Molecular Generation via Deep Learning

Corporate partner mentors from Merck: Xiang Yu, Peter Zhang – Data Mine student: Maria Berardi

PURDUE UNIVERSITY | The Data Mine

Project Outline

Motivation and Objectives

- Creating new molecules satisfying desired properties is a fundamental goal in drug discovery, hence Merck's interest in this project.
- Creating new chemical compounds is expensive in terms of time and cost, so researchers have found ways to automate part of the search process using Deep Learning.
- We examine two algorithms performing this task: Masked Graph Modeling (MGM) and Molecule Deep Q-Networks (MolDQN).
- We analyze their performance with the goal of proposing ways to combine the two approaches and suggest further developments.

Masked Graph Modeling (MGM)

An NLP inspired approach

- MGM is similar to masked language modeling in NLP, which consists of:
  1. Tokenizing text
  2. Masking a random selection of tokens
  3. Training the model to fill in the blanks
- Instead of tokenizing text, we "tokenize" a molecule into atoms and bonds, treated as components of a graph: nodes and edges.
- Graph components are iteratively masked and the model is trained to recover them.

Advantaged of MGN

- Representing a molecule as a graph allows for the use of graph specific features, such as distance between atoms, and other 3D features, which are not readily available using string representations as many other models do.
- A graph is sampled by repeatedly updating its components, rather than modeling the entire situation first.

Figure 1: Example of a possible generation trajectory, taken from [1], carried out using a 10% masking rate. Here you see a non-novel molecule (on the left, at step 0) and two novel molecules (at step 50 and 300 respectively).

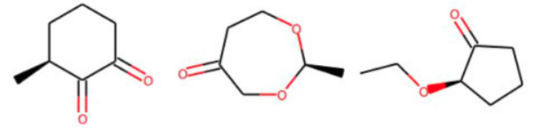

Molecule Deep Q-Networks (MolDQN)

Phrasing molecule generation as an MDP

- An MDP (Markov Decision Process) provides a mathematical framework to model a situation involving sequential decision making.
- An MDP consists of a set of *states, actions* and *rewards*. Actions can lead from one state to another. We assign a numerical reward to an action depending on its result.
- The MolDQN algorithm is based on an MDP where:
  - molecules are states;
  - allowed actions are atom addition, bond addition or bond removal;
  - reward is given based on one or more chemical properties that the molecule should satisfy at the end of the modification process.
- The model is trained to choose the actions that maximize the cumulative reward.

Figure 2: Valid actions on the state of cyclohexane.

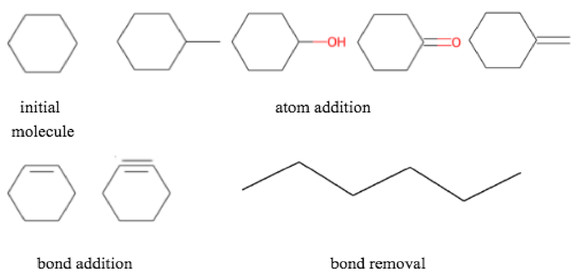

Advantages of MolDQN

- Only chemically valid actions are part of the action space, ensuring that all generated molecules are valid.
- The model is not pretrained on a data set, but starts from scratch and learns from experience, hence not inheriting potential bias from training data.
- There is potential for multi-objective optimization, if we need to train a model to optimize more than one molecular characteristic.

Training the MolDQN model on a single characteristic

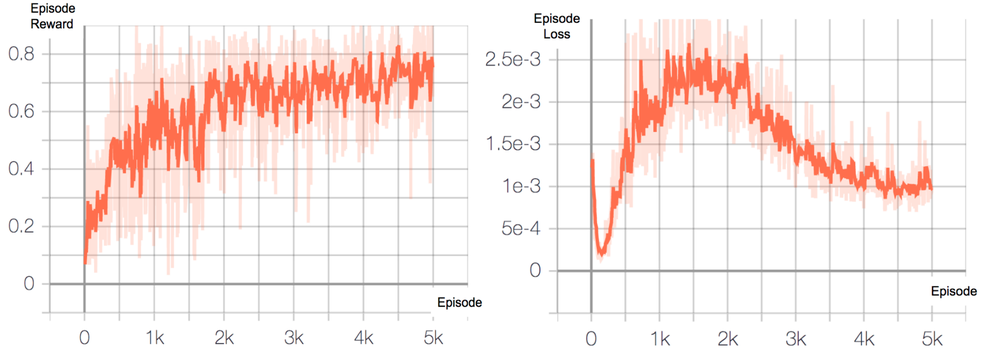

Figure 3: Episode reward (on the left) and episode loss (on the right) training the model for 5000 episodes on the single property QED (quantitative estimate of drug-likeness).

Possible Future Developments

- The search for a complete and nonredundant approach to the exploration of chemical space continues.
- We would like to be able to combine the two approaches examined this semester. In particular, we hope that using a graph representation of molecules in the context of a Q-learning algorithm would allow us to exploit the efficient learning method of deep Q-learning while being able to record more molecular features thanks to the graph structure.

Acknowledgements

I would like to thank my mentors, Xiang Yu and Peter Zhang, for all their help and guidance, and the Data Mine staff for their support.

References

1. Mahmood O, Mansimov E, Bonneau R, Cho K. Masked graph modeling for molecule generation. Nat Commun. 2021 May 26;12(1):3156. doi: 10.1038/s41467-021-23415-2. PMID: 34039973; PMCID: PMC8155025.
2. Zhou, Z., Kearnes, S., Li, L. *et al.* Optimization of Molecules via Deep Reinforcement Learning. *Sci Rep 9*, 10752 (2019). https://doi.org/10.1038/s41598-019-47148-x

The Data Mine Corporate Partners Symposium 2022