




Recurrent Graph Transformers for Algorithmic Problems

Previously, recurrent Graph Neural Networks (GNNs) were shown to be able to learn simplified algorithms on graphs. In this work, we want to investigate the upcoming Graph Transformer architecture and its ability to learn more complicated algorithms. Graph Transformers generalize the transformer architecture known for Natural Language Processing (NLP) from sequences to general graphs. In contrast to message-passing GNNs that exchange messages solely between neighboring nodes, these Graph Transformers employ a (global) attention mechanism that lets a node attend to any other node in the graph. This can potentially reduce the number of required layers for some of our tasks/algorithms and we want to answer the question whether this also works in practice.

Similar to sequence transformers, positional encodings are used to give a sense of locality to nodes. We want to compare existing encodings and develop new ones for our tasks. Lastly, some tasks will still require a number of layers that scales with the size of the graph at hand, so we want to extend existing Graph Transformers with a fixed number of layers to a recurrent design.

Besides the aforementioned practical goals of this thesis, we will also study the theoretical capabilities that come with these architectures. This can mean for example that we inquire into the expressivity of found positional encodings and compare them to what is necessary for certain algorithms. Generally, we strive for an architecture that is well motivated in a theoretical sense.



Requirements: Strong motivation, knowledge in graph theory and machine learning, as well as good coding skills. Prior experience with GNNs or Machine Learning is a big advantage. We will have weekly meetings to discuss open questions and determine the next steps.

Interested? Please contact us for more details!

Contact

- Florian Grötschla: fgroetschla@ethz.ch, ETZ G63
- Joël Mathys: jmathys@ethz.ch, ETZ G63