MOTIF-Driven Contrastive Learning of Graph Representations

Arjun Subramonian, UCLA, arjunsub@ucla.edu

AAAI-21 Undergraduate Consortium

Introduction

Graph representation learning aids numerous applications, like predicting properties of molecules. These applications are often low-resource because labels are expensive to acquire, which makes them amenable to self-supervised learning. There currently exist self-supervised methods for learning graph-level representations, including InfoGraph (Sun et al. 2019), GPT-GNN (Hu et al. 2020), and Context Prediction (Hu et al. 2020).

MOTIF

The aforementioned methods fail to leverage the commonality of significant subgraphs across graphs in a dataset, known as motifs, which often indicate graph-level semantic properties. We propose a self-supervised framework called MOTIF that simultaneously trains a graph neural network (GNN) to learn subgraph representations and clusters these representations based on their semantics to mine motifs (see **Figure 1**). Thereby, we pretrain the GNN for various graph-level downstream tasks, like molecular property prediction.

On each iteration of MOTIF, we obtain subgraph representations by sampling subgraphs and passing these subgraphs through the GNN. We subsequently compute an optimal soft assignment of each subgraph representation to clusters, whose centers represent motifs. Our goal is to optimize the cluster assignments by adjusting the cluster centers to maximize the similarity between subgraph representations and the centers. Then, the optimal cluster assignments provide us with pseudo-labels with which we can contrastively train the GNN.

MOTIF was ineffective when employing random walk or k-hop neighbor subgraph sampling, as these methods solely rely on topology without considering node representations learned by the GNN. Instead, to sample subgraphs from a graph, I computed the similarities between the representations of adjacent nodes and normalized them to produce an affinity matrix. I then applied spectral clustering with the affinity matrix to recursively bisect the graph into subgraphs of multiple sizes (see **Figure 2**).

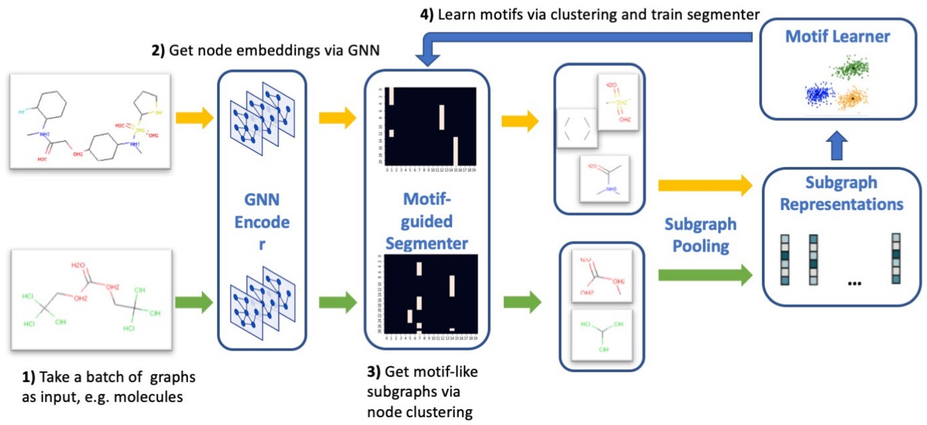

Figure 1. Illustration of MOTIF framework. MOTIF simultaneously trains a GNN to learn subgraph representations and clusters these representations based on their semantics to mine motifs. MOTIF also employs a dynamic, motif-guided segmenter that exploits node representations learned by the GNN to sample motif-like subgraphs.

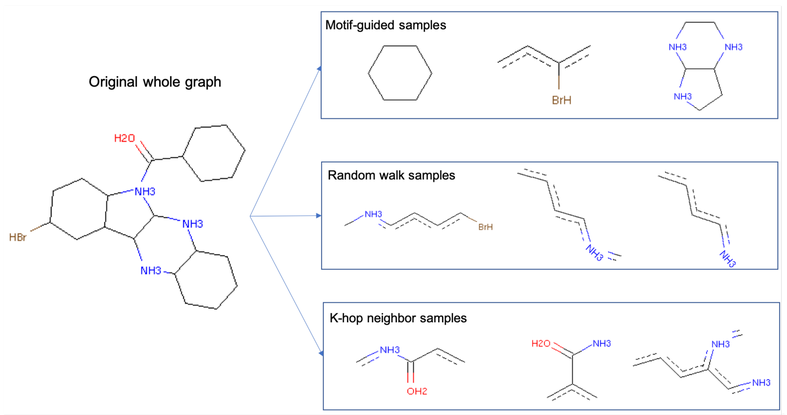

Figure 2. Motif-like subgraphs sampled by the dynamic, motif-guided segmenter vs. subgraphs sampled by random walk and k-hop neighbor samplers. As we train the GNN, the representations of nodes belonging to the same motif will become more similar, so the sampler will reinforce clustering and vice versa, thereby boosting MOTIF's performance.

Results and Discussion

We evaluated MOTIF on molecular property prediction datasets from the Open Graph Benchmark (OGB). We pretrained GNNs using MOTIF and state-of-the-art baselines on the large ogbg-molhiv dataset. Then, we emulated the scarcity of labels in the real world by finetuning the pretrained model on all the ogbg molecule datasets with only 10% of labels. The MOTIF-pretrained model consistently performed better than or on par with models pretrained with the baselines when transferring to a supervised finetune task with few labels. Furthermore, the MOTIF-pretrained model outperformed the non-pretrained model by 2.8% on average (see **Figure 3**).

These results suggest that MOTIF further enhances molecular property prediction with few labels, which can advance drug discovery and quantum chemistry. Furthermore, MOTIF is more interpretable, as one can inspect the motifs mined to intuit the graph representations learned. In the future, we will explore introducing a sampling loss term and end-to-end pooling for subgraph representations.

	bace	bbbp	clintox	hiv
Non-Pretrain	72.80 ± 2.12	82.13 ± 1.69	74.98 ± 3.59	73.38 ± 0.92
ContextPred	73.02 ± 2.59	80.94 ± 2.55	74.57 ± 3.05	73.85 ± 1.38
InfoGraph	76.09 ± 1.63	80.38 ± 1.19	**78.36 ± 4.04**	72.59 ± 0.97
GPT-GNN	75.56 ± 2.49	83.35 ± 1.70	74.84 ± 3.45	74.82 ± 0.99
MOTIF	**76.16 ± 2.51**	**83.78 ± 1.77**	77.50 ± 3.35	**75.51 ± 0.67**

	sider	tox21	toxcast	Average
Non-Pretrain	55.65 ± 1.35	76.10 ± 0.58	63.34 ± 0.75	71.19
ContextPred	54.15 ± 1.54	74.85 ± 1.28	63.19 ± 0.94	70.65 (-0.7%)
InfoGraph	56.88 ± 1.80	76.12 ± 1.11	64.40 ± 0.84	72.11 (+1.3%)
GPT-GNN	55.59 ± 1.58	76.34 ± 0.68	64.76 ± 0.62	72.18 (+1.4%)
MOTIF	**57.28 ± 1.09**	**76.68 ± 0.36**	**65.42 ± 0.62**	**73.19 (+2.8%)**

Figure 3. We evaluated transfer learning performance using ROC-AUC. For multi-task prediction, we averaged the ROC-AUC across all tasks. We report the test result of the best model on the validation set across 10 runs.

Acknowledgements

Research conducted in collaboration with Shichang Zhang, Ziniu Hu, and Yizhou Sun, all from the same institution.