LEARNING TO MINE APPROXIMATE NETWORK MOTIFS

Anonymous authors
Paper under double-blind review

ABSTRACT

Frequent and structurally related subgraphs, also known as network motifs, are valuable features of many graph datasets. However, the high computational complexity of motif mining has limited the scope of their use in many real-world datasets.
To facilitate the development of machine learning approaches aimed at motif mining, we propose a formulation of the motif mining problem as a node labeling task. In addition, we build benchmark datasets and evaluation metrics which test the ability of models to capture different aspects of motif discovery. Next, we propose `MotiFiesta`, a first attempt at solving this problem in a fully differentiable manner with promising results on challenging baselines and surpassing classical approaches. Finally, we demonstrate through `MotiFiesta` that this learning setting can be applied simultaneously to general-purpose data mining and interpretable feature extraction for graph classification tasks.