Halting Recurrent GNNs and the Graded μ-Calculus

Jeroen Bollen

Data Science Institute, UHasselt, Belgium
jeroen.bollen@uhasselt.be

Graph Neural Networks (GNNs) are machine-learning models that operate on graph-structured data. Their expressive power is intimately related to logics that are invariant under graded bisimilarity. Current proposals for recurrent GNNs either assume that the graph size is given to the model, or suffer from a lack of termination guarantees. Here, we propose a halting mechanism for recurrent GNNs. We prove that our model can express all node classifiers definable in graded modal μ-calculus, even for the standard GNN variant that is oblivious to the graph size. To prove our main result, we develop a new approximate semantics for graded μ-calculus, which we believe to be of independent interest. We leverage this new semantics into a new model-checking algorithm, called the counting algorithm, which is oblivious to the graph size. In a final step we show that the counting algorithm can be implemented on a halting-classifier recurrent GNN.

This research is jointly done with Jan Van den Bussche (UHasselt), Stijn Vansummeren (UHasselt), and Jonni Virtema (University of Glassgow) and published in the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR 2025).