Integration of property predictions with molecule generation using reinforcement learning for fragment-based drug design

Abstract
The project centres around ideas of fragment-based drug design trying to address challenges of going from X-ray fragment hit to potential drug-like molecule. It will combine experimental X-ray fragment data, with generative diffusion models, affinity predictions models and reinforcement learning. This will enable the design of a machine learning model capable of generating molecules from X-ray fragment data that are optimised towards desirable properties such as good binding affinity or low toxicity for a given protein targe. The goal will be to develop explore effective reward strategies combined with generative diffusion models for the goal of reliable molecule generation.

Introduction. Fragment-based drug design (FBDD) is a method used in drug discovery that involves screening molecules from a fragment library against a chosen biological target. A fragment library is comprised of motives that are typically found in drug-like molecules but are smaller than a typical drug**.** Once a fragment is identified as binding to the biological target, typically a protein, it can be optimized into a more potent drug-like compound through growing, linking, or merging of fragments [1]. Generative modelling, based on e.g., diffusion models such as MiDi [2] has shown promise in generating new drug-like molecules and can be refined based on fragment data through strategies such as, e.g. SILVR [3]. However, these models only generate molecules *randomly* without any desired properties and do not hold any information about how well they may interact with a target protein molecule. Over the last few years, the XChem beamline (I04-1) at the Diamond Light Source has been collecting large sets of structural data on >250 different protein targets. Being able to utilise the fragment data efficiently, to explore new drug-like molecules quickly is important in fast design, test, make cycles in a drug discovery project based on a fragment screen. The Mey group has been working on refinement strategies for diffusion models, as well as machine-learning based strategies for binding affinity predictions between drug-like molecules and protein targets [4]. Machine learning methods lend themselves well, to navigate 1000s of fragment data efficiently, provide binding affinity estimates and elaborate ideas on how to grow, link or merge fragments. Combining the expertise on generative models and fragment data from the Mey group with approaches for noisy label detection [5] and reinforcement learning [6] developed in the Storkey group will allow for new ways in which to explore generative models for X-ray fragment data.

Research Challenge. Getting from fragment hits across multiple pockets on a protein target to a drug-like candidate requires data processing and modelling strategies that yield reliable computational affinities. This project will combine state-of-the-art fragment-merging and linking algorithms with machine learning and simulation-based affinity prediction methods. The main challenge is that even with fragment hits the search space of going from fragment to drug-like molecule remains prohibitively large to be explored automatically with existing computational techniques. Here we will integrate molecule generation from fragments with affinity predictions using reinforcement learning strategies. Reinforcement learning is an

iterative machine learning technique to effectively explore large landscapes using reward strategies. The idea is to use reinforcement learning to generated molecules from fragments with desirable properties. These properties can be generated either from simulation-based methods [7], from machine learning models [5,8], or from the integrated biological assay data at XChem. This new pipeline will enable going from fragment to a high potency drug-candidates in an automated fashion. Figure 1 summarises the ideas, starting out with many X-ray fragment hits on the left. The goal is to then generate new molecules with desirable properties (e.g., good binding affinity) using a reinforcement approaches. Ultimately validating the properties of the newly designed molecules again the protein target retrospectively and potentially prospectively through experiments at XChem.

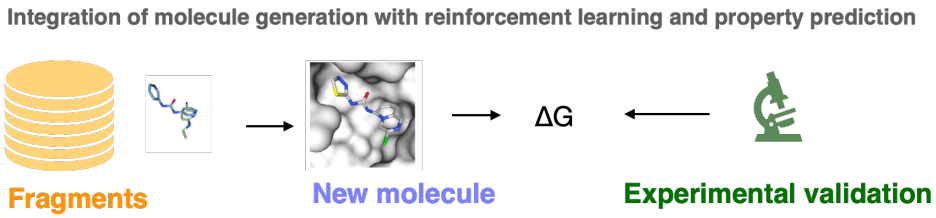

Figure 1: Fragments can be elaborated to drug molecules using reinforcement learning strategies to explore space that optimises a desired reward (e.g. binding affinity)

Data and methodology

The data required for the project is twofold. (1) Training data for generative models such as SPICE or GEOM-Drug have shown good promise at reliably training diffusion model based generative models. (2) Fragment data and affinity data used for exploring reinforcement learning strategies will come from existing datasets generated by for example the AI-driven Structure-enabled Antiviral Platform (ASAP) [9], as well as the XChem beamline directly. The first milestone will address combining generative model with an affinity model to be able to reliably score protein-ligand complexes. This will allow benchmarking of combined molecule generation and affinity predictions making sure the underlying models used in reinforcement learning strategies will be reliable. The second milestone will integrate reinforcement learning strategies with molecule generation, building on existing benchmarks in the Mey and Storkey research groups. Different strategies of varying complexity will be explored for defining rewards from both simulation-based an ML approaches. Once these two components have been integrated well, retrospective validation on ASAP data will be essential as well as prospectively using the tools on fragment screens from XChem.

Responsible AI and Ethics Considerations

Ethical concerns will touch multiple areas of the project and adequate solutions will be put in place to address these between the supervisory team and the student. Generative models can be used to design harmful compounds, such as toxic chemicals or biological agents. Mitigating towards dual-use of these models such that they cannot easily exploit the design of harmful agents will be essential. By implementing dual-use oversight measures, transparency in model access, and auditing usage we hope to address this risk. This will also encompass any reinforcement learning processes used and incorporating these with generative

models. All experimental data used in the project will have undergone any previous ethical approval for using XChem facilities infrastructure to run the experiments.

Expected output and impact will be: (1) An ML model to combine property predictions with molecule generation. (2) Integration of reinforcement learning strategies for molecule generation optimising different properties based on X-ray fragment data. (3) Validation of proposed protocols from 1 and 2 on biological targets investigated at the Diamond Lightsource XChem fragment data. If (3) is successful one route to impact will be for these methods to be deployed into the analysis suite of the XChem beamline which will be made available to all beamline users. Beyond XChem, all models will be made available broadly for anyone to be able to use.

[1] Scott et al. *Biochemistry* 51, 25, 4990-5003 (2021) [2] Vignac et al. arXiv:2302.09048 **[3]** Runcie and Mey, *J. of Chem. Inf Model.* 63.19, 5996-6005 (2023) **[4]** Gorantla et al. *J. Chem. Inf. Model.* 64, 1955-1965 (2024) **[5]** Jelley et al. arXiv:2406.13376 **[6]** Toner and Storkey https://doi.org/10.48550/arXiv.2409.06830 **[7]** Mey et al. *Living J. Comp. Mol. Sci*. 2,18378 (2020**) [8]** Gorantla et al. *J. Chem. Inf. Model.* 64, 2496-2507 (2024) **[9]** https://asapdiscovery.org