

Verification of Graph Learning Models and Analysis of Query Languages

Michael Benedikt¹

¹*University of Oxford, United Kingdom*

Abstract

Graph neural networks (GNNs) are the predominant architectures for a variety of learning tasks on graphs. Like other modern machine learning models, their success is accompanied by many fundamental concerns. Can we understand what exactly a trained model does? Can we have confidence that a trained model will not do anything bad? In this talk we will give one of the first approaches to verifying GNNs. We will work using the static analysis of techniques of database theory. Many of our results work by converting the GNN into a logic, and then showing that we can analyze the logic. This translation is useful not only for verification, but for explanation of the behavior of GNNs.

The work reported here is joint with Chia-Hsuan Lu, Boris Motik, and Tony Tan (Decidability of Graph Neural Networks via Logical Characterizations), it is closely related to work of Barcelo, Kostylev, Monet, Reutter, and Silva from ICLR 2020 (The Logical Expressiveness of Graph Neural Networks), and also to prior work on analysis of logics on graphs with arithmetic (On two-variable guarded fragment logic with expressive local Presburger constraints; Two Variable Logic with Ultimately Periodic Counting). The talk will be in the same spirit as the AMW summer school course on analysis of GNNs via query languages: but there will be no dependency between the two, and no overlap in the results.

AMW 2024: 16th Alberto Mendelzon International Workshop on Foundations of Data Management, September 30th–October 4th, 2024, Mexico City, Mexico

✉ michael.benedikt@cs.ox.ac.uk (M. Benedikt)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).