# Link-INVENT: Generative Linker Design with Reinforcement

Jeff Guo<sup>1,⊥</sup>, Franziska Knuth<sup>1,2,⊥</sup>, Christian Margreitter<sup>1</sup>, Jon Paul Janet<sup>3</sup>, Kostas Papadopoulos<sup>1</sup>, Ola Engkvist<sup>1,4</sup>, Atanas Patronov<sup>1,\*</sup>

<sup>1</sup> Molecular AI, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden

<sup>2</sup> Department of Physics, Norwegian University of Science and Technology, Trondheim, Norway

<sup>3</sup> Medicinal Chemistry, Research and Early Development, Cardiovascular, Renal and Metabolism (CVRM),

BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden

<sup>4</sup> Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg 41756, Sweden

\* Corresponding author: patronov@gmail.com

⊥ These authors have contributed equally

## Abstract

In this work, we present Link-INVENT as an extension to the existing *de novo* molecular design platform REINVENT. We provide illustrative examples on how Link-INVENT can be applied on fragment linking, scaffold hopping, and PROTACs design case studies where the desirable molecules should satisfy a combination of different criteria. With the help of Reinforcement Learning, the agent used by Link-INVENT learns to generate favourable linkers connecting molecular subunits that satisfy diverse objectives, facilitating practical application of the model for real-world drug discovery projects. We also introduce a range of linker-specific objectives in the scoring function of REINVENT. The code is freely available at <a href="https://github.com/MolecularAI/Reinvent">https://github.com/MolecularAI/Reinvent</a>.

### Introduction

Deep learning (DL) offers potential to accelerate drug design by efficiently traversing chemical space, defined as the set of all possible biologically relevant molecules, and estimated to be on the order of  $10^{23}$  to  $10^{60}$ .<sup>1–3</sup> The task is challenging as candidate drug molecules must satisfy a multi-parameter optimization (MPO) objective where parallel optimization of the individual objectives can be difficult. Thus, a brute-force molecular search approach is often infeasible. DL-augmented molecular design combined with computational oracles to approximate physico-chemical properties has enabled candidate drug molecules to be designed in an accelerated manner.<sup>4</sup> DL-based approaches to the drug design task include deep generative models with reinforcement learning (RL)<sup>5–10</sup>, learning a molecular latent space<sup>11</sup>, and genetic algorithms<sup>12–14</sup>, which generate molecular ideas as Simplified Molecular-Input Line-Entry System (SMILES) strings<sup>15</sup> or molecular graphs<sup>9,10</sup>.

Recently, the application of DL-based methods to join two molecular subunits via a chemical linker has gained considerable interest.<sup>16–19</sup> Generating suitable linkers is important for fragment-based drug discovery (FBDD)<sup>20,21</sup> and scaffold hopping<sup>22</sup>, and fundamental for the design of proteolysis targeting chimeras (PROTACs)<sup>23–25</sup>. The former two techniques are avenues to discover and optimize novel small molecule drugs, while the latter is a relatively new therapeutic modality able to achieve targeted protein degradation. Therefore, linker design represents a relevant problem in drug discovery.

FBDD is an alternative to traditional high-throughput screening (HTS) and virtual screening (VS) which screens 'Lipinski compliant' small molecules. In contrast, FBDD screens 'fragments', typically with a molecular weight (MW) under 260 Da. Although 'fragment' hits typically exhibit weaker binding affinities than small molecules, they often form polar interactions with the receptor and possess

favourable lipophilicity, limiting entropically-driven binding.<sup>20,21,26</sup> Thus, 'fragments' can be an advantageous starting point for drug design and techniques to optimize their potency and physicochemical properties include fragment growing and fragment linking.<sup>20,21,27,28</sup> The latter is of particular interest as proper linking of two 'fragments' such that the linked molecule does not perturb the constituents' interactions can lead to significant potency gain. This is attributed to favourable entropic effects and known as 'super-additivity'. In practice, fragment linking is challenging and 'super-additivity' is rarely achieved, owing to incompatible linkers disrupting the fragments' binding poses.<sup>27,28</sup> Thus, improvements in linker design are critical to unlock the full potential of FBDD.

Scaffold hopping refers to modifying the core structure of a molecule to improve physicochemical properties while retaining potency.<sup>22</sup> The task can be formulated as a linker design problem if the scaffold itself is defined as the linker between two molecular subunits. Scaffold hopping is challenging as retaining potency requires 3D structural awareness of the interactions formed between the molecule and its receptor. Similar to fragment linking, improvements in linker design can enhance the ability to generate novel scaffold ideas.

PROTACs are heterobifunctional molecules in which a linker joins a ligand binding to a protein of interest (POI), conferring specificity, and an E3 ubiquitin ligase. The formation of the ternary complex leads to subsequent ubiquitination, achieving POI degradation and thus, targeted knockdown.<sup>23–25</sup> While the unique mechanism of action provides promising therapeutic applicability beyond traditional small molecules, PROTACs design is challenging. PROTACs are comparably large molecules, typically existing beyond 'Lipinski's rules' and thereby posing a design challenge since experience is limited.<sup>29–31</sup> Moreover, linker design is challenging due to the relatively high conformational flexibility present in longer linkers and has mostly deferred to empirical structure-activity relationship (SAR) studies, often necessitating numerous iterations of design-make-test-analyze (DMTA) cycles.<sup>32</sup> Therefore, there is a need for improved linker design to improve overall PROTACs design.

Previously developed computational tools for linker design involve searching a database, making the generalizability of proposed linkers inherently limited.<sup>33–36</sup> While success has been demonstrated when using these methods combined with filtering steps, one would ideally want to generalize the task such that plausible linker ideas can be proposed given any molecular subunits.<sup>33–36</sup> Recently, DL-based linker design models have been proposed that circumvent database searches.<sup>16–19</sup> DeLinker is a graphbased model proposed by Imrie et al. which explicitly incorporates 3D information via the distance and angle between the molecular subunits to augment the feature vector.<sup>16</sup> SyntaLinker is a conditional transformer model proposed by Yang et al. which treats linker generation as a natural language processing (NLP) task using SMILES.<sup>15,17</sup> SyntaLinker was further extended by Hu et al. to perform kinase scaffold hopping after focusing the model via transfer learning.<sup>18</sup> Moreover, Langevin et al. proposed the Scaffold Constrained Molecular Generation (SAMOA) algorithm based on recurrent neural networks (RNNs) where one of the capabilities of the model is linker generation.<sup>19</sup> However, while these models are capable of generating linker ideas, a major drawback is the limited support to optimize explicitly for desired physico-chemical properties. The current models only allow users to control for the desired linker length<sup>16–19</sup> and a select number of physico-chemical properties, e.g., number of hydrogen-bond donors (HBD).<sup>17</sup> To encourage wide adoption of DL-based linker design, increased flexibility to define tailored MPO objectives and better generalizability is needed.

In this work, we present Link-INVENT as an extension to the existing *de novo* design platform REINVENT.<sup>6</sup> The suggested algorithm shares some similarities with the SAMOA algorithm as proposed by Langevin et al.<sup>19</sup> in that the code builds upon REINVENT's existing codebase and uses policy-based RL for MPO.<sup>6</sup> However, our algorithm has three crucial differences compared to earlier work. Firstly, the prior trained by Langevin et al. is based on ChEMBL compounds and follows the protocol as reported for REINVENT, which was purposed to sample small molecules as SMILES.<sup>6,19,37</sup> Consequently, in their linker generation solution, linkers are sampled when the "\*" token (the model's internal representation of characters in a SMILES string), denoting the attachment point, is reached, and based on the conditional

probabilities of the SMILES sequence so far. The limitation is that linkers should be generated in the context of both molecular subunits. In the extreme case, the SAMOA algorithm may struggle to generate plausible linkers if the SMILES sequence was "CC\*C..." where the length of the SMILES on the right side of the "\*" token is greater than the left side, as the conditional probabilities for linker generation would only be based on the sequence so far, i.e., "CC". In contrast, Link-INVENT is trained based on the conditional probabilities of observing a linker given both molecular subunits, similar to the SyntaLinker model reported by Yang et al.<sup>17</sup> Secondly, the data preparation to train the Link-INVENT prior was based on reaction-splicing of the ChEMBL compounds similar to the Lib-INVENT library design model we reported previously.<sup>37,38</sup> Our training set contains linkers that join molecular subunits ranging from a few atoms in size to larger moieties with rings. As a result, a single Link-INVENT prior is suited for diverse linker generation tasks. Lastly, Link-INVENT was built on the latest version of REINVENT and supports an extensive selection of physico-chemical properties that can be optimized through RL. Moreover, we have implemented additional linker specific properties that can be optimized (in the form of additional scoring function components), ranging from physico-chemical properties to flexibility and rigidity, allowing one to explicitly optimize linker properties. We demonstrate the use of Link-INVENT on fragment linking, scaffold hopping, and PROTACs design case studies. Through RL, the Link-INVENT agent learns to generate favourable linkers connecting molecular subunits that satisfy diverse MPO objectives, facilitating practical application of the model for real-world drug discovery projects. The code is freely available at <u>https://github.com/MolecularAI/Reinvent</u>.



**Fig. 1** Link-INVENT training and inference overview. ChEMBL data was processed to generate the training data for Link-INVENT which features an encoder-decoder architecture of recurrent neural networks (RNNs). Link-INVENT takes as input a pair of warheads and generates output linkers (highlighted green), yielding connected molecules. One crucial capability of Link-INVENT is the flexible Scoring Function where diverse multi-parameter optimization (MPO) objectives can be specified. Through reinforcement learning (RL), the agent's policy is iteratively updated to satisfy the target MPO objective and thus generating desirable connected molecules.

#### Methods

**Model Overview**. Link-INVENT takes as input a pair of warheads, i.e., two molecular subunits with exit vectors defined, generates a linker, and returns the linked molecule in the SMILES format (Fig. 1).<sup>15</sup> The model is adapted from Lib-INVENT, our previously reported generative model for library design by Fialková et al. which in turn is based on work by Arús-Pous et al.<sup>38,39</sup> Specifically, Link-INVENT features an encoder-decoder architecture consisting of identical RNNs with embedding size 256 and three hidden layers of 512 long short-term memory cells (LSTM).<sup>40</sup>

- 1. Initial Filtering: Filter the raw ChEMBL data to keep 'drug-like' compounds only (see Supporting Information for details). Lenient filtering criteria was applied such that the training data is effective for PROTACs applications where the warheads can be larger in size compared to traditional 'fragments'.<sup>30,31</sup>
- Reaction-based Slicing: Slice the filtered ChEMBL compounds following the protocol from our Lib-INVENT work using reaction SMIRKS.<sup>38</sup> The result is a dataset of tuples with the structure: (linker, warheads pair, full molecule)
- 3. **Sliced Data Filtering:** Filter the tuples to remove unrealistic data points, e.g., linkers with molecular weight greater than 500 Da.
- 4. Generate Training and Validation Sets: A validation set containing 287 Bemis-Murcko scaffolds was held out.<sup>41</sup>
- 5. SMILES Randomization: Data augmentation for the training and validation sets was performed via SMILES randomization. At each training epoch, the model is provided with datasets comprised of the same sliced tuples (linker, warheads pair, full molecule) but with a different SMILES representation. The purpose was to improve chemical space generalizability of the generative model as shown by Arús-Pous et al.<sup>42</sup>

For full details of the data preparation, see the Supporting Information.

**Model Training.** First, a vocabulary was generated that maps characters present in the training set (and validation set) SMILES to tokens (see the Supporting Information for token details). Querying Link-INVENT requires tokenization of the input warheads by the encoder and the output linker tokens from the decoder are then transformed into their SMILES equivalent.<sup>15</sup> The initial generative model, denoted the prior, was trained by maximizing the likelihood of generating a linker *conditioned* on the input pair of warheads. Teacher forcing was used such that the ground-truth labels were fed back to the model at each token sampling step to improve training stability.<sup>43</sup> The trained prior is a generative model that has learned the SMILES syntax and is thus capable of generating syntactically valid linkers given a pair of input warheads.

**Model Inference and Multi-Parameter Optimization.** Following REINVENT's protocol, the agent is initialized to have the same parameters as the prior and serves two purposes<sup>6,44</sup>: ensuring the agent is also capable of generating syntactically valid linker SMILES and anchoring the sampled linkers to relevant chemical space as defined by the training data derived from ChEMBL.<sup>15,37</sup> Subsequently, the agent is tasked to generate linkers that satisfy MPO objectives, given by the Scoring Function. The Scoring Function specifies all components to be optimized and is formulated as a weighted geometric mean in this work:

$$S(x) = \left(\prod_{i=1}^{n} C_{i}(x)^{w_{i}}\right)^{\frac{1}{\sum_{i=1}^{n} w_{i}}}$$
(1)

where x is a sampled linked molecule (or linker), n is the number of components in the Scoring Function S,  $C_i$  is the score for the *i*th component, and  $w_i$  is the weight for the *i*th component. The weighting allows the user to control the relative importance of certain components, where a greater weighting will result in a greater contribution of that component to the Scoring Function. One crucial advantage of

Link-INVENT is the flexibility in components that can be specified in the Scoring Function, ranging from physico-chemical properties, structural features, predictive models, and physics-based approximations of binding energy (see Supporting Information for a list of all linker components implemented in Link-INVENT). While the agent is initially identical to the prior, its policy is updated via RL such that the agent is steered to generate molecules that increasingly satisfy the desired MPO objective.<sup>6</sup> Correspondingly, Link-INVENT inference is as follows:

- 1. **Agent Sampling:** Generate batch size (128 in this work) number of linkers *conditioned* on an input pair of warheads. Thus, 128 linked molecules were generated at each epoch.
- 2. Assess Linked Molecules' Desirability: Combine the warheads and linkers to form the linked molecules and compute their *desirability* based on satisfaction of the Scoring Function.
- 3. Update Agent Policy: Compute the loss and update the agent's policy to steer sampling towards favourable linkers. The specific loss function used in Link-INVENT was previously introduced by Fialková et al. in our Lib-INVENT work and defined as the difference between the augmented and posterior likelihoods (DAP).<sup>38</sup> Correspondingly, the same loss function was used in this work and is constructed by first defining the augmented log likelihood:

$$\log \pi_{Augmented} = \log \pi_{Prior} + \sigma S(x) \tag{2}$$

where  $\pi$  denotes a policy, i.e., probabilities of sampling tokens conditioned on observing a token sequence, S(x) is the Scoring Function whose value is computed based on a generated SMILES sequence, x, and  $\sigma$  is a scalar factor. From equation 2, the augmented log likelihood is comprised of the prior log likelihood adjusted by the *desirability* of a sampled SMILES sequence based on the Scoring Function. The loss is then defined as follows:

$$J(\theta) = \left(\log \pi_{Augmented} - \log \pi_{Agent}\right)^2$$
(3)

through RL, the agent policy is updated at each epoch to minimize the loss,  $J(\theta)$ .

Steps 1-3 are repeated until the permitted number of epochs has elapsed. All favourable linkers (and the corresponding full molecules) that achieve a total score (computed by aggregating the scores achieved on each composite objective defined in the Scoring Function) exceeding a user-defined threshold (typically 0.4) are outputted. In this work, the threshold was set to 0 to store all molecules generated. The purpose of this was to compare the profiles of molecules generated towards the beginning of the experiment and how RL gradually guides the generation of favourable molecules.

**Balancing Chemical Space Exploration and Exploitation.** Link-INVENT offers full control over chemical space exploration and exploitation by leveraging Diversity Filters (DFs) as implemented in REINVENT.<sup>6,45</sup> *Buckets* can be defined with limited size that keep track of unique scaffolds. Agent sampling of molecules containing identical scaffolds populate the same *bucket*. If a *bucket* is full, further sampling of the scaffold will cause the agent to receive a score of 0, regardless of the corresponding molecule's *desirability*, as assessed by the Scoring Function. This mechanism encourages agent exploration of diverse minima. For more details regarding DFs, see the work by Blaschke et al.<sup>6,45</sup> The specific DF used in this work features *buckets* of size 25 and penalizes repeated sampling of Bemis-Murcko scaffolds.<sup>41</sup> We note that as Link-INVENT generates linkers *conditioned* on a pair of input warheads, the warheads

themselves are held constant. Thus, the DF effectively penalizes repeated sampling of the Bemis-Murcko

scaffolds of the linker themselves.41

# 

**Fig. 2** Link-INVENT Scoring Function selected linker specific properties. Attachment points are denoted by red asterisks. **a.** "Linker effective length", "linker maximum graph length", and "linker length ratio" provide direct control over the distance between attachment atoms in a linker and the degree of branching permitted. The maximum graph length bond traversal is highlighted green in the bottom example. **b.** "Linker ratio of rotatable bonds" provides control over the flexibility of proposed linkers.

**Scoring Function: Controlling Linker Properties.** In addition to the previously supported Scoring Function properties in REINVENT that operate on the full molecule, Link-INVENT offers control over the linker itself (Fig. 2). One can control the linker length and branching (Fig. 2a) via the following properties:

- 1. Linker Effective Length: the number of bonds between the attachment atoms
- 2. Linker Maximum Graph Length: the number of bonds encompassed in the longest molecular graph traversal path
- Linker Length Ratio: the ratio of the "linker effective length" over the "linker maximum graph length"

Moreover, one can control linker flexibility through the "linker ratio of rotatable bonds" component which is defined as the number of rotatable bonds (as calculated by RDKit<sup>46</sup>) over the total number of

bonds (Fig. 2b). We note that this treatment of flexibility is not the only valid definition and inherent limitations exist such as being completely agnostic to intra-molecular hydrogen bonds. Furthermore, RDKit's calculation of rotatable bonds does not consider bonds to terminal atoms rotatable as it depends on the hybridization of the atom it is attached to. Consequently, bonds to attachment points are always considered non-rotatable. This is exemplified in Fig. 2b where the butane linker receives a ratio of 60/100. Consequently, a linker can never achieve a ratio of rotatable bonds of 100 and to achieve a higher ratio, linkers must become increasingly longer which can lead to unrealistic ideas being proposed. In practice, this is not a limitation in guiding Link-INVENT towards flexible/rigid linkers as one can introduce appropriate score transformations that provide meaningful agent feedback (discussed in the Results section). For a full list of properties available in the Link-INVENT Scoring Function, see the Supporting Information.

#### Results

We demonstrate the application of Link-INVENT by devising the following experiments:

- 1. **Illustrative Example:** A simple experiment to illustrate how Link-INVENT gradually learns to satisfy MPO objectives.
- 2. Experiment 1: Fragment Linking: Link two fragment hits and satisfy a molecular docking constraint.
- 3. **Experiment 2: Scaffold Hopping:** Generate new scaffold ideas to improve physico-chemical properties while retaining potency by satisfying a molecular docking constraint.

4. **Experiment 3: PROTACs:** Demonstrate the flexibility of Link-INVENT to generate linkers with diverse properties. The focus in this section is to showcase the linker specific properties implemented for the Link-INVENT Scoring Function.

The same prior was used for all the experiments and demonstrates the versatility of the single trained generative model in addressing diverse tasks.



**Fig. 3** Illustrative Example. The experiment was run in triplicate. The curve shows the average score achieved by the batch of molecules sampled at a given epoch and the upper and lower bounds of the shaded region represent the maximum and minimum score, respectively. The objective is to link two benzene rings while limiting the number of hydrogen bond donors (HBDs) and the linker containing exactly one ring. Example structures (linkers highlighted) are superimposed on the plot at various stages of training to illustrate how Link-INVENT gradually learns to generate molecules that satisfy the desired MPO objective.

**Illustrative Example.** As an initial illustrative example, we devise an experiment to link two benzene rings with the objective of limiting the number of HBDs and the linker possessing exactly one ring (Fig. 3). Correspondingly, the Scoring Function contains two components:

- Linker Number of Hydrogen Bond Donors: maximum reward is given if the linker contains no HBDs. See Supporting Information Fig. S1 for the score transformation.
- 2. Linker Number of Rings: reward is only given if the linker contains exactly one ring.

Fig. 3 shows the Link-INVENT training progress over 20 epochs. The average score over triplicate runs shown in the curve is gradually increasing. Example molecules generated over the course of training are superimposed on the plot. The first molecule on the left possesses multiple HBDs and the linker does not contain a ring. Consequently, this molecule receives low reward. As training progresses, the example molecules start to satisfy our MPO objective. Towards the end of the 20 epochs, the example molecule not only possesses no HBDs, but the linker also has exactly one ring. The purpose of this experiment was to illustrate how the Link-INVENT agent learns via RL to generate molecules that increasingly satisfy the target objective.



Fig. 4 Experiment 1: Fragment linking strategy for casein kinase 2 inhibitors for the alpha catalytic site (CK2α). a. Initial fragment hits. The fragment structures are colour-coded: gray fragment PDB ID: 5CSV and green fragment PDB ID: 5CSH. The gray fragment binds by forming hydrogen-bond interactions with Lys68 and Asp175 while the green fragment binds via hydrophobic interactions. The fragment linking strategy was to leverage the nitrogen atoms on both fragments to design a linear linker, separated by 9.9 Å. b. Fragment linking led to the discovery of the linked molecule, CAM4066 (PDB ID: 5CU4). The constituent fragments are circled in the structure. The linear linker features amide bonds that modulate the linker flexibility and rigidity which the authors attribute to its binding potency.<sup>47,48</sup>

Experiment 1: Fragment Linking. Fusco, Brear, et al. applied a fragment linking strategy (Fig. 4) to design casein kinase 2 inhibitors for the alpha catalytic site (CK2α).<sup>47,48</sup> CK2 is overexpressed in cancers and causes apoptosis evasion, leading to poor cancer prognosis and is therefore a therapeutic target of interest. Fig. 4a shows the fragment linking strategy envisioned by Fusco, Brear, et al. The gray fragment binds by forming hydrogen-bond interactions with Lys68 and Asp175 and the green fragment binds primarily through hydrophobic interactions (Fig. 4a). The difference in the binding modes of the constituent fragments were judged to be suitable for linking. Following this strategy, a potent and selective CK2 $\alpha$  inhibitor (CAM4066) was discovered that retains the important Lys68 hydrogen-bond interaction (Fig. 4b). CAM4066 was demonstrated to exhibit *in vivo* efficacy and is thus an example of a successful fragment linking campaign.<sup>47,48</sup>

In this section, we adopt the fragment linking strategy devised by Fusco, Brear, et al. (Fig. 4a) and task Link-INVENT with generating plausible linked molecules that retain the Lys68 hydrogen-bond interaction.<sup>47,48</sup> Moreover, while Fusco, Brear, et al. exclusively evaluated linear linker ideas, we allow Link-INVENT to explore linkers with rings and branching (to a certain extent). Correspondingly, we devise a Scoring Function composed of the following components:

- DockStream: this component is a molecular docking package that is fully compatible with Link-INVENT. DockStream supports docking using a variety of backends. In this work, we use Glide and LigPrep which we previously identified to yield the best average performance over a variety of receptor targets.<sup>49–54</sup> A docking constraint was enforced to retain the Lys68 hydrogen-bond interaction.<sup>47,48</sup> See Supporting Information Fig. S2 for the docking score transformation.
- Linker Length Ratio ≥ 70: this component prevents linkers with branching that is significantly longer than the effective length (number of bonds between the linker attachment atoms). See Supporting Information Fig. S3 for the Scoring Function transformation.
- 3. Linker Molecular Weight ≤ 200 Da: this component also prevents linkers with extensive branching but more importantly, prevents the Link-INVENT agent in exploiting the weaknesses of molecular docking, e.g., generating linkers that possess a large number of HBDs which may achieve a favourable docking score but at the expense of limited permeability.<sup>55</sup> See Supporting Information Fig. S3 for the Scoring Function transformation.



**Fig. 5** Experiment 1: Fragment Linking Link-INVENT results. The experiment was run in triplicate. The curve shows the average score achieved by the batch of molecules sampled at a given epoch and the upper and lower bounds of the shaded region represent the maximum and minimum score, respectively. **a.** Glide LigPrep docking score optimization. The average docking score achieved by the batch of compounds generated by Link-INVENT gradually more favourable (lower score in the case of Glide). **b.** Violin plots showing the distribution of docking scores for the triplicate runs. 'N' is the number of molecules generated over 100 epochs. The distributions are nearly identical, demonstrating reproducible experimental outcome. The black dotted line shows the docking score possessed by the reference ligand (CAM4066, -15.20 kcal/mol). **c.** Venn diagram plots showing the overlap between unique Bemis-Murcko scaffolds in the triplicate runs. **d.** 

The binding pose of a selected generated molecule (green) superimposed with the reference ligand (gray) and the constituent fragments are circled. **PDB ID: 5CU4.** The structure of the generated molecule is similar to the reference ligand. The yellow and turquoise dotted lines show the interactions formed by the reference ligand and generated ligand, respectively. The generated molecule retains the Lys68 interaction as enforced by the docking constraint and forms the same polar interactions as the reference ligand, largely attributing to the extensive overlap between the binding poses and supporting plausibility.

The fragment linking experiment was run in triplicate and the results are shown in Fig. 5 (See Supporting Information Fig. S4 for all training plots). Over the course of 100 epochs, the average Glide docking score of the batch of molecules generated by Link-INVENT gradually becomes more favourable (Fig. 5a). The docking scores distributions of the triplicate runs are essentially identical and demonstrate reproducible experimental outcome (Fig. 5b). The relatively few molecules that possess a docking score of 0 do not satisfy the docking constraint and were generated towards the beginning of the Link-INVENT run at a timestep where the agent has received minimal feedback. Furthermore, some molecules proposed by Link-INVENT exhibit a more favourable docking score than the reference ligand (-15.20 kcal/mol, black dotted line in Fig. 5b). The majority of the remaining molecules dock similar to the reference ligand (approximately -14 kcal/mol) and demonstrates that Link-INVENT at the very least, proposes chemical ideas that can satisfy the docking constraint. Subsequently, the interplay between the agent and the DF is exemplified in Fig. 5c. The DF encourages balance between agent exploration and exploitation by penalizing repeated sampling of identical Bemis-Murcko scaffolds.<sup>41</sup> The triplicate runs yield a large number of unique scaffolds with minimal overlap, demonstrating diversity in the results and showing that replicate experiments explore different areas in chemical space (Fig. 5c). Next, the plausibility of generated molecules was investigated by comparing their binding poses with the reference ligand. Fig. 5d shows the binding pose of an example top scoring molecule (based on satisfaction of the composite Scoring Function) superimposed with the reference ligand (See Supporting Information Fig. S5 for more examples). Firstly, the proposed linker is similar to the ground-truth linker, differing only by a single atom shift of an amide bond and the presence of an additional nitrogen. It is important to note that information about the reference ligand was not available to the Link-INVENT agent during the generative process. Fusco, Brear, et al. posits that the flexibility and rigidity of the reference ligand linker is crucial to its potency.<sup>47,48</sup> The similarity in the linker proposed by Link-INVENT suggests that the docking constraint implicitly guides the agent towards 3D structural awareness, in agreement with our previous results.<sup>49</sup> This is further supported by the predicted polar interactions of the generated molecule (Fig. 5d turquoise dotted lines) being mostly identical to those of the reference ligand (Fig. 5d yellow dotted lines) with the only exception being His160. Consequently, the structural similarity between the linkers naturally results in significant overlap of the binding poses and is exemplified in the docking score in which the generated molecule is predicted to dock more favourably than the reference ligand. Taken together, the results in this section demonstrate that Link-INVENT is able to generate plausible chemical ideas spanning diverse minima and is easily tuned for bespoke applications via the Scoring Function.



**Fig. 6** Experiment 2: Scaffold Hopping strategy for dual leucine zipper kinase (DLK) inhibitor optimization. **a.** Initial inhibitor possessing poor physico-chemical properties causing *in vivo* high clearance (**PDB ID: 5CEO**). The two hydrogen-bonds in the hinge region with Cys193 are crucial for potency. The goal was to replace the pyridine core while retaining the Cys193 interactions. **b.** Scaffold hopping led to the discovery of a DLK inhibitor with a pyrazole core and with demonstrated *in vivo* efficacy (**PDB ID: 5CEQ**).<sup>56,57</sup> The retained molecular sub-units are circled in the structure.

**Experiment 2: Scaffold Hopping.** The c-Jun N-terminal kinase (JNK) pathway is implicated in neuronal injury and neurodegeneration and is a therapeutic target of interest. Patel et al. aimed to develop a small molecule inhibitor to modulate this pathway via targeting dual leucine zipper kinase (DLK) which is an upstream JNK regulator. Initial efforts led to the development of a potent and selective inhibitor but with high *in vivo* clearance (Fig. 6a).<sup>56,57</sup> Subsequently, in a more recent work, Patel et al. applied a scaffold hopping strategy to improve the physico-chemical properties of their initial inhibitor to achieve central nervous system (CNS) penetration (Fig. 6b). Scaffold hopping from the pyridine core to a pyrazole core led to the discovery of a DLK inhibitor with *in vivo* efficacy (Fig. 6b).<sup>57</sup>

In this section, we adopt the scaffold hopping strategy devised by Patel et al. and task Link-INVENT with generated novel core ideas with a focus on improving CNS properties. A docking constraint to enforce the Cys193 hydrogen-bond interactions is applied to retain predicted potency and the

b.

following specific physico-chemical properties, adopted from Patel et al., were enforced<sup>57</sup>: the number of HBDs must be less than 2, the topological polar surface area (tPSA) must be less than 90 Å<sup>2</sup>, and the CNS MPO score must be greater than or equal to 4. The CNS MPO is an algorithm developed from analysis of CNS drugs and candidates as a predictor for CNS efficacy and encompasses six physicochemical properties (ClogP, ClogD, MW, tPSA, number of HBDs, and pKa).<sup>58</sup> In the devised experiment, we do not account for all six CNS MPO properties and only enforce logP, MW, tPSA, and number of HBDs. Correspondingly, we define the Scoring Function with the following components:

- DockStream: this component is identical to the usage described in the Fragment Linking section. The only exception was that the docking constraint was enforced to retain the Cys193 hydrogenbond interactions in the hinge region.<sup>57</sup> See Supporting Information Fig. S6 for the docking score transformation.
- Number of Hydrogen Bond Donors < 2: this component is included in the CNS MPO algorithm and enforces the overall linked molecule to possess less than two HBDs. This quantity was specifically desired by Patel et al.<sup>57</sup> See Supporting Information Fig. S7 for the Scoring Function transformation.
- 3. Molecular Weight ≤ 450 Da: this component is included in the CNS MPO algorithm and is enforced to be in an interval in agreement with CNS penetration but with some leniency to allow more Link-INVENT exploration of chemical space.<sup>58</sup> See Supporting Information Fig. S7 for the Scoring Function transformation.

- 3 ≤ SlogP ≤ 4: this component is included in the CNS MPO algorithm and is enforced to be in an interval in agreement with CNS penetration.<sup>58</sup> See Supporting Information Fig. S7 for the Scoring Function transformation.
- 5. tPSA ≤ 90 Å<sup>2</sup>: this component is included in the CNS MPO algorithm and is enforced to be in an interval in agreement with CNS penetration.<sup>58</sup> The interval was also specifically desired by Patel et al.<sup>57</sup> See Supporting Information Fig. S7 for the Scoring Function transformation.
- 1 ≤ Linker Number of Aromatic Rings ≤ 2: this component was specifically desired by Patel et al. as the binding site geometry is most compatible with a planar ring present in the core scaffold.<sup>57</sup> See Supporting Information Fig. S7 for the Scoring Function transformation.

b.



-11.35 kcal/mol

**Fig. 7** Experiment 2: Scaffold Hopping Link-INVENT results. The experiment was run in triplicate. The curve shows the average score achieved by the batch of molecules sampled at a given epoch and the upper and lower bounds of the shaded region represent the maximum and minimum score, respectively. The analysis is identical to the Fragment Linking experiments. **a.** Glide LigPrep docking score optimization. The average docking score achieved by the batch of compounds generated by Link-INVENT. **b.** Violin plots showing the distribution of docking scores for the triplicate runs. 'N' is the number of molecules generated over 100 epochs that satisfy all the CNS criteria. The black dotted line shows the docking score possessed by the reference ligand (-11.26 kcal/mol). **c.** Venn diagram plots showing the overlap between unique Bemis-Murcko scaffolds in the triplicate runs. **d.** The binding pose of a selected generated molecule (green) superimposed with the reference ligand (gray) and the retained molecular sub-units are circled. **PDB ID: 5CEO.** The yellow and turquoise dotted lines

show the interactions formed by the reference ligand and generated ligand, respectively. The generated molecule retains the Cys193 interaction as enforced by the docking constraint and is predicted to form an additional hydrogen-bond interaction with Gln195. The extensive overlap between the binding poses of the generated ligand and the reference ligand supports plausibility.

The scaffold hopping experiment was run in triplicate and the results are shown in Fig. 7 (see Supporting Information Fig. S8 for all training plots). Over the course of 100 epochs, the average Glide docking score of the batch of molecules generated by Link-INVENT gradually becomes more favourable (Fig. 7a) and the similarity in the docking scores distributions demonstrate reproducible experimental outcome (Fig. 7b). In contrast to the fragment linking experiment, relatively few molecules possess a more favourable docking score than the reference ligand (shown by the black dotted line). Instead, the majority of molecules score slightly worse (approximately -9.5 kcal/mol). This is not completely unexpected as the MPO objective is significantly more challenging than the previous fragment linking case study. Consequently, the solution space is much narrower. It is important to note, however, that the objective of the scaffold hopping experiment is not strictly to propose novel cores that dock better than the initial inhibitor (Fig. 6a). Patel et al. note that their initial inhibitor, while potent, exhibits high in vivo clearance.<sup>56,57</sup> Therefore, an inhibitor with *sufficient* binding affinity and good CNS penetration could achieve in vivo efficacy. The narrower solution space in the scaffold hopping experiment is further supported by Fig. 7c where the absolute counts of unique Bemis-Murcko scaffolds is less than the fragment linking experiment.<sup>41</sup> This is not a limitation of Link-INVENT but rather the nature of the MPO objective. Nonetheless, the absolute counts for the generated scaffolds is still high and demonstrates Link-INVENT samples from diverse minima. Similar to the fragment linking results, minimal overlap between replicate runs show that replicate experiments explore different areas in chemical space (Fig. 7c). The plausibility of the proposed scaffolds was investigated by comparing their binding poses with the reference ligand. Fig. 7d shows the binding pose of an example top scoring molecule (based on satisfaction of the composite Scoring Function) superimposed with the reference ligand (see Supporting Information Fig. S9 for more examples). Firstly, the proposed scaffold features planar aromatic rings, as enforced by the Scoring Function, and as desired by Patel et al.<sup>57</sup> Secondly, the Cys193 hydrogen-bond interactions are retained, as enforced by the docking constraint. The proposed ligand is predicted to form an additional hydrogen-bond with Gln195, owing to the hydrocarbon chain that extends the spatial occupancy of the overall molecule (Fig. 7d). This suggests that the application of a docking constraint can guide the Link-INVENT agent towards 3D structural awareness, learning to exploit the binding site geometry and electronics. Lastly, the binding poses of the generated ligand and the reference ligand overlap significantly, supporting plausibility. Taken together, the results in this section demonstrate the flexibility of the Link-INVENT Scoring Function to optimize relatively complex MPO objectives and that the agent learns to propose plausible chemical ideas.



**Fig. 8** Experiment 3: PROTACs strategy for Bcl-2 and Mcl-1 dual degradation. **a.** Molecular dynamics (MD) simulated ternary complex of one of the developed PROTACs binding to Mcl-1 (orange) and cereblon (CRBN, blue). The ternary complex is stabilized by protein-protein interactions (PPIs). The linker is an alkyl chain and is circled in the structure. Mcl-1 **PDB ID: 2PQK** and CRBN **PDB ID: 4TZ4**. **b.** The naphthalimide-based μM inhibitor which was linked with pomalidomide to form the ternary complex. The red asterisks denote the linker attachment points. The two linkers that achieved potent and selective *in vitro* degradation are shown.

**Experiment 3: PROTACs.** B-cell lymphoma 2 (Bcl-2) and myeloid leukemia 1 (Mcl-1) are anti-apoptotic proteins which can inhibit intrinsic apoptosis, i.e., induced by mitochondrial stress, and are therapeutic targets of interests. Wang et al. designed linkers between a naphthalimide-based μM inhibitor with pomalidomide which is a cereblon (CRBN) binding ligand (Fig. 8).<sup>59</sup> The resulting ternary complex would undergo ubiquitination and lead to targeted dual degradation of Bcl-2 and Mcl-1. Due to challenges in obtaining crystal structures as a basis to form hypotheses on optimal linker lengths, PROTACs linker

design has mostly been empirical.<sup>32</sup> Wang et al. adopt an iterative SAR approach to investigate the effect of linker length on Bcl-2 and Mcl-1 dual degradation, whereby a linker too short would cause steric clash and prevent the formation of the stable ternary complex and a linker too long could result in too much conformational entropy to overcome.<sup>32,59</sup> Following this approach, Wang et al. successfully transform a low-affinity and non-selective ligand for Bcl-2 and Mcl-1 into PROTACs stabilized via protein-protein interactions (PPIs) and show potent and selective dual degradation *in vitro*.<sup>59</sup> Fig. 8a and 8b show the molecular dynamics (MD) simulated ternary complex of one of the discovered PROTACs performed by Wang et al. and the general linking strategy envisioned, respectively.<sup>59</sup>

In this section, we use the PROTACs design strategy by Wang et al. to demonstrate Link-INVENT's linker specific components for the Scoring Function. In select experiments, a fixed set of physicochemical properties was enforced and based on observed values from compiled PROTACs databases.<sup>30,31</sup> Correspondingly, we define the Scoring Function with the following components (see Supporting Information Fig. S10 for the Scoring Function Transformations):

#### 1. tPSA ≤ 250 Å<sup>2</sup>

- 2. 3.5 ≤ logP ≤ 6.0
- 3. Number of Hydrogen Bond Acceptors ≤ 16
- 4. Number of Hydrogen Bond Donors  $\leq 6$
- 5. Number of Rotatable Bonds < 25

We demonstrate control over the properties of generated linkers while keeping physico-chemical properties of the PROTAC within the specified intervals described above. Subsequently, we devise three Sub-Experiments:

- Sub-Experiment 1: Fix physico-chemical properties and control linker length. We show that Link-INVENT can generate linkers within a specified narrow length interval. In addition to including the physico-chemical properties listed above, the Scoring Function contains the following components:
  - Linker Effective Length = [4,6], [7,9], [10,12], or [13,15]: this component enforces linkers to possess an effective length within the specified intervals. See Supporting Information S11 for the Scoring Function transformation.
  - 2) Linker Length Ratio = 100 this component prevents linker branching.

The combination of components 1 and 2 enforce Link-INVENT to generate linkers without branching.

Sub-Experiment 2: Fix physico-chemical properties, linker length within the interval [7,9], and control linker linearity, i.e., linkers with and without rings. We show that Link-INVENT can generate linkers within a specified narrow length interval and control for the presence of rings. In addition to including the physico-chemical properties listed above, the Scoring Function contains the following component (see Supporting Information Fig. S17 for the Scoring Function Transformations):

- Linker Effective Length = [7,9]: this component enforces linkers to possess an effective length within the specified interval of [7,9].
- 2) Linker Length Ratio = 100 this component prevents linker branching.
- 3) Linker Number of Rings = 0 this component enforces linkers to possess no rings, i.e., the linker is linear. In the experiment where we want to generate linkers with rings, we simply omit this component in the Scoring Function.

Similar to Sub-Experiment 1, components 1 and 2 enforce Link-INVENT to generate linkers without branching.

- 3. **Sub-Experiment 3:** In this Sub-Experiment, no length or physico-chemical properties are enforced. Instead, we task Link-INVENT with generating linkers with variable flexibility which is defined by the "linker ratio of rotatable bonds" component, i.e., ratio between the number of rotatable bonds over total number of bonds. Correspondingly, the Scoring Function contains only one component:
  - Linker Ratio of Rotatable Bonds = [0,30], [40,60], [70,100]: the defined intervals correspond to "Low", "Moderate", and "High" flexibility (see Supporting Information S21 for the Scoring Function transformation).



**Fig. 9** Experiment 3: PROTACs Link-INVENT results. The experiments were run in triplicate. The curve in **c**. shows the average score achieved by the batch of molecules sampled at a given epoch and the upper and lower bounds of the shaded region represent the maximum and minimum score, respectively. **a**. Experiment that fixes physico-chemical properties and tasks Link-INVENT with generating linkers with effective length within the specified intervals: [4,6], [7,9], [10,12], and [13,15]. The baseline experiment does not enforce linker length. Consequently, the effective linker lengths resemble a broad distribution. The generated linkers from the triplicate runs which also satisfy all the physico-chemical properties criteria are aggregated in the plot. **b**. Experiment that fixes physico-chemical properties and effective linker length within the interval [7-9]. Link-INVENT is tasked with generating linear and cyclic linkers. The baseline experiment does not enforce linker length a 1:2 ratio of linear:cyclic is observed. The generated linkers from the triplicate runs which also satisfy all the physico-chemical properties criteria are aggregated in the plot. **c**. Experiment that tasks Link-INVENT with generating linkers containing ratio of rotatable bonds within the specified intervals: "Low" [0, 30], "Moderate" [40,60], "High" [70,100]. The agent implicitly

learns that linkers containing rings and sp2 hybridized atoms achieve a low ratio of rotatable bonds. Conversely, linear linkers with sp3 hybridized atoms achieve a high ratio of rotatable bonds.

**PROTACs Sub-Experiment 1: Controlling Linker Length.** Link-INVENT was tasked with generating linker ideas of variable length while keeping physico-chemical properties within a specified range (Fig. 9a, see Supporting Information Fig. S12-16 for all training plots). The baseline experiment does not enforce a specific effective linker length interval and the distribution of lengths span a large range (Fig. 9a). In contrast, one can enforce the Link-INVENT agent to explore effective linker lengths within a certain interval, as shown by the enrichments observed in Fig. 9a, e.g., the 'enforce 4-6' experiment enforced effective linker lengths in the interval [4-6] and the corresponding bar is enriched relative to other lengths. The purpose of this Sub-Experiment is to show the ease in which one can control effective linker length exploration, mimicking a real-world PROTAC linker design campaign.<sup>32,59</sup>

**PROTACs Sub-Experiment 2: Controlling Linker Linearity.** Link-INVENT was tasked with generating linker ideas with effective length in the interval [7,9] while keeping physico-chemical properties within a specified range and controlling linearity (Fig. 9b, see Supporting Information Fig. S18-20 for all training plots). The baseline experiment does not enforce linearity and the resulting ratio of linear linkers to cyclic linkers, i.e., linkers containing at least one ring, is approximately 1:2. In contrast, one can enforce the Link-INVENT agent to explore linear linkers or cyclic linkers, shown by the enrichments observed in Fig. 9b. The purpose of this Sub-Experiment is to further showcase the user flexibility in specifying desired linker properties.

**PROTACs Sub-Experiment 3: Controlling Linker Flexibility.** This Sub-Experiment showcases Link-INVENT's "linker ratio of rotatable bonds" component which can be specified in the Scoring Function. We note that while the component itself is meant to be a descriptor of linker flexibility, inherent

limitations exist, e.g., not accounting for intra-molecular hydrogen-bonding interactions which would rigidify the linker. Link-INVENT was tasked with generating linker ideas with variable ratios of rotatable bonds where we define 'Low', 'Moderate', and 'High' as the intervals [0,30], [40,60], and [70,100], respectively (Fig. 9c, see Supporting Information Fig. S22 for all training plots). Examples of linkers possessing variable degrees of flexibility are shown in Fig. 9c. The agent implicitly learns that linkers containing rings and sp2 hybridized atoms are more rigid. A clear transition from "Low" flexibility to "High" flexibility is marked by increasing linearity and sp3 hybridized atoms. Without enforcing any length constraints, proposed linkers become increasingly longer to achieve a high "linker ratio of rotatable bonds" value. This is exemplified in the example linker in the "High" experiment (Fig. 9c). Naturally, the linker shown is likely unrealistic and this Sub-Experiment was an extreme example to showcase the flexibility of Link-INVENT's Scoring Function. In practice, one could constrain the linker length within a specified interval as was done in Sub-Experiments 1 and 2 and explore variable flexibility. In this regard, the "linker ratio of rotatable bonds" provides some control over the conformational entropy of proposed linker ideas.

#### Conclusions

In this work, we introduced Link-INVENT as an extension to the *de novo* design platform, REINVENT.<sup>6</sup> Link-INVENT is a recurrent neural network (RNN)-based generative model trained to propose linker ideas given two input molecular subunits. In contrast to previous methods for linker design involving database searching which are inherently limited to a pre-defined collection of chemical ideas<sup>33–36</sup>, Link-INVENT builds linkers at the token level, proposing linkers as SMILES and can generalize in chemical space.<sup>15,42</sup> Moreover, Link-INVENT's Scoring Function provides users with the ability to optimize bespoke multiparameter optimization (MPO) objectives via reinforcement learning (RL), offering control that is not present in existing deep learning (DL)-based approaches to linker generation.<sup>16–19</sup> A vast number of molecular properties can be optimized, ranging from previously implemented components available in REINVENT that operate on the entire molecule to newly implemented linker specific components which provide control over linker properties.

We demonstrate the application of Link-INVENT on three case studies encompassing fragment linking<sup>20,21</sup>, scaffold hopping<sup>22</sup>, and PROTACs design<sup>23–25</sup>. The Scoring Functions for the experiments were devised based from the corresponding fragment linking<sup>47,48</sup>, scaffold hopping<sup>56,57</sup>, and PROTACs design<sup>59</sup> works. We illustrate the practical adoption of Link-INVENT to real-world drug discovery projects by showcasing how to translate experimental insights into an informative Scoring Function for Link-INVENT. Subsequently, the agent learned to satisfy the desired MPO objective via RL. Specifically, in the fragment linking experiment<sup>47,48</sup>, we showed that Link-INVENT can propose plausible linker ideas that satisfy a molecular docking constraint with additional constraint over the permitted linker spatial occupancy by controlling for branching. More than 5000 unique Bemis-Murcko scaffolds were generated by the Link-INVENT agent, demonstrating that diverse linker ideas were explored.<sup>41</sup> In the scaffold hopping experiment<sup>56,57</sup>, we showed that Link-INVENT can simultaneously optimize a relatively complex MPO objective encompassing a molecular docking constraint and favourable central nervous system (CNS) compatible physico-chemical properties. In this experiment, Link-INVENT navigated a narrow solution space and proposed plausible scaffold ideas which satisfy all desired properties and are diverse as shown by the number of unique Bemis-Murcko scaffolds.<sup>41</sup> In the PROTACs experiment<sup>59</sup>, we further showed Link-INVENT's extensive user control on the linker properties. We demonstrated the ability to enforce the Link-INVENT agent to explore effective linker lengths within a specified interval while keeping physico-chemical properties within a specified range. Moreover, linker linearity can be controlled, enforcing the agent to explore only linear linkers or linkers containing rings. Finally, we prove that linker flexibility can be controlled via the "linker ratio of rotatable bonds" component which provides users with the ability to modulate the conformational entropy of proposed linker ideas. These series of PROTACs Sub-Experiments mimic real-world PROTACs linker design which typically investigate linkers of variable length and flexibility.<sup>32,59</sup>

Link-INVENT is a ready-to-use generative model for linker design with the capability to optimize bespoke MPO objectives via the flexible Scoring Function. The case studies in this work show how Link-INVENT can be applied to real-world drug discovery projects and that the agent proposes plausible and diverse linker ideas. The code is freely available at <u>https://github.com/MolecularAl/Reinvent</u>.

# **Associated Content**

#### Supporting Information

- Details related to the data preparation
- Details on the vocabulary of the Link-INVENT model
- Details on the new linker specific components implemented in Link-INVENT
- Details on the docking protocol used including parameters
- Hardware information and experiment computation times
- All training plots for the experiments presented in this work
- More example binding poses for experiments 1 and 2

## References

- Schneider, P.; Walters, W. P.; Plowright, A. T.; Sieroka, N.; Listgarten, J.; Goodnow, R. A.; Fisher, J.; Jansen, J. M.; Duca, J. S.; Rush, T. S.; Zentgraf, M.; Hill, J. E.; Krutoholow, E.; Kohler, M.; Blaney, J.; Funatsu, K.; Luebkemann, C.; Schneider, G. Rethinking Drug Design in the Artificial Intelligence Era. *Nat. Rev. Drug Discov.* **2020**, *19* (5), 353–364. https://doi.org/10.1038/s41573-019-0050-3.
- (2) Jiménez-Luna, J.; Grisoni, F.; Weskamp, N.; Schneider, G. Artificial Intelligence in Drug Discovery: Recent Advances and Future Perspectives. *Expert Opin. Drug Discov.* 2021, 1–11. https://doi.org/10.1080/17460441.2021.1909567.
- (3) Polishchuk, P. G.; Madzhidov, T. I.; Varnek, A. Estimation of the Size of Drug-like Chemical Space Based on GDB-17 Data. J. Comput. Aided Mol. Des. 2013, 27 (8), 675–679. https://doi.org/10.1007/s10822-013-9672-4.

- (4) Beam Search for Automated Design and Scoring of Novel ROR Ligands with Machine Intelligence\*\* - Moret - 2021 - Angewandte Chemie International Edition - Wiley Online Library https://onlinelibrary.wiley.com/doi/full/10.1002/anie.202104405 (accessed 2022 -03 -19).
- (5) Popova, M.; Isayev, O.; Tropsha, A. Deep Reinforcement Learning for de Novo Drug Design. *Sci. Adv.* **2018**, *4* (7), eaap7885. https://doi.org/10.1126/sciadv.aap7885.
- Blaschke, T.; Arús-Pous, J.; Chen, H.; Margreitter, C.; Tyrchan, C.; Engkvist, O.; Papadopoulos, K.; Patronov, A. REINVENT 2.0: An AI Tool for De Novo Drug Design. J. Chem. Inf. Model. 2020, 60 (12), 5918–5922. https://doi.org/10.1021/acs.jcim.0c00915.
- (7) Thomas, M.; Smith, R. T.; O'Boyle, N. M.; de Graaf, C.; Bender, A. Comparison of Structure- and Ligand-Based Scoring Functions for Deep Generative Models: A GPCR Case Study. J. Cheminformatics 2021, 13 (1), 39. https://doi.org/10.1186/s13321-021-00516-0.
- (8) Li, Y.; Zhang, L.; Liu, Z. Multi-Objective de Novo Drug Design with Conditional Graph Generative Model. *J. Cheminformatics* **2018**, *10* (1), 33. https://doi.org/10.1186/s13321-018-0287-6.
- (9) Mercado, R.; Rastemo, T.; Lindelöf, E.; Klambauer, G.; Engkvist, O.; Chen, H.; Jannik Bjerrum, E. Graph Networks for Molecular Design. *Mach. Learn. Sci. Technol.* 2021, 2 (2), 025023. https://doi.org/10.1088/2632-2153/abcf91.
- (10) Atance, S. R.; Diez, J. V.; Engkvist, O.; Olsson, S.; Mercado, R. De Novo Drug Design Using Reinforcement Learning with Graph-Based Deep Generative Models. 2021. https://doi.org/10.26434/chemrxiv-2021-9w3tc.
- (11) Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; Aspuru-Guzik, A. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Cent. Sci.* 2018, 4 (2), 268–276. https://doi.org/10.1021/acscentsci.7b00572.
- (12) Bai, Q.; Tan, S.; Xu, T.; Liu, H.; Huang, J.; Yao, X. MolAICal: A Soft Tool for 3D Drug Design of Protein Targets by Artificial Intelligence and Classical Algorithm. *Brief. Bioinform.* 2021, 22 (3). https://doi.org/10.1093/bib/bbaa161.
- (13) Choi, J.; Lee, J. V-Dock: Fast Generation of Novel Drug-like Molecules Using Machine-Learning-Based Docking Score and Molecular Optimization. 2021. https://doi.org/10.33774/chemrxiv-2021-75t5k.
- (14) Nigam, A.; Pollice, R.; Aspuru-Guzik, A. JANUS: Parallel Tempered Genetic Algorithm Guided by Deep Neural Networks for Inverse Molecular Design. *ArXiv210604011 Cs* **2021**.
- (15) Weininger, D. SMILES, a Chemical Language and Information System. 1. Introduction to Methodology and Encoding Rules. J. Chem. Inf. Comput. Sci. 1988, 28 (1), 31–36. https://doi.org/10.1021/ci00057a005.
- (16) Imrie, F.; Bradley, A. R.; van der Schaar, M.; Deane, C. M. Deep Generative Models for 3D Linker Design. J. Chem. Inf. Model. **2020**, 60 (4), 1983–1995. https://doi.org/10.1021/acs.jcim.9b01120.
- (17) Yang, Y.; Zheng, S.; Su, S.; Zhao, C.; Xu, J.; Chen, H. SyntaLinker: Automatic Fragment Linking with Deep Conditional Transformer Neural Networks. *Chem. Sci.* **2020**, *11* (31), 8312–8322. https://doi.org/10.1039/D0SC03126G.
- (18) Hu, L.; Yang, Y.; Zheng, S.; Xu, J.; Ran, T.; Chen, H. Kinase Inhibitor Scaffold Hopping with Deep Learning Approaches. J. Chem. Inf. Model. 2021, 61 (10), 4900–4912. https://doi.org/10.1021/acs.jcim.1c00608.
- (19) Langevin, M.; Minoux, H.; Levesque, M.; Bianciotto, M. Scaffold-Constrained Molecular Generation. J. Chem. Inf. Model. 2020, 60 (12), 5637–5646. https://doi.org/10.1021/acs.jcim.0c01015.
- (20) Erlanson, D. A.; McDowell, R. S.; O'Brien, T. Fragment-Based Drug Discovery. *J. Med. Chem.* **2004**, *47* (14), 3463–3482. https://doi.org/10.1021/jm040031v.
- (21) Murray, C. W.; Rees, D. C. The Rise of Fragment-Based Drug Discovery. *Nat. Chem.* **2009**, *1* (3), 187–192. https://doi.org/10.1038/nchem.217.

- (22) Böhm, H.-J.; Flohr, A.; Stahl, M. Scaffold Hopping. *Drug Discov. Today Technol.* **2004**, *1* (3), 217–224. https://doi.org/10.1016/j.ddtec.2004.10.009.
- (23) Sakamoto, K. M.; Kim, K. B.; Kumagai, A.; Mercurio, F.; Crews, C. M.; Deshaies, R. J. Protacs: Chimeric Molecules That Target Proteins to the Skp1–Cullin–F Box Complex for Ubiquitination and Degradation. *Proc. Natl. Acad. Sci. U. S. A.* **2001**, *98* (15), 8554–8559. https://doi.org/10.1073/pnas.141230798.
- (24) Toure, M.; Crews, C. M. Small-Molecule PROTACS: New Approaches to Protein Degradation. *Angew. Chem. Int. Ed.* **2016**, *55* (6), 1966–1973. https://doi.org/10.1002/anie.201507978.
- (25) Sun, X.; Gao, H.; Yang, Y.; He, M.; Wu, Y.; Song, Y.; Tong, Y.; Rao, Y. PROTACs: Great Opportunities for Academia and Industry. *Signal Transduct. Target. Ther.* **2019**, *4* (1), 1–33. https://doi.org/10.1038/s41392-019-0101-6.
- (26) Chodera, J. D.; Mobley, D. L. Entropy-Enthalpy Compensation: Role and Ramifications in Biomolecular Ligand Recognition and Design. *Annu. Rev. Biophys.* 2013, 42, 121–142. https://doi.org/10.1146/annurev-biophys-083012-130318.
- (27) Bancet, A.; Raingeval, C.; Lomberget, T.; Le Borgne, M.; Guichou, J.-F.; Krimm, I. Fragment Linking Strategies for Structure-Based Drug Design. *J. Med. Chem.* **2020**, *63* (20), 11420–11435. https://doi.org/10.1021/acs.jmedchem.0c00242.
- (28) Yu, H. S.; Modugula, K.; Ichihara, O.; Kramschuster, K.; Keng, S.; Abel, R.; Wang, L. General Theory of Fragment Linking in Molecular Design: Why Fragment Linking Rarely Succeeds and How to Improve Outcomes. J. Chem. Theory Comput. 2021, 17 (1), 450–462. https://doi.org/10.1021/acs.jctc.0c01004.
- (29) Lipinski, C. A.; Lombardo, F.; Dominy, B. W.; Feeney, P. J. Experimental and Computational Approaches to Estimate Solubility and Permeability in Drug Discovery and Development Settings. Adv. Drug Deliv. Rev. 1997, 23 (1), 3–25. https://doi.org/10.1016/S0169-409X(96)00423-1.
- (30) Maple, H. J.; Clayden, N.; Baron, A.; Stacey, C.; Felix, R. Developing Degraders: Principles and Perspectives on Design and Chemical Space. *MedChemComm* 2019, *10* (10), 1755–1764. https://doi.org/10.1039/C9MD00272C.
- Weng, G.; Shen, C.; Cao, D.; Gao, J.; Dong, X.; He, Q.; Yang, B.; Li, D.; Wu, J.; Hou, T. PROTAC-DB: An Online Database of PROTACs. *Nucleic Acids Res.* 2021, 49 (D1), D1381–D1387. https://doi.org/10.1093/nar/gkaa807.
- Bemis, T. A.; La Clair, J. J.; Burkart, M. D. Unraveling the Role of Linker Design in Proteolysis Targeting Chimeras. J. Med. Chem. 2021, 64 (12), 8042–8052. https://doi.org/10.1021/acs.jmedchem.1c00482.
- (33) Cheeseright, T.; Mackey, M.; Rose, S.; Vinter, A. Molecular Field Extrema as Descriptors of Biological Activity: Definition and Validation. *J. Chem. Inf. Model.* **2006**, *46* (2), 665–676. https://doi.org/10.1021/ci050357s.
- (34) Maass, P.; Schulz-Gasch, T.; Stahl, M.; Rarey, M. Recore: A Fast and Versatile Method for Scaffold Hopping Based on Small Molecule Crystal Structure Conformations. *J. Chem. Inf. Model.* 2007, 47 (2), 390–399. https://doi.org/10.1021/ci060094h.
- (35) Thompson, D. C.; Denny, R. A.; Nilakantan, R.; Humblet, C.; Joseph-McCarthy, D.; Feyfant, E. CONFIRM: Connecting Fragments Found in Receptor Molecules. *J. Comput. Aided Mol. Des.* 2008, *22* (10), 761–772. https://doi.org/10.1007/s10822-008-9221-8.
- (36) Vainio, M. J.; Kogej, T.; Raubacher, F.; Sadowski, J. Scaffold Hopping by Fragment Replacement. *J. Chem. Inf. Model.* **2013**, *53* (7), 1825–1835. https://doi.org/10.1021/ci4001019.
- (37) Gaulton, A.; Hersey, A.; Nowotka, M.; Bento, A. P.; Chambers, J.; Mendez, D.; Mutowo, P.; Atkinson, F.; Bellis, L. J.; Cibrián-Uhalte, E.; Davies, M.; Dedman, N.; Karlsson, A.; Magariños, M. P.; Overington, J. P.; Papadatos, G.; Smit, I.; Leach, A. R. The ChEMBL Database in 2017. *Nucleic Acids Res.* 2017, 45 (D1), D945–D954. https://doi.org/10.1093/nar/gkw1074.

- (38) LibINVENT: Reaction-based Generative Scaffold Decoration for in Silico Library Design | Journal of Chemical Information and Modeling https://pubs.acs.org/doi/10.1021/acs.jcim.1c00469 (accessed 2022 -01 -01).
- (39) Arús-Pous, J.; Patronov, A.; Bjerrum, E. J.; Tyrchan, C.; Reymond, J.-L.; Chen, H.; Engkvist, O. SMILES-Based Deep Generative Scaffold Decorator for de-Novo Drug Design. *J. Cheminformatics* 2020, *12* (1), 38. https://doi.org/10.1186/s13321-020-00441-8.
- (40) Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9* (8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735.
- (41) Bemis, G. W.; Murcko, M. A. The Properties of Known Drugs. 1. Molecular Frameworks. *J. Med. Chem.* **1996**, *39* (15), 2887–2893. https://doi.org/10.1021/jm9602928.
- (42) Arús-Pous, J.; Johansson, S. V.; Prykhodko, O.; Bjerrum, E. J.; Tyrchan, C.; Reymond, J.-L.; Chen,
   H.; Engkvist, O. Randomized SMILES Strings Improve the Quality of Molecular Generative
   Models. J. Cheminformatics 2019, 11 (1), 71. https://doi.org/10.1186/s13321-019-0393-0.
- (43) Williams, R. J.; Zipser, D. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Comput.* **1989**, *1* (2), 270–280. https://doi.org/10.1162/neco.1989.1.2.270.
- (44) Olivecrona, M.; Blaschke, T.; Engkvist, O.; Chen, H. Molecular De-Novo Design through Deep Reinforcement Learning. J. Cheminformatics 2017, 9 (1), 48. https://doi.org/10.1186/s13321-017-0235-x.
- (45) Blaschke, T.; Engkvist, O.; Bajorath, J.; Chen, H. Memory-Assisted Reinforcement Learning for Diverse Molecular de Novo Design. J. Cheminformatics 2020, 12 (1), 68. https://doi.org/10.1186/s13321-020-00473-0.
- (46) RDKit: Open-Source Cheminformatics. Https://Www.Rdkit.Org, 2019.
- (47) De Fusco, C.; Brear, P.; legre, J.; Georgiou, K. H.; Sore, H. F.; Hyvönen, M.; Spring, D. R. A Fragment-Based Approach Leading to the Discovery of a Novel Binding Site and the Selective CK2 Inhibitor CAM4066. *Bioorg. Med. Chem.* **2017**, *25* (13), 3471–3482. https://doi.org/10.1016/j.bmc.2017.04.037.
- (48) Brear, P.; De Fusco, C.; Hadje Georgiou, K.; Francis-Newton, N. J.; Stubbs, C. J.; Sore, H. F.; Venkitaraman, A. R.; Abell, C.; Spring, D. R.; Hyvönen, M. Specific Inhibition of CK2α from an Anchor Outside the Active Site. *Chem. Sci.* **2016**, *7* (11), 6839–6845. https://doi.org/10.1039/C6SC02335E.
- (49) Guo, J.; Janet, J. P.; Bauer, M. R.; Nittinger, E.; Giblin, K. A.; Papadopoulos, K.; Voronov, A.; Patronov, A.; Engkvist, O.; Margreitter, C. DockStream: A Docking Wrapper to Enhance de Novo Molecular Design. J. Cheminformatics 2021, 13 (1), 89. https://doi.org/10.1186/s13321-021-00563-7.
- (50) Schrödinger Release 2019-4: Glide, Schrödinger, LLC, New York, NY, 2019.
- (51) Friesner, R. A.; Banks, J. L.; Murphy, R. B.; Halgren, T. A.; Klicic, J. J.; Mainz, D. T.; Repasky, M. P.; Knoll, E. H.; Shelley, M.; Perry, J. K.; Shaw, D. E.; Francis, P.; Shenkin, P. S. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy. *J. Med. Chem.* **2004**, *47* (7), 1739–1749. https://doi.org/10.1021/jm0306430.
- Halgren, T. A.; Murphy, R. B.; Friesner, R. A.; Beard, H. S.; Frye, L. L.; Pollard, W. T.; Banks, J. L. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 2. Enrichment Factors in Database Screening. *J. Med. Chem.* 2004, 47 (7), 1750–1759. https://doi.org/10.1021/jm030644s.
- (53) Friesner, R. A.; Murphy, R. B.; Repasky, M. P.; Frye, L. L.; Greenwood, J. R.; Halgren, T. A.; Sanschagrin, P. C.; Mainz, D. T. Extra Precision Glide: Docking and Scoring Incorporating a Model of Hydrophobic Enclosure for Protein-Ligand Complexes. J. Med. Chem. 2006, 49 (21), 6177– 6196. https://doi.org/10.1021/jm0512560.
- (54) Schrödinger Release 2019-4: LigPrep, Schrödinger, LLC, New York, NY, 2019.

- (55) Alex, A.; Millan, D. S.; Perez, M.; Wakenhut, F.; Whitlock, G. A. Intramolecular Hydrogen Bonding to Improve Membrane Permeability and Absorption in beyond Rule of Five Chemical Space. *MedChemComm* **2011**, *2* (7), 669–674. https://doi.org/10.1039/C1MD00093D.
- Patel, S.; Cohen, F.; Dean, B. J.; De La Torre, K.; Deshmukh, G.; Estrada, A. A.; Ghosh, A. S.;
  Gibbons, P.; Gustafson, A.; Huestis, M. P.; Le Pichon, C. E.; Lin, H.; Liu, W.; Liu, X.; Liu, Y.; Ly, C.
  Q.; Lyssikatos, J. P.; Ma, C.; Scearce-Levie, K.; Shin, Y. G.; Solanoy, H.; Stark, K. L.; Wang, J.; Wang,
  B.; Zhao, X.; Lewcock, J. W.; Siu, M. Discovery of Dual Leucine Zipper Kinase (DLK, MAP3K12)
  Inhibitors with Activity in Neurodegeneration Models. *J. Med. Chem.* 2015, *58* (1), 401–418.
  https://doi.org/10.1021/jm5013984.
- Patel, S.; Harris, S. F.; Gibbons, P.; Deshmukh, G.; Gustafson, A.; Kellar, T.; Lin, H.; Liu, X.; Liu, Y.; Liu, Y.; Ma, C.; Scearce-Levie, K.; Ghosh, A. S.; Shin, Y. G.; Solanoy, H.; Wang, J.; Wang, B.; Yin, J.; Siu, M.; Lewcock, J. W. Scaffold-Hopping and Structure-Based Discovery of Potent, Selective, And Brain Penetrant N-(1H-Pyrazol-3-YI)Pyridin-2-Amine Inhibitors of Dual Leucine Zipper Kinase (DLK, MAP3K12). *J. Med. Chem.* 2015, *58* (20), 8182–8199. https://doi.org/10.1021/acs.jmedchem.5b01072.
- (58) Moving beyond Rules: The Development of a Central Nervous System Multiparameter Optimization (CNS MPO) Approach To Enable Alignment of Druglike Properties | ACS Chemical Neuroscience https://pubs.acs.org/doi/pdf/10.1021/cn100008c (accessed 2022 -01 -31).
- Wang, Z.; He, N.; Guo, Z.; Niu, C.; Song, T.; Guo, Y.; Cao, K.; Wang, A.; Zhu, J.; Zhang, X.; Zhang, Z. Proteolysis Targeting Chimeras for the Selective Degradation of Mcl-1/Bcl-2 Derived from Nonselective Target Binding Ligands. J. Med. Chem. 2019, 62 (17), 8152–8163. https://doi.org/10.1021/acs.jmedchem.9b00919.