CHEMICAL STRUCTURE ELUCIDATION FROM TANDEM MASS SPECTROMETRY AND CHEMICAL FORMULA

A novel attempt at a Reinforcement Learning approach to De Novo Chemical Structure Elucidation from Mass Spectrometry and Chemical Formula.

Introduction

De Novo Chemical Structure Elucidation has the following **difficulties**:

- Mass spectrum cannot easily be **represented** (due to noise)
- Low amount (~60k samples) of **training data** (for a generative task)
- Generation of **invalid molecules** or with **incorrect chemical formula**
- Difficult to capture **permutation invariant** nature (GNN) while still allowing generation of molecules with **all sizes** (sequential).

In this investigation, we attempt to resolve all of the aforementioned difficulties, using a novel reinforcement learning approach.

Materials and Methods

To represent the mass spectrum, we perform **graph convolution on fragmentation trees**[1], which has not been done before, to the authors' knowledge.
For the generation, we will use a variation of Q-learning that only approximates **value functions** of each state given a goal (fragmentation tree embeddings), reducing this task to a **regressive** task, and allowing generation at **all sizes**.

We first **pretrain** a Fragmentation Tree GCN (FTree GCN) with a predictor to predict the original mass spectrum, so that the graph-level embedding **retains information**. Then, we used anomaly detection methods[2] to train the model on all states (exponential with respect to number of heavy atoms) that lead to the target to be considered 0, and others 1.

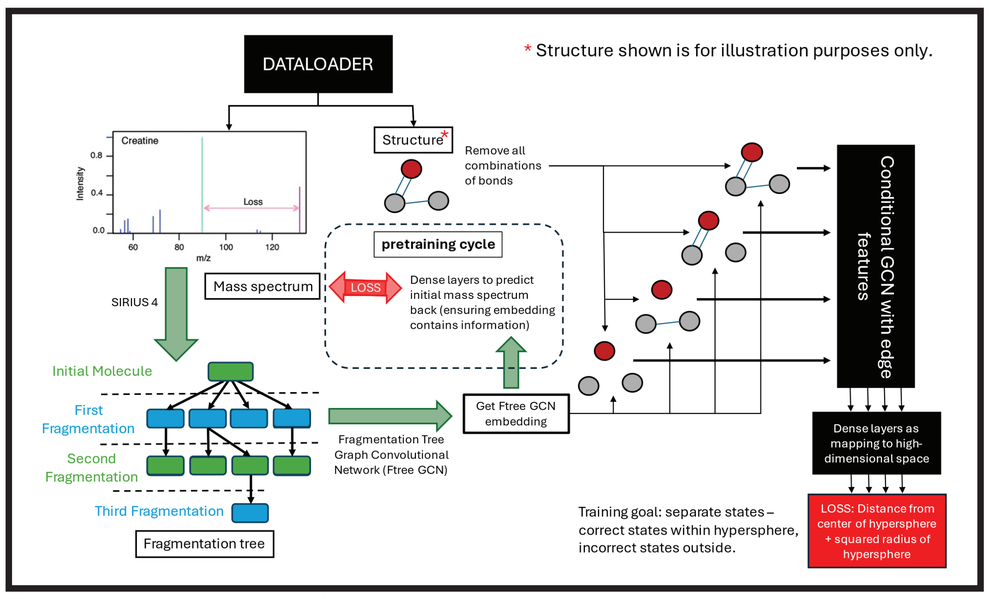

Figure 1. a visualization of the training process of the model, aiming to learn to tell between correct and incorrect states given only positive samples, using anomaly detection[2] methods. The fragmentation tree is obtained using the SIRIUS 4 software[3].

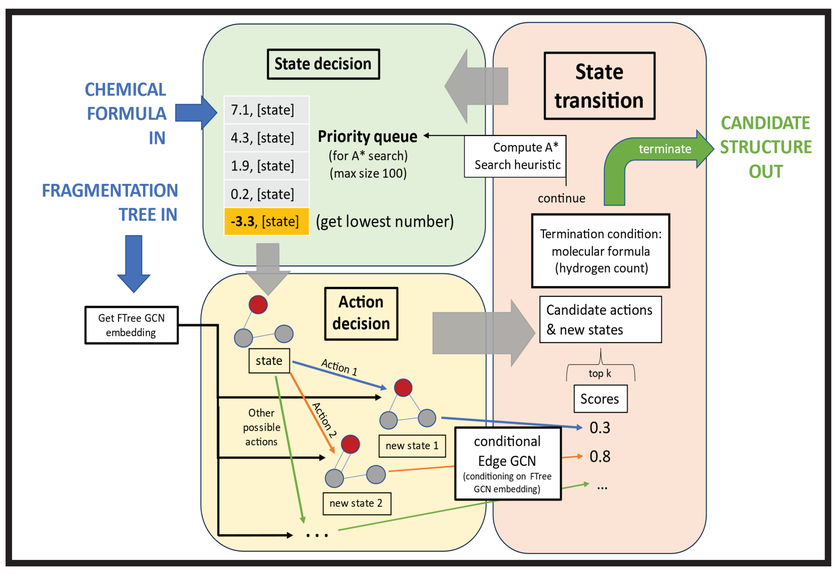

Figure 2. A visualization of the generation process used by our model, using A* Search. The inputs are the mass spectrum and chemical formula. The mass spectrum is preprocessed into an embedding at the left of the figure, used as the conditioning factor in the RL cycle at the right, that can generate multiple candidate structures.

The **heuristic** mentioned in figure 2 is for the A* search algorithm, defined by:

$$(1 - score_{anomaly}) * (H_t/2) - H_{left}$$

Where $score_{anomaly}$ is the anomaly score by the model, H_t is the total number of hydrogens in the initial state subtracted by the target state, and H_{left} is the number of hydrogens in the current state minus the target state.

The number in the priority queue shown is the negative of the heuristic, as Python's priority queue implementation takes the smallest value.

Results and Discussion

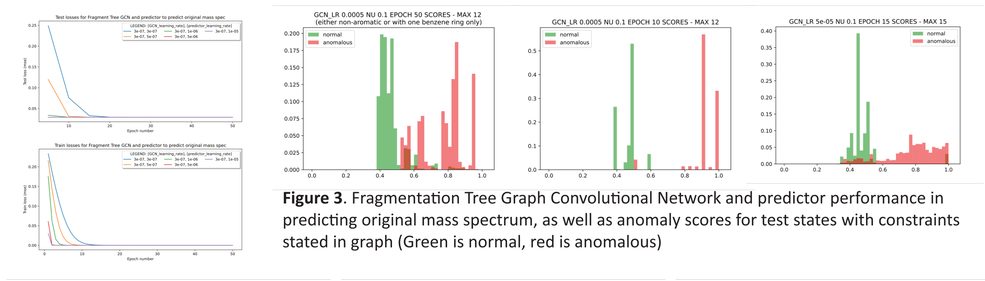

Figure 3. Fragmentation Tree Graph Convolutional Network and predictor performance in predicting original mass spectrum, as well as anomaly scores for test states with constraints stated in graph (Green is normal, red is anomalous)

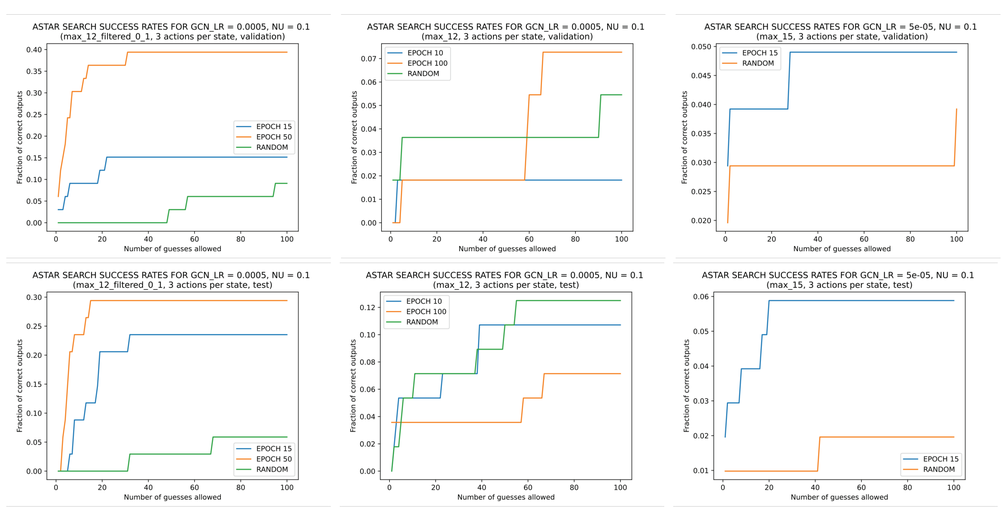

Figure 4. A* search success rates in different scenarios, starting from initial state with only aromatic bonds, taking up to 3 actions per state.

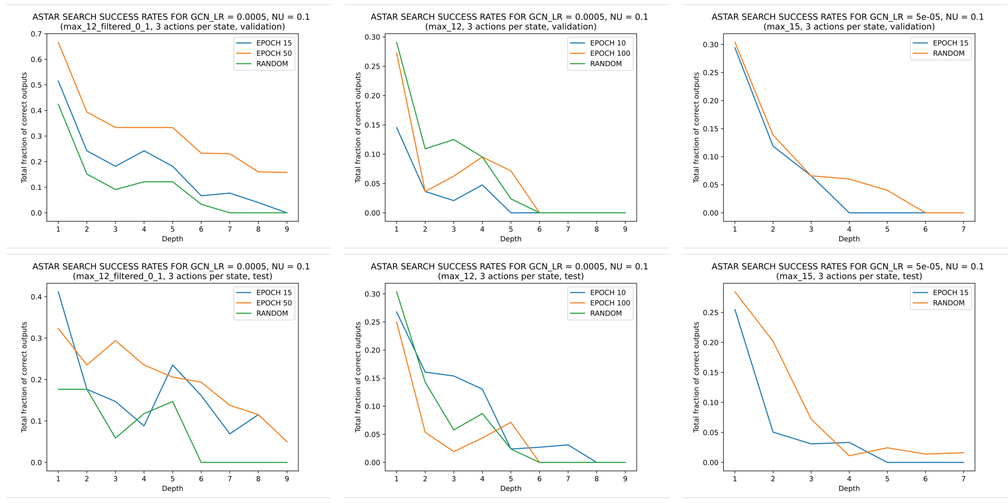

Figure 5. A* search success rates in different scenarios, starting [depth] bonds away from the target state, taking up to 3 actions per state. The green line is the performance when taking random actions.

We note, from figure 3, that the model did succeed in separating states.

From figure 5, we can see that, from depths of around 5, our models start to fail to succeed. However, in figure 4, we see that starting from the maximum depth, our model does achieve nonzero accuracy.
It appears that 2 overlapping factors cause the training to only be effective with extreme values for depths:

- The **low amount of training data** on states at **low depths**, causing training to be **ineffective with lower depths**, only relevant at higher depths.
- There being **few restrictions on the actions** that can be taken at **low depths** as compared to **high depths**, improving performance only with lower depths.

Due to these, perhaps the performance will be lower with an intermediate depth, while higher with depths closer to 1 or the maximum.

Future Work

- Weigh training cases to cause effective learning of states at low depths.
- Determine the reason why taking random actions achieves comparable results to trained models (figures 4 and 5) despite the models succeeding in separation (figure 3), as well as why random actions did better without restricting the molecules to non-aromatic or benzene ring molecules only.

References

[1]Böcker S, Rasche F (2008) Towards de novo identification of metabolites by analyzing tandem mass spectra. Bioinformatics 24:I49–I55
[2]Wang, X., Jin, B., Du, Y., Cui, P., Yang, Y.. (2021) *One-Class Graph Neural Networks for Anomaly Detection in Attributed Networks*. Neural Computing & Applications, volume 33, pages 12073–12085
[3]Dührkop, K., Shen, H., Meusel, M., Rousu, J. & Böcker, S. (2015) *Searching molecular structure databases with tandem mass spectra using CSI:FingerID*. Proc. Natl Acad. Sci. USA 112, 12580–12585.

Member:
Yu Simu, NUS High School
Mentor:
Dr Shen Bingquan, DSO National Laboratories

YDSP Young Defence Scientists Programme
DSTA Defence Science & Technology Agency
DSO NATIONAL LABORATORIES