Topology Matters: Understanding How Network Topological Features Impact Graph Neural Networks

Catherine Tong, Vivek Kothari and Nicholas Lane

To date, GNNs have been evaluated on a handful of benchmark graph datasets with limited variations in topological features; There has also yet been a systematic empirical study on how graph topology may be consequential in model performance.

We apply common GNNs on graphs from several domains for a transductive task of node classification. We analyze topological attributes of these graphs such as betweenness, centrality, and clustering and their impact on the performance of the GNN algorithms. Our experiments find strong and significant correlations (averaging $|r|=0.74$) between test performance and topological characteristics. We further observe statistically significant degradation of model performance on graphs with multiple components and isolated nodes, directed graphs and noise. Collectively we provide a series of empirical observations and insights that we expect to be of value in guiding future, topologically-conscient, developments in the area.

Our results show that network topology does matter and has a significant impact on the performance of Graph Neural Networks. However, we also find this relationship between topology and GNNs to be complex and nuanced.