

Learning Deep Generative Models of Graphs

Monday 17 September 2018 11:05 (30 minutes)

Abstract: Graphs are fundamental data structures which concisely capture the relational structure in many important real-world domains, such as knowledge graphs, physical and social interactions, language, and chemistry. Here we introduce a powerful new approach for learning generative models over graphs, which can capture both their structure and attributes. Our approach uses graph neural networks to express probabilistic dependencies among a graph's nodes and edges, and can, in principle, learn distributions over any arbitrary graph. In a series of experiments our results show that once trained, our models can generate good quality samples of both synthetic graphs as well as real molecular graphs, both unconditionally and conditioned on data. Compared to baselines that do not use graph-structured representations, our models often perform far better. We also explore key challenges of learning generative models of graphs, such as how to handle symmetries and ordering of elements during the graph generation process, and offer possible solutions. Our work is the first general approach for learning generative models over arbitrary graphs, and opens new directions for moving away from restrictions of vector- and sequence-like knowledge representations, toward more expressive and flexible relational data structures.

Bio: Yujia Li is a research scientist at DeepMind. He obtained his Ph.D. at University of Toronto, studying machine learning and deep neural networks. His current research focuses on developing structured neural models and solving problems on structured data. He also has experience in other aspects of machine learning, including structured prediction models, generative models, semi-supervised learning and application domains including computer vision and temporal signal processing.

Presenter: Dr LI, Yujia (DeepMind)

Session Classification: Day 1