




Benchmarking Graph Neural Networks

Neural networks have become the go-to approach for learning from many different types of data, be it text, images, medical scans or audio. Recently, they have also been designed and applied with increasing success to graph data. Graph neural networks have achieved great success in knowledge graphs, reinforcement learning, chemistry, or physics simulations. These models are mostly based on the message passing framework, where nodes exchange messages over the edges and update their states based on all of the aggregated incoming messages. There is a whole menagerie of different models, each with different in-built biases, strengths and weaknesses. However currently there is a lack of a set of rigorous, well-understood benchmark tasks of increasing complexity, which would allow for a comprehensive comparison of the models.



In this thesis we aim to introduce a set of controlled, synthetic, graph-based tasks of increasing complexity that can help tease apart different GNN architectures. In addition we aim to better understand the current roster of benchmark tasks. Finally we will carry out an extensive comparison of current state-of-the-art GNN architectures.

Requirements: Strong motivation, knowledge in deep learning, or a solid background in machine learning. Experience with Python and TensorFlow or PyTorch is an advantage as well as knowledge in graph theory, distributed computing and graph neural networks. We will have weekly meetings to address questions, discuss progress and think about future ideas.

Interested? Please contact us for more details!

Contact

- Béni Egressy: begressy@ethz.ch, ETZ G94
- Florian Grötschla: fgroetschla@ethz.ch, ETZ G61.3