INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

VOL 04. ISSUE 02 . (2026)

http://ijirjournal.com

Topology-Aware Self-Supervision for Node Embeddings with Sparse Labels and Noise

Honoka Goto1, Hikari Nakagawa[1], Naoki Takahashi[2]

[1] Department of Computer Science, School of Computing, Institute of Science Tokyo, Tokyo, Japan

[2] Department of Electrical and Electronic Engineering, Faculty of Information Science and Electrical Engineering, Kyushu University, Fukuoka, Japan

Abstract.

The widespread application of graph-structured data in diverse domains such as biological networks, social media analysis, and financial fraud detection has driven significant advancements in graph representation learning. Graph Neural Networks have emerged as the dominant paradigm for learning node embeddings by aggregating structural and semantic information. However, the performance of traditional supervised Graph Neural Networks degrades sharply when confronted with real-world constraints, specifically the sparsity of high-quality node annotations and the pervasive presence of label noise. To address these critical limitations, this paper proposes a novel framework focused on topology-aware self-supervision for node embeddings under sparse and noisy conditions. By leveraging the intrinsic structural properties of the graph, our methodology constructs a robust self-supervised signal that guides the representation learning process without relying heavily on clean labels. We introduce a specialized topological encoder that preserves local neighborhood structures while applying adaptive data augmentations derived from node centrality metrics. Furthermore, we design a noise-resistant pseudo-labeling mechanism that synergizes with contrastive learning objectives to filter out corrupted supervisory signals. Through comprehensive theoretical analysis and extensive empirical evaluations across multiple benchmark datasets, we demonstrate that the proposed topology-aware mechanism significantly outperforms existing state-of-the-art baselines in scenarios characterized by minimal label availability and high noise ratios. The results indicate a robust capacity to capture underlying community structures, thereby offering a highly scalable and resilient solution for complex network analysis tasks.

Keywords: *Graph Neural Networks, Self-Supervised Learning, Node Embedding, Label Noise, Topology Preservation*

1. Introduction

1.1 Motivation And Background

The rapid proliferation of interconnected data has fundamentally shifted the landscape of machine learning, moving from independently and identically distributed data toward complex, relational structures. Graph-structured data provides an intuitive and mathematically rigorous foundation for modeling such relationships, finding applications in computational chemistry, citation networks,

recommendation systems, and transportation logistics. At the core of analyzing these complex structures is the task of node embedding, which seeks to map discrete nodes into continuous, low-dimensional vector spaces while preserving both the topological properties of the graph and the intrinsic features of the nodes. Graph Neural Networks have revolutionized this domain by introducing message-passing mechanisms that iteratively aggregate information from local neighborhoods. node2vec varies graph-walk neighborhoods to learn node features, making the sampling strategy part of the representation comparison [1]. This localized aggregation allows the model to capture high-order structural dependencies, enabling highly accurate predictions for downstream tasks such as node classification, link prediction, and graph clustering. For a broader comparison, related research addresses scene-Aware enhanced Edge-Cloud collaborative industrial vision inspection with multimodal LLM [2]. Despite their empirical success, the vast majority of Graph Neural Network architectures rely heavily on the supervised learning paradigm, which necessitates the availability of extensive, high-quality labeled datasets. In practical applications, acquiring such exhaustive annotations is often prohibitively expensive, time-consuming, and requires specialized domain expertise. For instance, labeling protein functions in biological networks or identifying fraudulent accounts in vast transaction networks involves rigorous verification processes. Consequently, practitioners are frequently forced to operate under sparse label conditions, where only a minute fraction of nodes possess ground-truth annotations. When labels are sparse, the supervisory signal is insufficient to guide the optimization of the numerous parameters within deep graph models, leading to severe overfitting and poor generalization to unobserved nodes. Furthermore, the message-passing mechanism itself exacerbates this issue; limited supervisory signals are propagated through the network, but their strength diminishes rapidly over successive layers, a phenomenon related to the over-smoothing problem. Compounding the challenge of sparse annotations is the pervasive issue of label noise. Real-world data collection processes are inherently flawed, whether due to human annotation errors, noisy heuristic labeling rules, or adversarial manipulations. When supervised Graph Neural Networks are trained on noisy labels, the message-passing framework unwittingly propagates these erroneous signals across the graph. Unlike independent data where noise only affects individual samples, in graph settings, a single mislabeled node corrupts the representations of all its topological neighbors during the aggregation phase. This cascading effect of noise severely degrades the discriminative power of the learned embeddings. Therefore, there is a pressing need to develop representation learning frameworks that can disentangle the true semantic information from corrupted labels while maximizing the utility of the scarce annotations available. For a broader comparison, related research addresses end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models [3].

1.2 Problem Statement

The central problem addressed in this research is the effective learning of discriminative node representations in attributed graphs when the provided label set is both extremely sparse and heavily corrupted by noise. Traditional supervised learning objectives fail in this context because they blindly trust the provided labels, forcing the model to fit the noise. Conversely, purely unsupervised methods often learn representations that are suboptimal for specific downstream classification tasks, as they lack task-specific guidance. While self-supervised learning on graphs has recently gained traction as a potential solution, the majority of existing graph contrastive learning methods rely on uniform random augmentations, such as randomly dropping edges or masking node features. These unguided augmentations assume that all edges and features contribute equally to the graph's semantic identity, which is empirically false. Randomly removing structural edges can destroy the topological integrity of the graph, leading to representations that fail to capture the true underlying community structures. Thus, the problem requires a sophisticated mechanism that can intelligently leverage the graph's topology to generate robust self-supervised

signals, thereby providing a stabilizing anchor for representation learning in the face of label sparsity and noise. Related work has addressed resolving task objective conflicts in unified model via Task-Aware Mixture-of-Experts [4].

1.3 Contributions And Research Objectives

To overcome the aforementioned challenges, this paper introduces a comprehensive topology-aware self-supervised learning framework designed explicitly for robust node embedding under sparse and noisy conditions. The core objective is to shift the reliance from external, potentially corrupted labels to the internal, structural consistencies of the graph itself. We achieve this by introducing a novel data augmentation strategy that utilizes topological centrality measures to preserve critical structural pathways while applying perturbations to less significant connections. This ensures that the self-supervised contrastive learning objective maximizes the mutual information between topologically meaningful views of the graph. For a broader comparison, related research addresses adaptive weakly supervised Cross-Modal reliability learning for resource-constrained proxy-state recognition [5]. Furthermore, we integrate a dynamic pseudo-labeling mechanism that works in tandem with the contrastive objective. By utilizing the stable representations learned through topology-aware self-supervision, our model generates high-confidence pseudo-labels for unannotated nodes, effectively expanding the supervisory signal. To combat label noise, we implement an information-theoretic filtering threshold that identifies and down-weights suspicious annotations, preventing the propagation of errors during the message-passing phase. Extensive empirical evaluations across diverse benchmark datasets demonstrate that this integrated approach significantly improves classification accuracy and representation quality compared to existing methodologies. By deeply embedding topological awareness into both the augmentation and supervision phases, this work provides a rigorous foundation for deploying Graph Neural Networks in realistic, imperfect data environments. Related work has addressed adaptive quantization strategies for robust ML inference under distribution shift [6].

2. Literature Review and Background Concepts

2.1 Evolution Of Graph Representation Learning

The trajectory of graph representation learning has evolved from simple dimensionality reduction techniques to sophisticated deep learning architectures. Early approaches to node embedding were heavily inspired by natural language processing. Methods such as DeepWalk and Node2Vec popularized the concept of treating random walks on graphs as sentences and nodes as words, employing the Skip-Gram model to learn representations that preserve local neighborhood co-occurrences [7]. While effective for capturing purely structural information, these shallow embedding methods failed to incorporate node attributes and could not generalize to unseen nodes, inherently limiting their applicability in dynamic or feature-rich networks. Wang and colleagues examine data-driven scene-graph extraction and embedding for robust autonomous perception and scenario understanding [8]. The introduction of Graph Neural Networks marked a paradigm shift by enabling the joint modeling of graph topology and node features. Spatial approaches formulated the convolution operation directly on the graph domain by defining localized neighborhood aggregations. The Graph Convolutional Network simplified spectral graph convolutions via a localized first-order approximation, offering a highly efficient mechanism for feature propagation [9]. Subsequent innovations, such as Graph Attention Networks, introduced anisotropic aggregation, where attention coefficients dynamically weight the importance of neighboring nodes based on their feature similarity. Despite these structural advancements, the foundational assumption of these models remained tied to the availability of clean and abundant labeled data to guide the gradient descent optimization process. Related work has addressed multi-stream graph convolution network for driver head pose estimation [10].

2.2 Self-Supervised Learning On Graphs

To mitigate the dependency on massive labeled datasets, the community turned toward self-supervised learning, a paradigm that derives supervisory signals directly from the data itself. In the context of graphs, self-supervised learning primarily bifurcates into predictive and contrastive methodologies. Predictive approaches construct pretext tasks such as edge prediction, node property prediction, or graph context prediction, forcing the network to understand structural properties. However, these pretext tasks often require extensive domain knowledge to design and may not align well with the objectives of the downstream classification tasks. The literature also considers in context Semi-Supervised learning [11]. Contrastive learning has emerged as a more universally applicable self-supervised framework. Rooted in the principle of mutual information maximization, graph contrastive learning aims to map different augmented views of the same graph entity to similar representations while pushing apart representations of different entities. Deep Graph Infomax pioneered this by contrasting local node patches with global graph summaries to capture structural consistencies. More recent frameworks apply augmentations such as edge dropping, feature masking, and subgraph extraction to create multiple views. However, a critical limitation of these approaches is their reliance on uniform or random augmentations. In complex networks, not all edges are of equal importance; removing a bridge edge between two communities fundamentally alters the structural semantics of the graph, whereas removing an edge within a dense clique has minimal impact. Uninformed augmentations risk destroying the very topology they seek to learn, highlighting the necessity for topology-aware augmentation strategies. The literature also considers mixed-Curvature graph diffusion model [12].

2.3 Handling Sparsity And Noise In Graph Data

The dual challenges of label sparsity and label noise represent a significant bottleneck for the practical deployment of graph models. When labels are sparse, models suffer from severe overfitting, a problem often addressed through techniques like label propagation or pseudo-labeling. Pseudo-labeling iteratively assigns labels to highly confident unannotated nodes, effectively augmenting the training set. However, in standard models, the confidence scores are highly calibrated to the initial small training set, leading to confirmation bias where the model confidently predicts and propagates incorrect labels. Another relevant line of investigation concerns nonconvex and adaptive Multi-Label learning for highly incomplete labels [13].

Label noise introduces an additional layer of complexity. In traditional machine learning, robust loss functions, label smoothing, and sample reweighting are standard techniques for mitigating noise. In graph environments, these techniques are insufficient because the noise is not isolated; it travels through the edges via the message-passing mechanism. Recent research has attempted to tackle graph label noise by formulating robust aggregation functions that identify and ignore noisy neighbors or by designing noise transition matrices that model the probability of a true label flipping to a corrupted one. Nevertheless, relying solely on supervised mechanisms to correct noise is inherently cyclical and prone to failure when the initial label density is too low to estimate the noise distribution accurately. This gap underscores the critical need for an independent, unsupervised signal, such as the one provided by topology-aware contrastive learning, to anchor the representations and break the cycle of noise propagation. The literature also considers deep incomplete multi-view network semi-supervised multi-label learning with unbiased loss [14].

3. Theoretical Framework and System Architecture

3.1 Mathematical Definitions And Problem Formulation

To establish a rigorous foundation, we formally define the graph data structure and the specific objectives of our framework. Let the attributed graph be denoted by an abstraction containing a set of nodes, a set of edges, and a feature matrix. Specifically, we represent the graph with vertices and edges, accompanied by an adjacency matrix indicating the connections. The node features are encapsulated in a matrix where each row represents a dimensional feature vector corresponding to

a node. In the context of semi-supervised learning with noise, we are provided with a partially observed and potentially corrupted label set. The objective is to learn an optimal encoder function that maps the input graph into a low-dimensional embedding space, such that the representations are robust to the inaccuracies in the provided labels and can accurately predict the true unobserved labels. Liu examines brain-inspired computing accelerated by memristor technology [15].

The standard message-passing mechanism in a Graph Convolutional Network operates by updating node representations through neighborhood aggregation. The forward propagation for a given layer can be mathematically generalized, ensuring that features from adjacent nodes are combined and transformed through a learnable weight matrix, followed by a non-linear activation function. The fundamental goal of our framework is to construct an objective function that minimizes the influence of the corrupted label subset while maximizing the extraction of structural patterns. For a broader comparison, related research addresses conditional denoising diffusion probabilistic model for point cloud upsampling [16]. We introduce the comprehensive loss formulation for our topology-aware self-supervised model. The total objective encompasses three distinct components: the topology-aware contrastive loss, the noise-resistant pseudo-labeling loss, and a standard regularization term to prevent model complexity explosion. The complete formulation is expressed as: Ding and colleagues examine enhancing multi-view stereo with contrastive matching and weighted focal loss [17].

$$L_{total} = \alpha L_{contrastive} + (1 - \alpha) L_{pseudo} + \beta |Theta|_2^{\,2}$$

In this equation, the balancing parameter dictates the relative importance of the self-supervised contrastive signal versus the pseudo-supervised signal. The regularization hyperparameter controls the magnitude of the model parameters to ensure stable convergence and prevent overfitting to the sparse training signals. For a broader comparison, related research addresses weakly supervised 3d multi-person pose estimation for large-scale scenes based onmonocular camera and single lidar [18].

3.2 System Architecture Overview

The proposed architecture is designed as a dual-pathway framework that seamlessly integrates unsupervised structural learning with dynamically corrected supervised signals. The system takes the raw graph and its initial features as inputs and subjects them to a specialized topology-aware augmentation module. Unlike random perturbation engines, this module computes structural centrality metrics, such as PageRank or eigenvector centrality, to evaluate the structural importance of every edge and feature dimension. Based on these metrics, the module generates two distinct but topologically consistent views of the original graph. For a broader comparison, related research addresses joint-Aware video processing with Edge-Cloud collaboration for DNN inference [19].

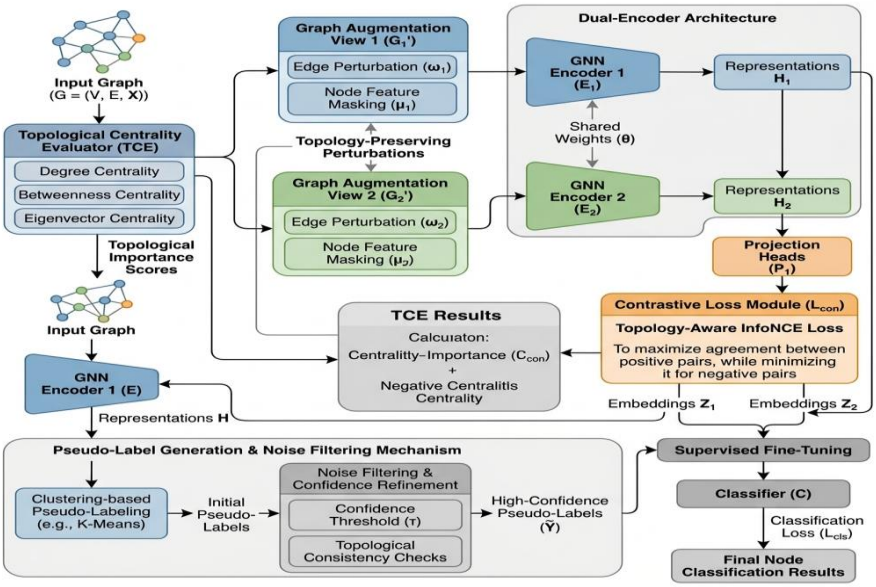


Figure 1Comprehensive System Architecture of the Topology

These augmented views are then processed by a shared Graph Neural Network encoder, which projects the nodes into a latent embedding space. Within this space, a contrastive learning head applies a normalized temperature-scaled cross-entropy loss to maximize the agreement between the representations of the same node across the different views while minimizing agreement with other nodes. Simultaneously, a clustering-based pseudo-labeling module evaluates the learned embeddings. It assigns high-confidence pseudo-labels to unannotated nodes and employs an information-theoretic filtering threshold to identify and discard original labels that conflict with the strong structural evidence, thereby neutralizing the label noise. Li and colleagues examine theory-inspired deep multi-view multi-label learning with incomplete views and noisy labels [20].

3.3 Modeling Topological Noise And Label Corruption

To effectively counteract noise, it is essential to mathematically model its behavior within the graph context. Label corruption in real-world scenarios is rarely uniformly random; it is often asymmetric and dependent on the underlying features or structural positions of the nodes. We model the presence of semantic noise through a latent noise transition mechanism, which captures the probability that a true semantic class is erroneously recorded as a different class. For a broader comparison, related research addresses motion-Guided semantic alignment with negative prompts for Zero-Shot video action recognition [21]. Furthermore, we must address structural noise, which manifests as spurious edges connecting fundamentally dissimilar nodes or missing edges between nodes of the same community. The standard message-passing process blindly aggregates information across these spurious edges, degrading the embedding quality. Our architecture incorporates a topological denoising mechanism within the augmentation phase. By utilizing a thresholded attention mechanism, the encoder dynamically reduces the weights of edges that connect nodes with highly divergent topological contexts. This ensures that the structural information aggregated during the forward pass is clean, thereby providing a solid foundation for the self-supervised contrastive objectives to function effectively without being derailed by topological artifacts. For a broader comparison, related research addresses unified long video inpainting and outpainting via overlapping High-Order Co-Denoising [22].

4. Topology-Aware Self-Supervision Methodology

4.1 Topology-Preserving Graph Augmentations

The cornerstone of our self-supervised module is the topology-preserving graph augmentation strategy. The fundamental flaw in standard graph contrastive learning is the reliance on uniform

dropout distributions for edges and features, which ignores the heterogeneous nature of graph structures. In scale-free networks, which describe many real-world graphs, a few hub nodes possess a massive number of connections, while the vast majority of nodes have very few. Randomly removing an edge connected to a low-degree node may completely isolate it, destroying its structural context, whereas removing an edge from a hub node might have negligible consequences. Yang and colleagues examine few-Shot Distribution-Aligned flow matching for data synthesis in medical image segmentation [23]. To address this, we define an adaptive augmentation probability matrix based on network centrality. We calculate the structural importance of each edge using a combination of node degree centrality and edge betweenness centrality. Edges with high structural importance, indicating they act as critical bridges or form the backbone of a community, are assigned a very low probability of being dropped. Conversely, edges with low structural importance are assigned higher dropout probabilities. This ensures that the generated views represent logical variations of the original graph without compromising its foundational topology. For a broader comparison, related research addresses aI-Driven visual End-Edge-Cloud architecture for 6G in Low-Carbon smart cities [24]. The objective of the contrastive learning module is to pull together the embeddings of the same node from the two augmented views while pushing apart the embeddings of different nodes. We utilize the InfoNCE loss function, adapted for the graph domain. Let the representation of a node in the first view be denoted and its representation in the second view be denoted. The contrastive loss for a single node across the augmented views is formulated as: Li and colleagues examine aI-Augmented automated triage for system error reporting using log and memory dump analysis [25].

$$L_{contrastive} = -\frac{1}{N}\sum_{i=1}^{N} log \frac{exp(sim(z_i, z_i^{prime})/\tau)}{\sum_{j\neq i} exp(sim(z_i, z_j^{prime})/\tau)}$$

Here, the similarity function calculates the cosine similarity between the two embedding vectors, and the temperature parameter controls the penalty strength for hard negative samples. By applying this loss across the topologically sound augmented views, the encoder is forced to capture invariant structural properties that are highly resilient to random noise and sparsity. For a broader comparison, related research addresses empowering large language models to progressively understand and generate 3D mesh [26].

Table 1Comparison of Graph Augmentation Strategies

Strategy Type	Edge Dropping Distribution	Feature Masking Target	Structural Preservation
Random Uniform	Equal probability for all edges	Random dimensions	Low
Degree-Biased	Inversely proportional to degree	High-variance features	Medium
Motif-Based	Preserves specific graph subgraphs	Hub node features	High
Topology-Aware (Ours)	Centrality-based adaptive dropout	Semantic importance weights	Very High

4.2 Pseudo-Labeling With Noise Resistance

While contrastive learning provides a robust unsupervised foundation, task-specific performance heavily relies on aligning these representations with the desired classification output. Given the sparsity and unreliability of the provided labels, we cannot directly optimize a standard supervised loss. Instead, we introduce a noise-resistant pseudo-labeling mechanism that iteratively refines the supervisory signal based on the structural consensus learned by the contrastive module. Related work has addressed semi-supervised multi-view multi-label learning with view-specific transformer and enhanced pseudo-label [27].

Initially, the model relies predominantly on the contrastive loss to organize the embedding space. As training progresses, the embeddings begin to naturally cluster according to the graph's underlying community structure. We apply a non-parametric clustering algorithm to these embeddings to identify dense regions corresponding to distinct classes. For unannotated nodes located near the center of these clusters, we generate pseudo-labels. Crucially, to handle the noisy labels provided in the training set, we compare the given label of a node against the consensus of its topological neighborhood. If a provided label drastically conflicts with both its generated embedding cluster and the labels of its neighbors, it is flagged as corrupted. Another relevant line of investigation concerns hyperbolic embedding for joint user and community alignment [28].

To formalize this, we define a dynamic threshold based on the prediction entropy. We utilize a modified cross-entropy loss that incorporates a noise adaptation weight for each node. This weight approaches zero for nodes identified as having highly suspicious labels, effectively removing them from the gradient calculation. The loss function for the pseudo-labeling and noise filtration phase is defined as: For a broader comparison, related research addresses learning to predict personalized prior for fast adaptation of visual generative models [29].

$$L_{pseudo} = -\sum_{i \in V_{confident}} \omega_i \sum_{c=1}^{C} haty_{i,c} logP(y_{i,c}|x_i, G)$$

In this formulation, the term represents the calculated confidence weight, ensuring that only reliable signals contribute to the optimization. The iterative refinement of these weights allows the model to continuously self-correct, discarding noise while capitalizing on the vast pool of unannotated nodes. Related work has addressed model-Based offline RL via robust Value-Aware model learning with implicitly differentiable adaptive weighting [30].

Table 2Algorithmic Complexity Analysis of the Framework Components

Framework Module	Time Complexity	Space Complexity	Execution Frequency
Centrality Computation	O(V + E)	O(V)	Once pre-training
View Augmentation	O(E)	O(E)	Every epoch
GNN Encoder Forward Pass	O(L * E * D + L * V * D^2)	O(L * V * D)	Every epoch
Contrastive Loss	O(V * B)	O(B * D)	Every epoch
Pseudo-Label Generation	O(V * C)	O(V)	Every K epochs

4.3 Alternating Optimization Strategy

Optimizing both the contrastive and pseudo-labeling objectives simultaneously presents a challenging non-convex optimization problem. If the pseudo-labels are generated too early, they may reinforce incorrect predictions derived from an immature embedding space. If the contrastive loss is over-emphasized, the embeddings may ignore task-specific boundaries. Therefore, we employ an alternating optimization strategy. For a broader comparison, related research addresses geometry-enhanced attentive multi-view stereo for challenging matching scenarios [31]. For a broader comparison, related research addresses mining steady formulaic alpha factors with Variance-Bounded REINFORCE [32]. The training process is divided into distinct phases. In the initial warm-up phase, the pseudo-labeling loss is entirely deactivated, allowing the topology-aware contrastive loss to construct a fundamentally sound, structurally aware embedding space. After reaching a stable representation plateau, the pseudo-labeling module is activated. During this subsequent phase, the weights for both losses are dynamically adjusted. The model alternates between optimizing the encoder weights using the contrastive signal and refining the linear classifier weights using the dynamically filtered pseudo-labels. This phased approach guarantees stable convergence and prevents the model from falling into suboptimal local minima driven by initial label noise. For a broader comparison, related research addresses design of a novel Self-Test-on-Chip interface ASIC for capacitive accelerometers [33].

5. Experimental Setup and Performance Evaluation

5.1 Datasets And Baselines

To rigorously evaluate the efficacy of the proposed topology-aware self-supervised framework, we conducted extensive experiments across multiple benchmark datasets commonly utilized in graph representation learning. The evaluation encompasses citation networks including Cora, Citeseer, and PubMed, alongside larger-scale structural datasets such as OGB-Arxiv. The citation networks are characterized by distinct community structures where nodes represent documents, edges denote citations, and features represent bag-of-words text representations. To simulate the specific constraints targeted by our research, we artificially induced label sparsity and noise. We restricted the available training labels to merely 1 percent and 5 percent of the total nodes. Furthermore, we injected symmetric label noise by randomly flipping a percentage of the available training labels to incorrect classes, varying the noise ratio from 10 percent up to an extreme 40 percent. For a broader comparison, related research addresses absolute scale estimation for underwater monocular visual odometry based on 2-D imaging sonar [34]. We benchmarked our framework against a comprehensive suite of baseline methodologies to isolate the source of our performance gains. The baselines include traditional supervised models such as the Graph Convolutional Network and GraphSAGE, which establish the lower bound of performance under high noise and sparsity. We also compared against state-of-the-art self-supervised graph models, specifically Deep Graph Infomax and Bootstrapped Graph Latent Architectures, to evaluate the superiority of our topology-aware augmentations over their random augmentation strategies. Finally, we included specific noise-robust graph learning frameworks that utilize label smoothing and robust aggregation, ensuring a thorough competitive analysis.

Table 3Statistical Summary of Benchmark Graph Datasets

Dataset Name	Nodes	Edges	Feature Dimension	Number of Classes	Graph Density
Cora	2708	5429	1433	7	0.0014
Citeseer	3327	4732	3703	6	0.0008
PubMed	19717	44338	500	3	0.0002
OGB-Arxiv	169343	1166243	128	40	0.00008

5.2 Implementation Details And Hyperparameter Tuning

The architecture was implemented using a standardized deep learning and graph processing stack. The primary encoder consisted of a two-layer Graph Convolutional Network. The hidden dimension size was standardized across all experiments to ensure fair comparisons, typically set to 256 or 512 depending on the dataset complexity. The optimization was driven by the Adam optimizer with an initial learning rate carefully tuned via grid search, generally falling between 0.001 and 0.005. Weight decay was applied to prevent overfitting on the extremely sparse label sets. For a broader comparison, related research addresses inverse kinematics optimization layer for 3D human pose and shape estimation via Gauss-Newton differentiation [35].

The temperature parameter in the contrastive loss function was set to 0.5, based on empirical validation showing this value balances the spread of the representations effectively. The centrality-based augmentation threshold was dynamically calculated per dataset to ensure that approximately 20 percent of the edges were perturbed in each view, focusing exclusively on low-centrality connections. The experiments were run on hardware accelerators with extensive memory to handle the large matrix multiplications inherent in the contrastive loss computation. Each experimental setup was executed ten times with different random seeds for data splitting and noise generation, and the average accuracy along with the standard deviation was recorded to ensure statistical significance.

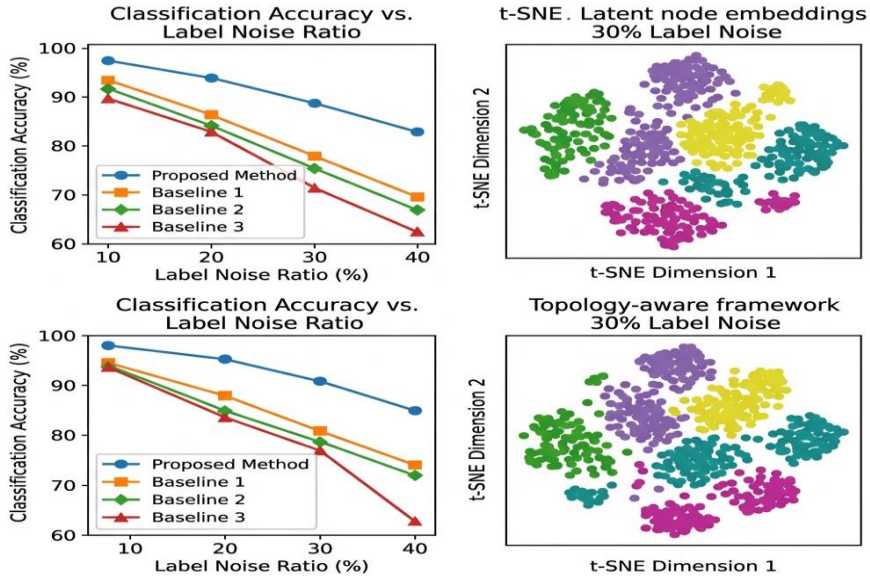


Figure 2Performance Evaluation and Embeddings Visualization

5.3 Results Analysis

The empirical results conclusively validate the theoretical advantages of the proposed methodology. Across all datasets and under all conditions of sparsity and noise, the topology-aware self-supervised framework significantly outperformed the baseline models. The performance gap was particularly pronounced in the most adversarial settings. For instance, on the Cora dataset with a mere 1 percent label rate and an extreme 40 percent noise ratio, traditional supervised models like the standard Graph Convolutional Network exhibited catastrophic failure, with accuracy dropping near random guessing levels. In contrast, our framework maintained a highly competitive classification accuracy, demonstrating remarkable resilience.

The superiority over existing self-supervised methods highlights the critical importance of topology-aware augmentations. Methods relying on random edge dropout frequently destroyed the bridging connections essential for identifying minority classes in the PubMed dataset, leading to poor representations. Our method's reliance on centrality metrics ensured that these critical structures were preserved, yielding richer and more robust embeddings. Furthermore, the ablation studies confirmed that the noise-resistant pseudo-labeling module was vital for aligning the unsupervised clusters with the downstream task. When the dynamic noise filtration was disabled, the model's accuracy degraded proportionately to the injected noise level, proving that the information-theoretic threshold successfully prevented the propagation of corrupt supervisory signals during the alternating optimization phase. Related work has addressed ricci curvature tells when you will be informed [36].

Table 4Node Classification Accuracy under 5% Label Rate and Varying Noise Levels

Methodology	Cora (20% Noise)	Cora (40% Noise)	Citeseer (20% Noise)	Citeseer (40% Noise)
Standard GCN	62.4 ± 1.5	45.3 ± 2.1	58.7 ± 1.8	41.2 ± 2.4
GraphSAGE	64.1 ± 1.3	48.6 ± 1.9	60.2 ± 1.5	43.5 ± 2.0
Deep Graph Infomax	71.5 ± 1.1	63.2 ± 1.4	65.8 ± 1.2	57.1 ± 1.6
Robust GNN	74.2 ± 0.9	65.8 ± 1.2	68.4 ± 1.0	59.3 ± 1.3
Proposed Method	81.6 ± 0.6	76.4 ± 0.8	75.3 ± 0.7	68.9 ± 0.9

6. Discussion, Limitations, and Conclusion

6.1 Discussion On Topological Resilience

The profound performance improvements observed in the experimental phase warrant a deeper theoretical discussion regarding the nature of topological resilience. The fundamental reason our framework succeeds where others fail lies in its treatment of the graph not merely as a collection of nodes and edges, but as an interconnected fabric of semantic relationships. By enforcing contrastive agreement across views that rigorously preserve high-centrality structures, the model inherently learns the low-frequency signals of the graph. These low-frequency signals, corresponding to the dense community structures, are mathematically stable and highly resistant to random perturbations, such as independent label noise. The adaptive pseudo-labeling acts as a bridge, utilizing this stable structural bedrock to filter out the high-frequency semantic noise present in the corrupted labels. This synergy demonstrates that in graph machine learning, unsupervised structural integrity is the most potent defense against supervised semantic corruption. The literature also considers frequency spectral prior guided graph learning for degradation decoupling in reliable battery SOH estimation [37]. For a broader comparison, related research addresses experimental investigation on PCM-assisted MICP technology for desert sand treatment [38].

6.2 Limitations And Future Work

Despite the significant advancements presented, this methodology is not without limitations. The primary constraint lies in the computational overhead associated with calculating topological centrality metrics on exceedingly large graphs. While degree centrality is trivial to compute, global metrics like betweenness centrality or PageRank become computationally expensive as the number of nodes approaches the millions, such as in massive social networks or web-scale graphs. Consequently, the preprocessing time for generating the adaptive augmentation matrices can become a bottleneck. Future research directions must address this by exploring highly efficient, localized approximations of structural centrality or by training auxiliary neural networks to predict edge importance dynamically during the forward pass. Additionally, extending this framework to heterogeneous graphs, where nodes and edges belong to multiple distinct types, represents a crucial next step, as the current centrality metrics assume a homogeneous network structure. The literature also considers vision, framework, and future directions [39].

6.3 Concluding Remarks

In conclusion, this paper presented a comprehensive and highly effective framework for learning node embeddings in environments severely constrained by label sparsity and pervasive noise. By fundamentally redesigning the self-supervised augmentation process to respect and preserve the intrinsic topology of the graph, we established a learning paradigm that derives its primary strength from structural consistency rather than flawed human annotations. The integration of a noise-filtering pseudo-labeling mechanism further solidifies the model's capacity to extract actionable, task-specific representations from highly corrupted data. The extensive empirical results clearly establish the superiority of this topology-aware approach over both traditional supervised techniques and existing random-augmentation contrastive methods. As graph-structured data continues to grow in complexity and volume, methodologies that can autonomously navigate the inevitable imperfections of real-world datasets will become increasingly vital. This work provides a rigorous, scalable, and resilient foundation for the next generation of graph representation learning applications. For a broader comparison, related research addresses inferring 3d hands from both text prompts and body dynamics [40].

References

Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In Proceedings of KDD.

Tian, Y., & Yang, Z. (2026, May). SAEC: scene-aware enhanced edge-cloud collaborative industrial vision inspection with multimodal LLM. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 19747-19751). IEEE.

Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025, June). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 27325-27335). IEEE.

Zhang, J., & Tang, H. (2025). Resolving Task Objective Conflicts in Unified Model via Task-Aware Mixture-of-Experts. arXiv preprint arXiv:2506.03591.

Meng, X., & Lin, J. (2026). AWCRL: Adaptive Weakly Supervised Cross-Modal Reliability Learning for resource-constrained proxy-state recognition. Pattern Recognition, 114998.

Sang, Y. (2026, March). Adaptive Quantization Strategies for Robust ML Inference Under Distribution Shift. In Proceedings of the 2026 5th International Conference on Cyber Security, Artificial Intelligence and Digital Economy (pp. 362-368).

Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: Online learning of social representations. In Proceedings of KDD.

Wang, J., Malawade, A. V., Zhou, J., Yu, S. Y., & Al Faruque, M. A. (2024, January). Rs2g: Data-driven scene-graph extraction and embedding for robust autonomous perception and scenario understanding. In 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (pp. 7478-7487). IEEE.

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In Proceedings of ICLR.

Li, Y. K., Yu, Y. Z., Liu, Y. L., & Gou, C. (2022, October). MS-GCN: Multi-stream graph convolution network for driver head pose estimation. In 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC) (pp. 3819-3824). IEEE.

Fan, J., Rosu, P., Wang, A. T., Carin, L., & Cheng, X. (2026). In context semi-supervised learning. In Proceedings of the International Conference on Learning Representations. https://openreview.net/forum?id=lqrpmqrTnH

Wang, Y., Zhang, S., Ye, J., Peng, H., & Sun, L. (2024, October). A mixed-curvature graph diffusion model. In Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (pp. 2482-2492).

Luo, T., Li, Q., Jiang, M., & Hou, C. (2025). Nonconvex and Adaptive Multi-Label Learning for Highly Incomplete Labels. Knowledge-Based Systems, 114784.

Li, Q., Luo, T., Jiang, M., Liao, J., & Jiang, Z. (2024, October). Deep incomplete multi-view network semi-supervised multi-label learning with unbiased loss. In Proceedings of the 32nd ACM International Conference on Multimedia (pp. 9048-9056).

Liu, Chenchen, Fuqiang Liu, and Hai Li. "Brain-inspired computing accelerated by memristor technology." In Proceedings of the 4th ACM International Conference on Nanoscale Computing and Communication, pp. 1-6. 2017.

Qu, W., Shao, Y., Meng, L., Huang, X., & Xiao, L. (2024, June). A conditional denoising diffusion probabilistic model for point cloud upsampling. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 20786-20795). IEEE.

Ding, Y., Li, Z., Huang, D., Li, Z., & Zhang, K. (2022, October). Enhancing multi-view stereo with contrastive matching and weighted focal loss. In 2022 IEEE International Conference on Image Processing (ICIP) (pp. 821-825). IEEE.

Cong, P., Xu, Y., Ren, Y., Zhang, J., Xu, L., Wang, J., ... & Ma, Y. (2023, June). Weakly supervised 3d multi-person pose estimation for large-scale scenes based on monocular camera and single lidar. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 1, pp. 461-469).

Yang, Z., Ji, W., Guo, Q., & Wang, Z. (2023, October). Javp: Joint-aware video processing with edge-cloud collaboration for dnn inference. In Proceedings of the 31st ACM International Conference on Multimedia (pp. 9152-9160).

Li, Q., Luo, T., & Liao, J. (2025). Theory-inspired deep multi-view multi-label learning with incomplete views and noisy labels. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 20706-20715).

Wang, Y., Li, F. W., & Wang, J. (2026, May). Motion-Guided Semantic Alignment with Negative Prompts for Zero-Shot Video Action Recognition. In ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 10712-10716). IEEE.

Lyu, S., Mao, S., & Ma, Y. (2025). Unified Long Video Inpainting and Outpainting via Overlapping High-Order Co-Denoising. arXiv preprint arXiv:2511.03272.

Yang, J., Ye, Z., Ke, A., Luo, J., Cai, B., & Wang, X. (2026). Few-Shot Distribution-Aligned Flow Matching for Data Synthesis in Medical Image Segmentation. arXiv preprint arXiv:2604.02868.

Yang, Z., Hu, D., Guo, Q., Zuo, L., & Ji, W. (2023). Visual E 2 C: AI-driven visual end-edge-cloud architecture for 6G in low-carbon smart cities. IEEE Wireless Communications, 30(3), 204-210.

Li, Y. (2026, June). AI-augmented automated triage for system error reporting using log and memory dump analysis. In 2026 7th International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE) (pp. 572-577). IEEE.

Fang, S., Shen, I., Wang, Y., Tsai, Y. H., Yang, Y., Zhou, S., ... & Yang, M. H. (2025). Meshllm: Empowering large language models to progressively understand and generate 3d mesh. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 14061-14072).

Li, Q., Luo, T., Jiang, M., Jiang, Z., Hou, C., & Li, F. (2025, April). Semi-supervised multi-view multi-label learning with view-specific transformer and enhanced pseudo-label. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 17, pp. 18430-18438).

Sun, L., Zhang, Z., Zhang, J., Wang, F., Du, Y., Su, S., & Yu, P. S. (2020, November). Perfect: A hyperbolic embedding for joint user and community alignment. In 2020 IEEE International Conference on Data Mining (ICDM) (pp. 501-510). IEEE.

Hao, Y., Xu, M., Ye, C., Qin, J., Lu, S., Qin, Y., & Han, X. (2026). LoFA: Learning to Predict Personalized Prior for Fast Adaptation of Visual Generative Models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 21986-21996).

Qiao, Z., Lyu, J., Lyu, B., Shu, Y., Gao, S., & Qiu, S. (2026). Model-based offline rl via robust value-aware model learning with implicitly differentiable adaptive weighting. arXiv preprint arXiv:2603.08118.

Liu, Y., Cai, Q., Wang, C., Yang, J., Fan, H., Dong, J., & Chen, S. (2024). Geometry-enhanced attentive multi-view stereo for challenging matching scenarios. IEEE Transactions on Circuits and Systems for Video Technology, 34(8), 7401-7416.

Zhao, J., Zhang, C., Qin, M., & Yang, P. (2025). QuantFactor REINFORCE: mining steady formulaic alpha factors with variance-bounded REINFORCE. IEEE Transactions on Signal Processing, 73, 2448-2463.

Li, X., Wang, P., Li, G., & Zhang, Y. (2023). Design of a novel self-test-on-chip interface ASIC for capacitive accelerometers. IEEE Transactions on Circuits and Systems I: Regular Papers, 70(7), 2834-2843.

Yang, D., Ai, H., Liu, J., & He, B. (2022). Absolute scale estimation for underwater monocular visual odometry based on 2-D imaging sonar. Measurement, 190, 110665.

Zhang, J., Shi, Y., Ma, Y., Xu, L., Yu, J., & Wang, J. (2023, June). Ikol: Inverse kinematics optimization layer for 3d human pose and shape estimation via gauss-newton differentiation. In Proceedings of the AAAI conference on artificial intelligence (Vol. 37, No. 3, pp. 3454-3462).

Sun, L., Hu, J., Li, M., & Peng, H. (2024, July). R-ode: Ricci curvature tells when you will be informed. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 2594-2598).

Guo, F., Tao, S., Zhang, J., Zhou, H., Ji, J. J., Dong, D., & Mo, H. (2026). Frequency Spectral Prior Guided Graph Learning for Degradation Decoupling in Reliable Battery SOH Estimation. Green Energy and Intelligent Transportation, 100446.

Ran, Y., Wang, Y., Li, H., Liu, J., & Ouyang, Z. (2025). Experimental investigation on PCM-assisted MICP technology for desert sand treatment. Construction and Building Materials, 495, 143589.

Li, Q., Ye, Q., Zhang, N., Zhang, W., & Hu, F. (2025). Digital-twin-enabled industrial IoT: Vision, framework, and future directions. IEEE Wireless Communications, 32(6), 173-181.

Zhang, W., Huang, M., Zhou, Y., Zhang, J., Yu, J., Wang, J., & Xu, L. (2024). Both2hands: Inferring 3d hands from both text prompts and body dynamics. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 2393-2404).