Scaling Efficiency of Hybrid Graph Neural Networks with Synthetic Augmentation and Contrastive Learning

Assignee Research

June 2, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does the computational efficiency of hybrid graph neural networks combining synthetic graph augmentation and contrastive learning scale with graph size compared to variational inference. We present a novel edge-level ego-network encoding for learning on graphs that can boost Message Passing Graph Neural Networks (MP-GNNs) by providing additional node and edge features or extending message-passing formats. The proposed encoding is sufficient to distinguish. 11 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Improving Subgraph-GNNs via Edge-Level Ego-Network Encodings. Research question: How does the computational efficiency of hybrid graph neural networks combining synthetic graph augmentation and contrastive learning scale with graph size compared to variational inference baselines, as measured by inference throughput and memory usage on OGBN-MAG and OGBN-Proteins benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

13 papers retrieved. 11 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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
Elene and Elene-L are evaluated in various graph-level settings, including expressivity, proximity, real-world graphs, a	$\times$	0.12
Elene-L (ED) achieves parity with Pure Graph Transformers when including edge-level signals.	$\times$	0.10
Elene denotes Eq. 2 as additional features, while Elene-L denotes the representations of Eq. 6 and Eq. 7.	$\times$	0.02
Experiments were conducted on a shared server with a 48GB Quadro RTX 8000 GPU, 40 CPU cores, and 502GB RAM.	$\times$	0.03
Each individual job has a limit of 96GB RAM and 8 CPU cores.	$\times$	0.00
Experiments were also reproduced on a SLURM cluster with nodes equipped with 22GB Quadro GPUs.	$\times$	0.02
Scalability experiments ran on Tesla T4 GPUs with 15.11GB of VRAM.	$\times$	0.02
Elene hyper-parameters were explored via grid search with k $\in$ {0, 1, 2, 3, 5} parameter ranges.	$\times$	0.04
GIN+ELENE-L (ED) achieves 100% accuracy on SR25 and 0.023 MAE on Counting Substructures (Triangles).	$\times$	0.01
GIN+ELENE-L (ND) achieves 100% accuracy on SR25 and 0.012 MAE on Counting Substructures (Triangles).	$\times$	0.01
Elene-L (ND) outperforms strong baselines from SPNNs and Graphormer on h-Proximity tasks.	$\times$	0.06

References

- `http://arxiv.org/abs/2505.15103v2`
- `http://arxiv.org/abs/2312.05905v2`
- `http://arxiv.org/abs/2206.07869v1`