Graph Neural Networks: Everything is Connected

Matthias Fey

Kumo.AI, Inc, Mountain View, California, United States
matthias@kumo.ai

Abstract. In this talk, we introduce the concept of Graph Neural Networks and its general framework of neural message passing. We thoroughly analyze the expressive power of GNNs and show-case how they relate and generalize concepts of Convolutional Neural Networks and Transformers to arbitrarily structured data. In particular, we argue for the injection of structural and compositional inductive biases into deep learning models. Despite recent trends in neural networks regarding LLMs, such models manifest our understanding of a structured world, require less computational budget, and are easier to understand and explain.

Keywords: Graph Neural Networks · Deep Learning

Our world is highly rich in structure, composed of objects, their relations and hierarchies. Despite the ubiquity of graphs in our world, most modern machine learning methods fail to properly handle such rich structural representations. Recently, a universal class of neural networks emerged that can seamlessly operate on graph-structured data, summarized under the umbrella term *Graph Neural Networks (GNNs)*.

Graph Neural Networks generalize the concepts of Convolutional Neural Networks and Transformers by following a *neural message passing scheme* (see Fig. 1), in which the computation graph of the network is no longer fixed but instead part of the input. This makes GNNs broadly applicable to a wide range of applications, even on highly irregular input structures. In addition, it allows us to actually witness them as a new paradigm on how to define neural networks. Despite recent successes with LLMs on large-language corpuses, GNNs utilize their built-in structural inductive bias to derive accurate predictions even on only sparsely available data with low computational budget.

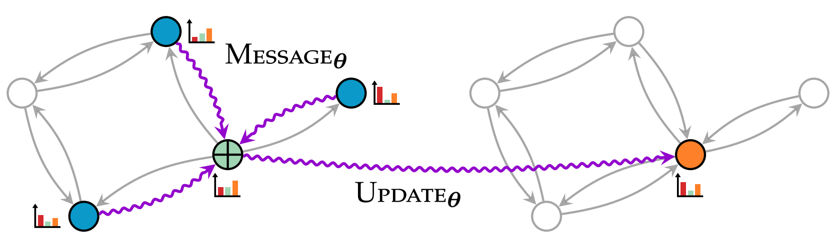

Fig. 1. Graph Neural Networks follow a neural message passing scheme.