

SOVEREIGN: Does GraphMETRO’s mixture-of-experts design improve cross-domain robustness on node classification tasks under

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 29, 2026

Abstract

Multi-view graph refining-based clustering (MGRC) methods aim to facilitate the clustering of data via Graph Neural Networks (GNNs) by learning optimal graphs that reflect the underlying topology of the data. However, current MGRC approaches are limited by their disjoint two-stage process, where the graph structure is learned in the first stage before the GNN messages are propagated in the subsequent stage. Additionally, current approaches neglect the importance of cross-view structural consistency and semantic-level information and only consider intra-view embeddings. To address these issues,

1 Introduction

Analysis of: UNAGI: Unified neighbor-aware graph neural network for multi-view clustering. Research goal: Does GraphMETRO’s mixture-of-experts design improve cross-domain robustness on node classification tasks under distribution shifts compared to standard GNN architectures with identical FLOPs?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score < 6.5.

3 Results

1 papers retrieved. 0 claims extracted, 0 verified. Tribunal: 5.0/10 → RE-
VISE (revision_round=1). Policy: ESCALATE_TO_OWNER.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv
Relevance ranking is query-dependent. Tribunal consensus is LLM-based
and prompt-sensitive.

References

- <https://www.semanticscholar.org/paper/1ed50ff05ac9cb8546be07ff6f3c5685a44a998a>