Novel Reinforcement Learning Framework for Synthesizable Drug Optimization

Tian, Aaron (School: Massachusetts Academy of Math and Science at WPI)

Synthesizability—the ability to build a compound using purchasable starting materials—is a vital consideration in computer-aided drug discovery; virtually created molecules are of no value if they are unsynthesizable. Although past research has found deep reinforcement learning to be effective at optimizing the drug-likeness of lead molecules, it was later discovered that the generated candidates could not be synthesized. The objective of this project was to develop a novel computational framework for lead optimization that generates synthesizable drug compounds. The framework, SynthNet, generates optimized drug candidates by iteratively performing chemical reactions on a set of purchasable reactants: an implicit guarantee of synthesizability. A deep reinforcement learning model was then implemented for reactant selection and trained with the soft actor-critic algorithm. SynthNet was assessed on its ability to optimize the quantitative estimate of drug-likeness (QED, a heuristic score of pharmacokinetic properties), and lipophilicity (clogP) of lead molecules. While the model performed marginally lower on the optimization objectives compared to the state-of-the-art, the guarantee of synthesizability is a substantial advantage. Hence, the generation technique used in this project greatly enhances the practicality of computational lead optimization models. Further research on improving the model's ability to perform multi-objective optimization, as well as the integration of more interpretable PK properties in the optimization objective, will enable the framework to supply a diverse pipeline of drug candidates that are ready for in-vitro and in-vivo testing.