




Beyond Homophily in Graph Neural Networks

Graph neural networks (GNNs) bring deep learning to structured data. Their example applications are predicting properties of molecules or understanding social networks better. Many real-world graphs have a property which is called homophily – that means that most nodes are similar to their neighbors.

As of now, most GNN models were trained on benchmarks with high homophily. Therefore, their performance on datasets without this property is rather poor. The goal of this project is to improve GNN models on datasets with low homophily in a way that retains good performance on traditional benchmarks with high homophily. There is also an opportunity to investigate ways to improve GNN performance on heterogeneous graphs, that have different kinds of vertices (i.e. authors, institutions, publications).



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

- Karolis Martinkus: martinkus@ethz.ch, ETZ G95
- Lukas Faber: lfaber@ethz.ch, ETZ G60.1