

GNNs on Evolving Graphs: A Benchmark of Incremental Updates and Meta-Learning Approaches

Abstract

The field of graph machine learning has seen significant growth with the success of Graph Neural Networks (GNNs). However, most traditional GNNs are designed for static graphs. In the real world, graphs are constantly evolving—new users join social networks, molecules change shape, and data streams into a network. Re-computing a GNN's embeddings for the entire graph every time a small change occurs is computationally expensive and inefficient. This research explores two more efficient approaches: a standard incremental update method and a novel meta-learning approach, which are then benchmarked to compare their performance.

Introduction

- ▶ Traditional Graph Convolutional Networks (GCNs) require full-graph inference to update embeddings.
- ▶ This becomes inefficient for large, frequently changing graphs.
- ▶ Even small updates cause embeddings of connected nodes to become "stale."
- ▶ Naively re-running full inference does not scale well in practice.

Objectives

- ▶ Develop a method to maintain accurate, up-to-date embeddings for dynamic graphs.
- ▶ Minimize computational cost compared to full re-computation.
- ▶ Achieve the efficiency of incremental updates while preserving the accuracy of full-graph methods.

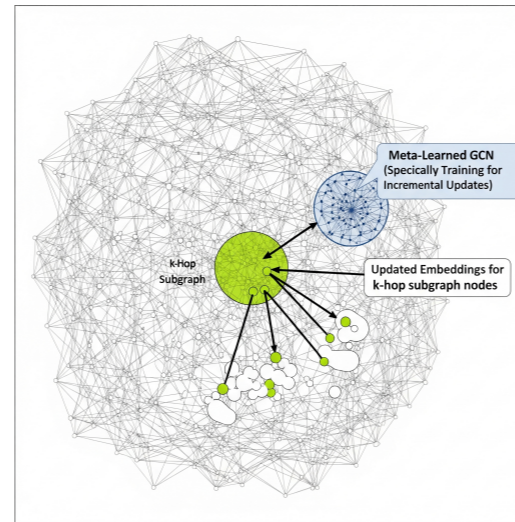
Related Work

- ▶ Early Graph Neural Networks (GNNs) were primarily designed for static graphs, limiting their application in dynamic, real-world settings.
- ▶ Several approaches have explored dynamic embeddings, focusing on re-computing or updating only parts of the graph.
- ▶ Recent work [Zhong et al.(2024)] proposes algorithms for efficient embedding computation on evolving graphs.
- ▶ Our research extends this direction by benchmarking a standard incremental update method against a novel meta-learning approach.

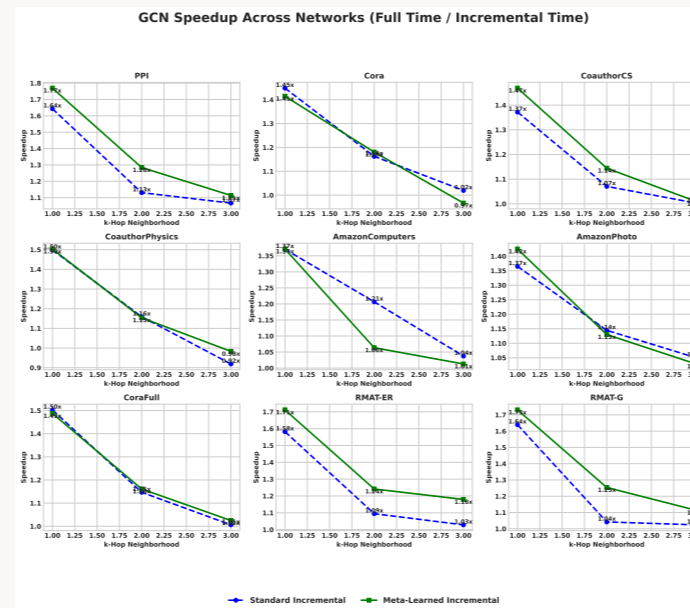
Methodology

The training process for the meta-learned model involves:

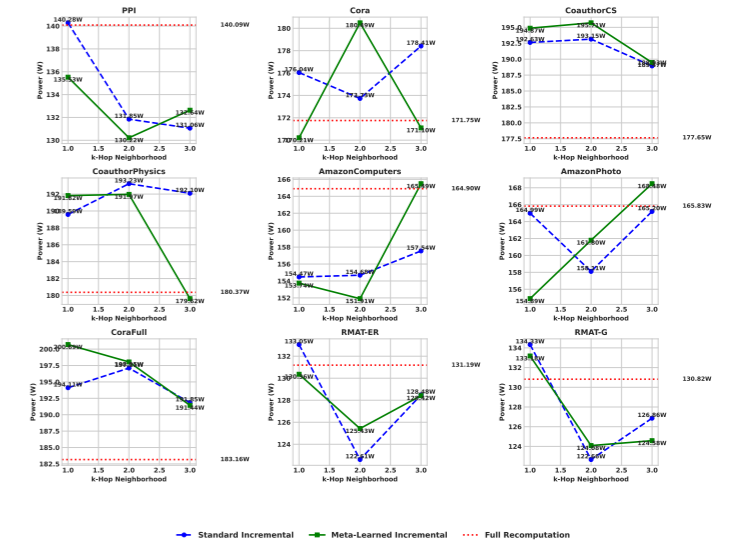
- ▶ Training the GNN to produce embeddings for a local subgraph that closely match those from the entire graph.
- ▶ Computing a set of "ground truth" embeddings for the full graph.
- ▶ Using these ground truth embeddings as the target for the meta-learner.



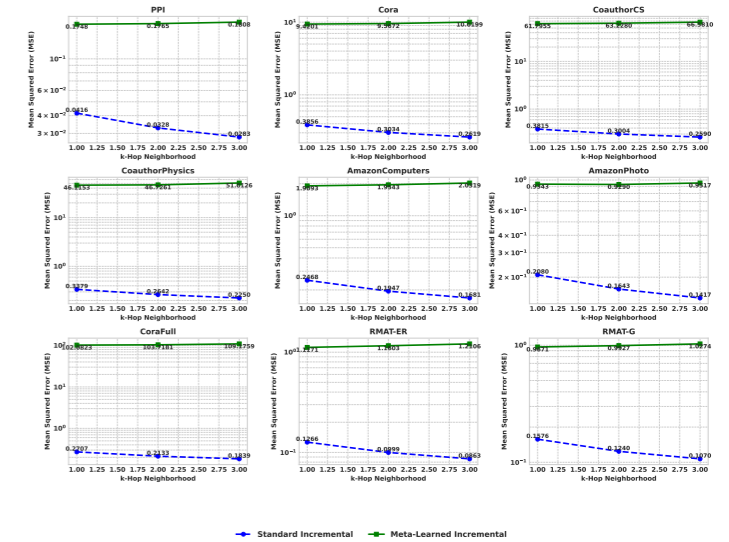
Results



GCN Power Consumption Across Networks



GCN Inference Accuracy (MSE) Across Networks



Reference

Yongjian Zhong, Hieu Vu, Tianbao Yang, and Bijaya Adhikari. 2024. Efficient and Effective Implicit Dynamic Graph Neural Network. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Barcelona, Spain) (KDD '24)*. Association for Computing Machinery, New York, NY, USA, 4595–4606. <https://doi.org/10.1145/3637528.3672026>

Conclusion

The Standard Incremental and Meta-Learned Incremental approaches are significantly faster and consume less power than the Full Re-computation approach. The Speedup plots across various networks further confirm that the incremental methods are more scalable.