Multi-View GNN Inference Efficiency Scaling with Graph Size and Latency

Assignee Research

June 2, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: How does the inference efficiency of multi-view GNNs scale with increasing graph size compared to single-view models when evaluated on benchmark datasets like ACM or DBLP, measured by latency per. 16 claims were extracted from source literature; 7 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Graph Neural Networks with Dynamic Temporal Encoding for Large-Scale Social Network Analytics. Research question: How does the inference efficiency of multi-view GNNs scale with increasing graph size compared to single-view models when evaluated on benchmark datasets like ACM or DBLP, measured by latency per query?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.7/10.

3 Results

4 papers retrieved. 16 claims extracted; 7 independently verified. Quality review score: 6.7/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
DyTEGNN is a framework that unifies continuous-time temporal embeddings with multi-relational graph convolutions.	✓	0.25
DyTEGNN incorporates a Hawkes-process-inspired temporal event encoder.	✓	0.22
DyTEGNN uses a heterogeneous message-passing scheme with edge-type-aware attention.	✓	0.21
DyTEGNN includes a hierarchical graph pooling module supporting million-node graphs through neighbor sampling.	✓	0.24
The framework was evaluated on three benchmark datasets: OGB-Twitter, Reddit-Threads, and DBLP-Citation.	✓	0.22
The framework was evaluated across four tasks: node classification, link prediction, community detection, and influence	✓	0.25
DyTEGNN achieved a macro F1-score of 91.3% on node classification on the OGB-Twitter dataset.	×	0.14
DyTEGNN achieved a macro F1-score of 89.7% on node classification on the Reddit-Threads dataset.	×	0.12
DyTEGNN achieved a macro F1-score of 88.4% on node classification on the DBLP-Citation dataset.	×	0.14
DyTEGNN outperformed seven baselines by up to 7.8% in node classification tasks.	×	0.11
DyTEGNN achieved an AUC of 0.974 on link prediction on the OGB-Twitter dataset.	✓	0.16
Ablation studies showed that temporal encoding contributed a +4.1% increase in F1 score.	×	0.13
Ablation studies showed that heterogeneous attention contributed a +3.2% increase in F1 score.	×	0.09
DyTEGNN scales sub-linearly with graph size.	×	0.13
DyTEGNN handled graphs with 50 million edges.	×	0.08
DyTEGNN achieved a 3.7x speedup over full-batch GNNs.	×	0.13

References

- https://www.semanticscholar.org/paper/e2992b8f491ef2f9b605c1cba5aabf81a1dcf246
- https://www.semanticscholar.org/paper/7a1771bff33729a0c6bf292d563653a5145c9b3f
- https://www.semanticscholar.org/paper/a0137db8f9858b64ec16f0ecc5076521c4010ccd