An Introduction to Graph Neural Networks

Alina Lazar
Youngstown State University
Youngstown, OH, USA
alazar@ysu.edu

Abstract

Graph Neural Networks (GNNs) are considered a subset of deep learning methods designed to extract important information and make useful predictions on graph representations. Researchers have been working to adapt neural networks to operate on graph data for more than a decade. Most practical applications come from the areas of physics simulations, object detection and recommendation systems. Given the extended application areas, GNNs are one of fastest growing and most active research topic, that attracts increasing attention not only from the machine learning and data science community, but from the larger scientific community as well. The materials for this tutorial will be selected and organized for researchers with no prior knowledge of GNNs. Further reading, applications and most popular software packages and frameworks will also be discussed.

Tutorial Description

Graph structured data exist everywhere in the real world. Almost any problem can be modeled using graph representations. From social networks to molecular structures and particle tracking, the range of practical applications is vast. In this context, it is becoming important to design and evaluate advanced learning methods on graph structured data. GNNs [2, 3], that extend the well-known deep neural network models to graph representations, offer researchers a new way to learn graph representations at the node, edge, and graph levels. For each of those levels, different challenges could be faced, therefore specific algorithms must be designed:

1. Node Classification: This task relates to node classification or regression. Usually implemented using Graph Convolutional Networks (GCN) [1], these models predict the category of each node.
2. Link Prediction: This task relates to the link prediction. Given two nodes' internal representations in a GNN, the model can be utilized to predict how likely is that two nodes are going to be connected in the future.
3. Graph Classification: This task relates to the graph-level characteristics of a graph. In this case pooling operations are used to capture a compact representation at the graph level.

This tutorial will cover relevant GNN-related topics, including the basics of learning on graph structured data, graph embeddings, attention networks, aggregation functions and examples of applications (node classification, predicting missing links, detecting communities, and graph matching). For these applications, GNNs have achieved impressive performance on relatively small graph datasets. Unfortunately, most real-world problems rely on large graphs that do not fit into the available GPU memory of current hardware systems. We will also discuss ways to design, evaluate and scale GNN training and inference methods.

Outline

Part 1:

- Graphs and graph structured data (15 min)
- Node, edge and graph level tasks (10 min)
- Graph embeddings (10 min)
- Graph convolutional neural networks (GCNNs) (10 min)

Part 2:

- Graph Attention Networks (10 min)
- Spectral-based GNN layers (10 min)
- Aggregation Functions in GNNs (10 min)
- GNN Applications (10 min)
- GraphGym: design and evaluate GNNs (5 min)

References

[1] Kipf, T. N., and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

[2] Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2008. The graph neural network model. *IEEE transactions on neural networks* 20(1):61–80.

[3] Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Philip, S. Y. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32(1):4–24.