Gradient Checkpointing in GNNs: Memory and Throughput Trade-offs Across Graph Densities

Assignee Research

June 1, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: What is the effect of gradient checkpointing on memory efficiency and training throughput in Graph Neural Networks (GNNs) when applied to node classification tasks across datasets with varying graph. Graph Neural Networks (GNNs) have emerged as an efficient alternative to convolutional approaches for vision tasks such as image classification, leveraging patch-based representations instead of raw pixels. These methods construct graphs where image patches serve as nodes, and. 12 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Explaining Vision GNNs: A Semantic and Visual Analysis of Graph-based Image Classification. Research question: What is the effect of gradient checkpointing on memory efficiency and training throughput in Graph Neural Networks (GNNs) when applied to node classification tasks across datasets with varying graph densities?.

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 2.5/10.

3 Results

4 papers retrieved. 12 claims extracted; 0 independently verified. Quality review score: 2.5/10.

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
Experiments were conducted on the ViG-Small variant of the Vision GNN.	×	0.05
In early layers (1-4) on the ImageNet validation set, visual similarity (Sl_vis) is greater than 0.6.	×	0.04
In early layers (1-4) on the ImageNet validation set, embedding similarity (Sl_emb) is greater than 0.9.	×	0.02
In early layers (1-4) on the ImageNet validation set, spatial distances (Dl) are less than 4.	×	0.02
In the final layers on the ImageNet validation set, spatial distances (Dl) are approximately 8.8.	×	0.02
In the final layers on the ImageNet validation set, visual similarity (Sl_vis) is approximately 0.3.	×	0.04
Embedding similarity (Sl_emb) increases suddenly during the final two layers.	×	0.02
Classification accuracy reaches 68.6% in the final layers on the ImageNet validation set.	×	0.08
Graph modularity scores (Ql) for adversarial images (ImageNet-a) start at 0.095.	×	0.03
Graph modularity scores (Ql) for standard images (ImageNet validation) start at 0.236.	×	0.02
Top-1 accuracy drops from 68.6% on standard images to 2.7% on adversarial images.	×	0.02
Correct class probability (pl) in the final layers falls from 0.486 on standard images to 0.018 on adversarial images.	×	0.02

References

- http://arxiv.org/abs/2603.27156v1
- http://arxiv.org/abs/2312.05905v2
- http://arxiv.org/abs/2504.19682v1