape-guided docking for pose prediction. *J Chem Inf Model* 55(8):1771–1780. <https://doi.org/10.1021/acs.jcim.5b00142>
44. Clark RD, Webster-Clark DJ (2008) Managing bias in ROC curves. *J Comput Aided Mol Des* 22(3):141–146. <https://doi.org/10.1007/s10822-008-9181-z>

45. Zwillinger, D. and Kokoska, S. (2000). CRC standard probability and statistics tables and formulae. Chapman & Hall: New York; 2000. Section 14.7.
46. Kendall MG (1945) The treatment of ties in ranking problems. *Biometrika* 33(3):239–251. <https://doi.org/10.1093/biomet/33.3.239>
47. Blaschke T, Engkvist O, Bajorath J, Chen H (2020) Memory-assisted reinforcement learning for diverse molecular de novo design. *J Cheminformatics* 12(1):68. <https://doi.org/10.1186/s13321-020-00473-0>
48. Arús-Pous J, Johansson SV, Prykhodko O, Bjerrum EJ, Tyrchan C, Reymond J-L, Chen H, Engkvist O (2019) Randomized SMILES strings improve the quality of molecular generative models. *J Cheminformatics* 11(1):71. <https://doi.org/10.1186/s13321-019-0393-0>
49. Cho K, van Merriënboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, Bengio Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *ArXiv14061078 Cs Stat* 2014.
50. Gaulton A, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, Mutowo P, Atkinson F, Bellis LJ, Cibrián-Uhalte E, Davies M, Dedman N, Karlsson A, Magariños MP, Overington JP, Papadatos G, Smit I, Leach AR (2017) The ChEMBL database in 2017. *Nucleic Acids Res* 45(D1):D945–D954. <https://doi.org/10.1093/nar/gkw1074>
51. Korshunova M, Huang N, Capuzzi S, Radchenko DS, Savych O, Moroz YS, Wells C, Willson TM, Tropsha A, Isayev O. A bag of tricks for automated de novo design of molecules with the desired properties: application to EGFR inhibitor discovery. 2021. Doi: <https://doi.org/10.26434/chemrxiv.14045072.v1>.
52. Arnott JA, Planey SL (2012) The influence of lipophilicity in drug discovery and design. *Expert Opin Drug Discov* 7(10):863–875. <https://doi.org/10.1517/17460441.2012.714363>
53. Baell JB, Holloway GA (2010) New substructure filters for removal of pan assay interference compounds (PAINS) from screening libraries and for their exclusion in bioassays. *J Med Chem* 53(7):2719–2740. <https://doi.org/10.1021/jm901137j>
54. Kenny PW, Sadowski J. Structure modification in chemical databases. In: *Chemoinformatics in drug discovery*. John Wiley & Sons, Ltd; 2005. p. 271–285. Doi: <https://doi.org/10.1002/3527603743.ch11>.
55. Daylight Theory: SMIRKS—a reaction transform language <https://www.daylight.com/dayhtml/doc/theory/theory.smirks.html>. Accessed 14 Jun 2021.
56. Lipinski CA, Lombardo F, Dominy BW, Feeney PJ (1997) Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv Drug Deliv Rev* 23(1):3–25. [https://doi.org/10.1016/S0169-409X\(96\)00423-1](https://doi.org/10.1016/S0169-409X(96)00423-1)
57. Bickerton GR, Paolini GV, Besnard J, Muresan S, Hopkins AL (2012) Quantifying the chemical beauty of drugs. *Nat Chem* 4(2):90–98. <https://doi.org/10.1038/nchem.1243>
58. Wu Y, Choma N, Chen A, Cashman M, Prates ÉT, Shah M, Vergara VGM, Clyde A, Brettin TS, de Jong WA, Kumar N, Head MS, Stevens RL, Nugent P, Jacobson DA, Brown JB. Spatial graph attention and curiosity-driven policy for antiviral drug discovery. *ArXiv210602190 Cs Q-Bio*. 2021.
59. Rolnick D, Ahuja A, Schwarz J, Lillicrap TP, Wayne G. Experience replay for continual learning. *ArXiv181111682 Cs Stat*. 2019.
60. Kurumbail RG, Stevens AM, Gierse JK, McDonald JJ, Stegeman RA, Pak JY, Gildehaug D, Iyashiro JM, Penning TD, Seibert K, Isakson PC, Stallings WC (1996) Structural basis for selective inhibition of cyclooxygenase-2 by anti-inflammatory agents. *Nature* 384(6610):644–648. <https://doi.org/10.1038/384644a0>
61. Deb PK, Mailabaram RP, Al-Jaidi B, Saadh M (2017) Molecular basis of binding interactions of NSAIDs and computer-aided drug design approaches in the pursuit of the development of cyclooxygenase-2 (COX-2) selective inhibitors. *Nonsteroidal Anti Inflamm Drugs*. 23(6):101–121. <https://doi.org/10.5772/intechopen.68318>
62. Ferreira L, Hitchcock DB (2009) A comparison of hierarchical methods for clustering functional data. *Commun Stat Simul Comput* 38(9):1925–1949. <https://doi.org/10.1080/03610910903168603>
63. Argiriadi MA, Ericsson AM, Harris CM, Banach DL, Borhani DW, Calderwood DJ, Demers MD, DiMauro J, Dixon RW, Hardman J, Kwak S, Li B, Mankovich JA, Marcotte D, Mullen KD, Ni B, Pietras M, Sadhukhan R, Sousa S, Tomlinson MJ, Wang L, Xiang T, Talanian RV (2010) 2,4-diaminopyrimidine MK2 inhibitors. Part I: observation of an unexpected inhibitor binding mode. *Bioorg Med Chem Lett* 20(1):330–333. <https://doi.org/10.1016/j.bmcl.2009.10.102>
64. Bressi JC, Jennings AJ, Skene R, Wu Y, Melkus R, Jong RD, O'Connell S, Grimshaw CE, Navre M, Gangloff AR (2010) Exploration of the HDAC2 foot pocket: synthesis and SAR of substituted N-(2-aminophenyl)benzamides. *Bioorg Med Chem Lett* 20(10):3142–3145. <https://doi.org/10.1016/j.bmcl.2010.03.091>
65. Nahoum V, Perez E, Germain P, Rodriguez-Barrios F, Manzo F, Kammerer S, Lemaire G, Hirsch O, Royer CA, Gronemeyer H, de Lera AR, Bourguet W (2007) Modulators of the structural dynamics of the retinoid X receptor to reveal receptor function. *Proc Natl Acad Sci* 104(44):17323–17328. <https://doi.org/10.1073/pnas.0705356104>
66. Wang G-H, Jiang F-Q, Duan Y-H, Zeng Z-P, Chen F, Dai Y, Chen J-B, Liu J-X, Liu J, Zhou H, Chen H-F, Zeng J-Z, Su Y, Yao X-S, Zhang X-K (2013) Targeting truncated retinoid X receptor- α by CF31 induces TNF- α -dependent apoptosis. *Cancer Res* 73(1):307–318. <https://doi.org/10.1158/0008-5472.CAN-12-2038>
67. Almog R, Waddling CA, Maley F, Maley GF, Roey PV (2001) Crystal structure of a deletion mutant of human thymidylate synthase Δ (7–29) and its ternary complex with tomudex and DUMP. *Protein Sci* 10(5):988–996. <https://doi.org/10.1110/ps.47601>
68. McInnes L, Healy J, Melville J. UMAP: uniform manifold approximation and projection for dimension reduction. *ArXiv180203426 Cs Stat* 2020.
69. Rogers D, Hahn M. Extended-connectivity fingerprints. <https://doi.org/10.1021/ci100050t>. Accessed 14 Jun 2021.
70. Durrant JD, McCammon JA (2011) Molecular dynamics simulations and drug discovery. *BMC Biol* 9(1):71. <https://doi.org/10.1186/1741-7007-9-71>
71. Kurkinen ST, Lätti S, Pentikäinen OT, Postila PA (2019) Getting docking into shape using negative image-based rescoring. *J Chem Inf Model* 59(8):3584–3599. <https://doi.org/10.1021/acs.jcim.9b00383>
72. Gu S, Smith MS, Yang Y, Irwin JJ, Shoichet BK (2021) Ligand Strain energy in large library docking. *bioRxiv*. <https://doi.org/10.1101/2021.04.06.438722>
73. Eberhardt J, Santos-Martins D, Tillack A, Forli S. AutoDock Vina 1.2.0: new docking methods, expanded force field, and python bindings. 2021. <https://doi.org/10.26434/chemrxiv.14774223.v1>.
74. Kudo N, Ito A, Arata M, Nakata A, Yoshida M (2018) Identification of a novel small molecule that inhibits deacetylase but not deacetylase reaction catalysed by SIRT2. *Philos Trans R Soc B Biol Sci*. <https://doi.org/10.1098/rstb.2017.0070>
75. Madhavi Sastry G, Adzhigirey M, Day T, Annabhimoju R, Sherman W (2013) Protein and ligand preparation: parameters, protocols, and influence on virtual screening enrichments. *J Comput Aided Mol Des* 27(3):221–234. <https://doi.org/10.1007/s10822-013-9644-8>
76. Schrödinger release 2019-4: protein preparation wizard; Epik, Schrödinger, LLC, New York, NY, 2016; Impact, Schrödinger, LLC, New York, NY, 2016; Prime, Schrödinger, LLC, New York, NY, 2019.
77. Roos K, Wu C, Damm W, Reboul M, Stevenson JM, Lu C, Dahlgren MK, Mondal S, Chen W, Wang L, Abel R, Friesner RA, Harder ED (2019) OPLS3e: extending force field coverage for drug-like small molecules. *J Chem Theory Comput* 15(3):1863–1874. <https://doi.org/10.1021/acs.jctc.8b01026>
78. Eastman P, Friedrichs MS, Chodera JD, Radmer RJ, Bruns CM, Ku JP, Beauchamp KA, Lane TJ, Wang L-P, Shukla D, Tye T, Houston M, Stich T, Klein C, Shirts MR, Pande VS (2013) OpenMM 4: a reusable, extensible, hardware independent library for high performance molecular simulation. *J Chem Theory Comput* 9(1):461–469. <https://doi.org/10.1021/ct300857j>
79. O'Boyle NM, Banck M, James CA, Morley C, Vandermeersch T, Hutchison GR (2011) Open babel: an open chemical toolbox. *J Cheminformatics* 3(1):33. <https://doi.org/10.1186/1758-2946-3-33>
80. Schrödinger Release 2019-4: Maestro, Schrödinger, LLC, New York, NY, 2019.

81. Groom CR, Bruno IJ, Lightfoot MP, Ward SC (2016) The Cambridge structural database. *Acta Crystallogr Sect B Struct Sci Cryst Eng Mater.* 72(2):171–179. <https://doi.org/10.1107/S2052520616003954>
82. Rappe AK, Casewit CJ, Colwell KS, Goddard WA III, Skiff WM. UFF, a full periodic table force field for molecular mechanics and molecular dynamics simulations. <https://doi.org/10.1021/ja00051a040>. Accessed 14 Jun 2021.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions