Course name: Graphs in Machine Learning

Course description:
Graphs are a ubiquitous data structure and employed extensively within computer science and related fields. A lot of real world applications can be readily modeled as graphs such as recommender systems, social networks, single molecular and protein interactions networks. Graphs are not only useful as structured knowledge repositories: they also play a very important role in modern machine learning. This course will cover both conventional algorithms and the most recent research on analysis of graphs from a machine learning perspective. The topics that we will cover include: basic graph algorithms such as graph clustering, summarization, anomaly detection and more advanced research topics such as network embedding, graphical neural networks and deep reinforcement learning on graphs.

Classes consist of both instructor and student presentations (see **Syllabus**). For the first part, the instructor will introduce the basic concepts, algorithms and computational tools. For the second part, students are expected to read and present research papers on most recent deep learning research on graphs and present their own projects.

Prerequisites:
Students are expected to have the following background:

- Basic programming skills to write a reasonably non-trivial computer program in Python or C (e.g., CS 38003 or equivalent are recommended).
- Undergraduate or graduate level machine learning courses (e.g., CS 37300 and CS 578 are sufficient).
- Students without this background should discuss their preparation with the instructor. The introduction session in the first week of the class will give an overview of the expected background.

Target Students
This course is targeted primarily to senior undergraduate and graduate students in computer science departments. It is also useful to graduate students who are interested in applying machine learning methods to their own research areas including Life Sciences, Health Sciences, Material Sciences and Social Sciences.

Optional textbooks
“Networks: An Introduction” by MarkNewman
“Python for Graph and Network Analysis” by Mohammed Zuhair, Al-Taie, and Seifedine Kadry
“Complex Network Analysis in Python” by Dmitry Zinoviev
“Deep Learning” by Ian Goodfellow, Yoshua Bengio, and Aaron Courville

Presentations (50% of grade)
Please make an appointment with the instructor one week before your presentation. Besides the models and algorithms of the paper, please clearly state what is the role of 'GRAPHS' in the work you present. If you want to change your presentation date, please arrange a swap with another student and notify me at least two weeks in advance.

Final projects (50% of grade)
Each student will pick a project related to graphs and machine learning. Here are some possible directions:

1. An interesting mathematical problems around a paper.
2. Adopting a developed computational framework on a new dataset.
3. Following-up experiments of an existing work to understand its important properties.
4. Simple extension of an existing machine learning model, such as unsigned network to signed network, trees to DAGs, shallow neural networks to deep neural networks.
5. Using graphs to model a problem in your own research area.

Please turn in a two-page project proposal before Mar 20th 23:59pm. I will give feedback on the proposal as soon as possible. Please turn in a final report before May 1st 23:59pm.

Syllabus (tentative)

1. Introduction
2. Basic graph theory
3. Cytoscape and Python
4. Biological networks
5. Network motifs
6. Network alignment algorithms
7. Random Walks and PageRank
8. Graph clustering and communities
9. Network regularization
10. Graph Similarity
11. Graph Summarization
12. Anomaly Detection on Graphs
13. Graphical Model I
14. Graphical Model II
15. Knowledge graph
16. Large-scale Graphs (student presentations)
17. Basic concepts and techniques of Deep Learning I
18. TensorFlow and Pytorch
19. Network embedding I (student presentations)
20. Network embedding II (student presentations)
21. Knowledge graph embedding (student presentations)
22. Graphical Neural Networks I (student presentations)
23. Graphical Neural Networks II (student presentations)

24. Graphical Neural Networks III (student presentations)
25. Graphical Neural Networks III (student presentations)
26. Deep reinforcement learning
27. Deep reinforcement learning on Graphs I (student presentations)
28. Deep reinforcement learning on Graphs II (student presentations)
29. Project presentations
30. Project presentations
31. Project presentations
32. Project presentations