Novelty Assessment Report

Paper: Rethinking the Gold Standard: Why Discrete Curvature Fails to Fully Capture Over-squashing in GNNs?

PDF URL: https://openreview.net/pdf?id=QYtmqCoilk

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

As a topological invariant for discrete structures, discrete curvature has been widely adopted in the study of complex networks and graph neural networks. A prevailing viewpoint posits that edges with highly negative curvature will induce graph bottlenecks and the over-squashing phenomenon. In this paper, we critically re-examine this view and put forward our central claim: **high negative curvature is a sufficient but not a necessary condition for over-squashing**. We first construct a family of counterexamples demonstrating the failure of discrete curvature, where some edges are severely squashed, but the curvature still appears positive. Furthermore, extensive experiments demonstrate that the most commonly used discrete curvature measure --- Ollivier–Ricci curvature --- fails to detect as many as 30%~40% of over-squashed edges. To alleviate this limitation, we propose Weighted Augmented Forman-3 Curvature ($\mathsf{WAF3}$), which significantly improves the detection of over-squashed edges. Additionally, we develop a highly efficient approximation algorithm for $\mathsf{WAF3}$, enabling curvature computation on graphs with five million edges in only 23.6 seconds, which is 133.7 times faster than the existing algorithm with the lowest complexity for curvatures.

Core Task Landscape

This paper addresses: **Detecting Over-Squashing in Graph Neural Networks Using Discrete Curvature**

A total of **30 papers** were analyzed and organized into a taxonomy with **19 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Theoretical Foundations and Analysis**
- **Rewiring and Structural Modification**
- **Curvature-Enhanced GNN Architectures**
- **Curvature-Based Structural Encodings**
- **Domain-Specific Applications**
- **Surveys and Comparative Studies**
- **Related Methodologies**

Complete Taxonomy Tree

- Detecting Over-Squashing in Graph Neural Networks Using Discrete Curvature Survey Taxonomy
- Theoretical Foundations and Analysis
  - Curvature-Based Over-Squashing Characterization ★ (4 papers)
  - [0] Rethinking the Gold Standard: Why Discrete Curvature Fails to Fully Capture Over-squashing in GNNs? (Anon et al., 2026) View paper
  - [10] Understanding over-squashing and bottlenecks on graphs via curvature (Topping, 2022) View paper
  - [11] Oversquashing in gnns through the lens of information contraction and graph expansion (Pradeep Kr. Banerjee, 2022) View paper
  - [25] Curvature and over-squashing in Graph Neural Networks (Giovanni, n.d.) View paper
  - Over-Squashing and Over-Smoothing Trade-offs (3 papers)
  - [16] Revisiting Over-smoothing and Over-squashing Using Ollivier-Ricci Curvature (Nguyen Khang, 2022) View paper
  - [22] On the Trade-off between Over-smoothing and Over-squashing in Deep Graph Neural Networks (Jhony H. Giraldo, 2023) View paper
  - Geometric and Topological Perspectives (2 papers)
  - [12] A Remedy for Over-Squashing in Graph Learning via Forman-Ricci Curvature based Graph-to-Hypergraph Structural Lifting (Banf, 2025) View paper
  - [14] Over-Squashing in Riemannian Graph Neural Networks (Balla, 2023) View paper
- Rewiring and Structural Modification
  - Curvature-Guided Rewiring Methods (3 papers)
  - [1] Rewiring networks for graph neural network training using discrete geometry (Bober, 2023) View paper
  - [6] The effectiveness of curvature-based rewiring and the role of hyperparameters in gnns revisited (Floriano Tori, 2024) View paper
  - [9] Mitigating Over-Smoothing and Over-Squashing using Augmentations of Forman-Ricci Curvature (Fesser, 2023) View paper
  - Curvature-Based Pooling and Aggregation (1 papers)
  - [15] Curvature-based Pooling within Graph Neural Networks (Roth, 2023) View paper
  - Random and Stochastic Rewiring Approaches (1 papers)
  - [4] Effects of Random Edge-Dropping on Over-Squashing in Graph Neural Networks (J Singh, 2025) View paper

- Curvature-Enhanced GNN Architectures
  - Curvature-Constrained Message Passing (3 papers)
  - [5] Curvature constrained MPNNs: Improving message passing with local structural properties (Hugo Attali, 2025) View paper
  - [13] How curvature enhance the adaptation power of framelet gcns (Shi Dai, 2023) View paper
  - [30] Curvature MPNNs: Improving Message Passing with Local Structural Properties (H Attali, n.d.) View paper
  - Curvature-Based Regularization and Dropout (1 papers)
  - [2] Curvdrop: A ricci curvature based approach to prevent graph neural networks from over-smoothing and over-squashing (Yang Liu, 2023) View paper
- Curvature-Based Structural Encodings
  - Local Curvature Profile Encodings (1 papers)
  - [23] Effective Structural Encodings via Local Curvature Profiles (Fesser, 2023) View paper
  - Curvature-Informed Graph Transformations (1 papers)
  - [24] A Rewiring Contrastive Patch PerformerMixer Framework for Graph Representation Learning (Zhongtian Sun, 2023) View paper
- Domain-Specific Applications
  - Neuroscience and Brain Network Analysis (1 papers)
  - [7] Signed Curvature Graph Representation Learning of Brain Networks for Brain Age Estimation (Jingming Li, 2024) View paper
  - Action Recognition and Temporal Graphs (1 papers)
  - [3] Skeleton-based Human Action Recognition using Ricci Curvature and Graphs Neural Networks: MVL Ribeiro et al. (MVL Ribeiro, 2025) View paper
  - Physical System Simulation (1 papers)
  - [20] PIORF: Physics-Informed Ollivier-Ricci Flow for Long-Range Interactions in Mesh Graph Neural Networks (Choi, 2025) View paper
  - Combinatorial Optimization and SAT Solving (1 papers)
  - [18] On the Hardness of Learning GNN-based SAT Solvers: The Role of Graph Ricci Curvature (Skenderi, 2025) View paper
- Surveys and Comparative Studies
  - Over-Squashing Survey Literature (2 papers)
  - [8] Over-Squashing in Graph Neural Networks: A Comprehensive survey (Akansha Singh, 2023) View paper
  - [26] Discrete Curvature and Applications in Graph Machine Learning (Weber, n.d.) View paper
- Related Methodologies
  - Information-Theoretic Approaches (1 papers)
  - [21] Discrete Curvature Graph Information Bottleneck (Fu, 2025) View paper
  - Alternative Bottleneck Detection Methods (2 papers)
  - [19] Adaptive Causal Discovery in Dynamic Web Architectures: A Real-Time Framework for Performance Bottleneck Identification (M Xu, 2025) View paper
  - [27] Are GNNs fundamentally bottlenecked? (Alexandru, n.d.) View paper
  - Curvature-Based Explainability (1 papers)
  - [28] Robust Explanations of Graph Neural Networks via Graph Curvatures (Y Liu, n.d.) View paper
  - Fractal and Alternative Graph Structures (1 papers)
  - [29] Fractal-Inspired Message Passing Neural Networks with Fractal Nodes (J Choi, n.d.) View paper

Narrative

Core task: detecting over-squashing in graph neural networks using discrete curvature. The field has organized itself around several complementary directions. Theoretical Foundations and Analysis establishes the mathematical underpinnings, connecting discrete curvature measures to information bottlenecks and characterizing when message-passing architectures fail to propagate signals effectively. Rewiring and Structural Modification develops algorithms that adjust graph topology—adding edges or reweighting connections—to alleviate bottlenecks identified by curvature diagnostics. Curvature-Enhanced GNN Architectures and Curvature-Based Structural Encodings explore how to integrate curvature information directly into model design, either by modulating message-passing operations or by treating curvature as an auxiliary feature. Domain-Specific Applications demonstrate these ideas in contexts ranging from skeleton-based action recognition to brain network analysis, while Surveys and Comparative Studies synthesize the landscape and Related Methodologies connect curvature-based approaches to broader graph learning techniques.

Within the theoretical branch, a handful of works have focused on formalizing the relationship between negative curvature and over-squashing. Curvature Bottlenecks[10] and Information Contraction Expansion[11] provide foundational analyses linking geometric properties to information flow, establishing that regions of low curvature can create severe bottlenecks. Rethinking Gold Standard[0] sits squarely in this cluster, revisiting earlier characterizations and proposing refined diagnostics for identifying problematic graph structures. Compared to Curvature Bottlenecks[10], which emphasizes worst-case bounds, Rethinking Gold Standard[0] appears to focus on practical detection criteria that balance sensitivity and computational cost. Meanwhile, works like Curvature Over-squashing[25] explore alternative curvature notions, highlighting ongoing debates about which geometric measures best predict message-passing failures. The interplay between rigorous theory and empirical validation remains a central open question across these studies.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Understanding over-squashing and bottlenecks on graphs via curvature

Authors: Topping, Jake, Jake Topping, Di Giovanni, Francesco, et al. (17 authors total) | **Year/Venue**: 2022 | **URL**: View paper

Abstract

Most graph neural networks (GNNs) use the message passing paradigm, in which node features are propagated on the input graph. Recent works pointed to the distortion of information flowing from distant nodes as a factor limiting the efficiency of message passing for tasks relying on long-distance interactions. This phenomenon, referred to as 'over-squashing', has been heuristically attributed to graph bottlenecks where the number of K-hop neighbors grows rapidly with K. We provide a precise d...

Relationship Analysis

Both papers belong to the Curvature-Based Over-Squashing Characterization category, providing formal analysis of how discrete curvature measures relate to bottlenecks and over-squashing in GNNs. The original paper critically examines the limitations of discrete curvature as a necessary condition for detecting over-squashing, constructing counterexamples where edges exhibit severe over-squashing despite positive curvature values, and proposes WAF3 as an improved curvature measure. In contrast, the candidate paper establishes the foundational connection between negative curvature and over-squashing, introducing Balanced Forman curvature and proving that negatively curved edges are sufficient to cause bottlenecks, while also proposing curvature-based graph rewiring methods to alleviate the problem.

---

2. Oversquashing in gnns through the lens of information contraction and graph expansion

Authors: Pradeep Kr. Banerjee, Pradeep Banerjee, Kedar Karhadkar, P. Banerjee, Yu Guang Wang, et al. (7 authors total) | **Year/Venue**: 2022 | **URL**: View paper

Abstract

The quality of signal propagation in message-passing graph neural networks (GNNs) strongly influences their expressivity as has been observed in recent works. In particular, for prediction tasks relying on long-range interactions, recursive aggregation of node features can lead to an undesired phenomenon called âoversquashingâ. We present a framework for analyzing oversquashing based on information contraction. Our analysis is guided by a model of reliable computation due to von Neumann that...

Relationship Analysis

Both papers belong to the Curvature-Based Over-Squashing Characterization category, analyzing how discrete curvature measures relate to bottlenecks and over-squashing in GNNs. While the original paper critically examines the limitations of discrete curvature as a necessary condition for detecting over-squashing and proposes WAF3 as an improved curvature measure, the candidate paper approaches over-squashing through an information-theoretic lens, analyzing it as signal decay in noisy computation graphs and proposing expander-based graph rewiring algorithms (G-RLEF) rather than new curvature definitions.

---

3. Curvature and over-squashing in Graph Neural Networks

Authors: F Di Giovanni | **URL**: View paper

Abstract

â¦ â¶ In principle, this can be fixed by choosing message passing functions that do not act as low-â¦ Intuition â¶ Swapping geodesics for edges, we can take discrete analogues on graphs â¦

Relationship Analysis

Both papers belong to the Curvature-Based Over-Squashing Characterization category, providing formal analysis of how discrete curvature measures characterize bottlenecks and over-squashing in GNNs. The candidate paper (a presentation/tutorial) overlaps with the original by discussing the relationship between negative curvature edges and over-squashing, introducing Balanced Forman curvature, and proposing graph rewiring solutions. However, the original paper critically challenges the necessity of negative curvature for over-squashing through counterexamples and proposes the novel WAF3 curvature with efficient approximation algorithms, while the candidate focuses on explaining existing curvature concepts and the SDRF rewiring method without questioning curvature's fundamental limitations.

Contributions Analysis

Overall novelty summary. The paper challenges the prevailing assumption that negative discrete curvature is both necessary and sufficient for detecting over-squashing in graph neural networks. It sits within the Curvature-Based Over-Squashing Characterization leaf, which contains four papers providing formal analysis of how discrete curvature measures characterize bottlenecks. This is a moderately populated research direction within the broader Theoretical Foundations and Analysis branch, suggesting active but not overcrowded investigation into the fundamental relationship between geometric properties and information flow in GNNs.

The taxonomy reveals closely related work in sibling leaves: Over-Squashing and Over-Smoothing Trade-offs examines interactions between competing failure modes, while Geometric and Topological Perspectives explores alternative embedding spaces beyond standard graphs. Neighboring branches address practical interventions—Curvature-Guided Rewiring Methods develops topology modification algorithms, and Curvature-Constrained Message Passing integrates curvature into neural architectures. The paper's theoretical focus on characterization boundaries distinguishes it from these application-oriented directions, though its proposed WAF3 measure bridges toward the rewiring and architecture branches.

Among twelve candidates examined, none clearly refute the three main contributions. The theoretical claim that negative curvature is insufficient examined two candidates without finding overlapping prior work. The MOSR metric examined ten candidates, again with no refutations detected. The WAF3 curvature measure examined zero candidates, suggesting either limited semantic overlap in the search or a genuinely novel formulation. Given the small search scope—twelve papers from a field with thirty surveyed works—these statistics indicate potential novelty but cannot rule out relevant prior work outside the top-ranked semantic matches.

The analysis covers a focused slice of the literature: top-ranked semantic neighbors plus citation expansion within a thirty-paper taxonomy. The absence of refutations across contributions suggests the work introduces fresh perspectives on curvature-based diagnostics, particularly the counterexamples demonstrating insufficiency and the WAF3 approximation algorithm. However, the limited search scope means closely related theoretical results or alternative curvature formulations may exist in venues or subfields not captured by this retrieval strategy.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Theoretical demonstration that discrete curvature is insufficient for detecting over-squashing

Description: The authors construct counterexample graphs and prove that edges can exhibit severe over-squashing while maintaining highly positive discrete curvature values across eight popular curvature definitions, demonstrating that curvature cannot fully capture all over-squashing phenomena in GNNs.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. The effectiveness of curvature-based rewiring and the role of hyperparameters in gnns revisited

URL: View paper

Brief Assessment

Curvature Rewiring Revisited[6] focuses on empirically evaluating whether edges selected by curvature-based rewiring satisfy theoretical conditions in real-world datasets, rather than proving theoretical insufficiency through counterexamples as the original paper does.

---

2. Understanding over-squashing and bottlenecks on graphs via curvature

URL: View paper

Brief Assessment

Curvature Bottlenecks[10] establishes that negatively curved edges cause over-squashing (sufficiency), but does not address whether curvature can miss over-squashed edges (necessity).

Contribution 2: Missed Over-Squashing Ratio (MOSR) metric

Description: A new metric that quantifies the proportion of over-squashed edges that are not identified by curvature-based methods. Extensive experiments show that Ollivier Ricci curvature can miss over 30% of over-squashed edges, revealing significant limitations in widely-used discrete curvature measures.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Over-Squashing in Graph Neural Networks: A Comprehensive survey

URL: View paper

Brief Assessment

Over-Squashing Survey[8] provides a comprehensive review of over-squashing in GNNs but does not introduce a metric quantifying the proportion of over-squashed edges missed by curvature-based methods. The survey discusses various mitigation strategies and theoretical foundations but does not propose MOSR or similar quantitative metrics for evaluating curvature method limitations.

2. FoSR: First-order spectral rewiring for addressing oversquashing in GNNs

URL: View paper

Brief Assessment

FoSR[31] focuses on spectral rewiring methods to address oversquashing through graph modification, not on developing metrics to quantify missed over-squashed edges by curvature-based methods.

3. PIORF: Physics-Informed Ollivier-Ricci Flow for Long-Range Interactions in Mesh Graph Neural Networks

URL: View paper

Brief Assessment

PIORF[20] focuses on physics-informed graph rewiring for mesh-based fluid dynamics simulations, not on developing metrics to quantify over-squashed edges missed by curvature methods in general GNNs.

4. On the Trade-off between Over-smoothing and Over-squashing in Deep Graph Neural Networks

URL: View paper

Brief Assessment

Smoothing Squashing Tradeoff[22] focuses on the spectral gap relationship between over-smoothing and over-squashing, proposing a curvature-based rewiring method. It does not introduce metrics to quantify missed over-squashed edges by curvature methods.

5. Curvdrop: A ricci curvature based approach to prevent graph neural networks from over-smoothing and over-squashing

URL: View paper

Brief Assessment

Curvdrop[2] focuses on using Ricci curvature for dropout sampling to address over-smoothing and over-squashing, but does not propose any metric to quantify the proportion of over-squashed edges missed by curvature-based methods.

6. Understanding over-squashing and bottlenecks on graphs via curvature

URL: View paper

Brief Assessment

Curvature Bottlenecks[10] does not propose any metric for quantifying the proportion of over-squashed edges missed by curvature-based methods.

7. DeepRicci: Self-supervised Graph Structure-Feature Co-Refinement for Alleviating Over-squashing

URL: View paper

Brief Assessment

DeepRicci[32] focuses on self-supervised graph structure-feature co-refinement using Riemannian geometry and Ricci flow, not on developing metrics to quantify missed over-squashed edges by curvature methods.

8. Mitigating Over-Smoothing and Over-Squashing using Augmentations of Forman-Ricci Curvature

URL: View paper

Brief Assessment

Forman-Ricci Augmentations[9] focuses on characterizing over-squashing and over-smoothing using augmented Forman-Ricci curvature and proposes rewiring techniques, but does not introduce a metric quantifying the proportion of over-squashed edges missed by curvature-based methods.

9. A Remedy for Over-Squashing in Graph Learning via Forman-Ricci Curvature based Graph-to-Hypergraph Structural Lifting

URL: View paper

Brief Assessment

Forman-Ricci Hypergraph[12] focuses on structural lifting from graphs to hypergraphs using Forman-Ricci curvature to address over-squashing, but does not propose a metric to quantify the proportion of over-squashed edges missed by curvature-based methods. The paper addresses over-squashing through architectural changes rather than evaluating curvature detection limitations.

10. Revisiting Over-smoothing and Over-squashing Using Ollivier-Ricci Curvature

URL: View paper

Brief Assessment

Ollivier-Ricci Revisited[16] focuses on establishing theoretical connections between Ollivier-Ricci curvature and over-smoothing/over-squashing phenomena, and proposes a rewiring algorithm (BORF). It does not introduce a metric quantifying the proportion of over-squashed edges missed by curvature-based methods.

Contribution 3: Weighted Augmented Forman-3 Curvature (WAF3) with efficient approximation algorithm

Description: A new discrete curvature definition that weights node contributions by degree to better detect over-squashed edges, achieving significantly lower MOSR values. The authors also develop a MinHash-based approximation algorithm that reduces time complexity to linear level, enabling computation on graphs with five million edges in 23.6 seconds with 133.7× speedup.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Appendix: Text Similarity Detection

Textual similarity detection checked 13 papers and found 2 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Understanding over-squashing and bottlenecks on graphs via curvature

Detected in: Core Task (sibling), Contribution: contribution_1, Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Rethinking the Gold Standard: Why Discrete Curvature Fails to Fully Capture Over-squashing in GNNs? View paper
- [1] Rewiring networks for graph neural network training using discrete geometry View paper
- [2] Curvdrop: A ricci curvature based approach to prevent graph neural networks from over-smoothing and over-squashing View paper
- [3] Skeleton-based Human Action Recognition using Ricci Curvature and Graphs Neural Networks: MVL Ribeiro et al. View paper
- [4] Effects of Random Edge-Dropping on Over-Squashing in Graph Neural Networks View paper
- [5] Curvature constrained MPNNs: Improving message passing with local structural properties View paper
- [6] The effectiveness of curvature-based rewiring and the role of hyperparameters in gnns revisited View paper
- [7] Signed Curvature Graph Representation Learning of Brain Networks for Brain Age Estimation View paper
- [8] Over-Squashing in Graph Neural Networks: A Comprehensive survey View paper
- [9] Mitigating Over-Smoothing and Over-Squashing using Augmentations of Forman-Ricci Curvature View paper
- [10] Understanding over-squashing and bottlenecks on graphs via curvature View paper
- [11] Oversquashing in gnns through the lens of information contraction and graph expansion View paper
- [12] A Remedy for Over-Squashing in Graph Learning via Forman-Ricci Curvature based Graph-to-Hypergraph Structural Lifting View paper
- [13] How curvature enhance the adaptation power of framelet gcns View paper
- [14] Over-Squashing in Riemannian Graph Neural Networks View paper
- [15] Curvature-based Pooling within Graph Neural Networks View paper
- [16] Revisiting Over-smoothing and Over-squashing Using Ollivier-Ricci Curvature View paper
- [17] Revisiting Over-smoothing and Over-squashing using Ollivier's Ricci Curvature View paper
- [18] On the Hardness of Learning GNN-based SAT Solvers: The Role of Graph Ricci Curvature View paper
- [19] Adaptive Causal Discovery in Dynamic Web Architectures: A Real-Time Framework for Performance Bottleneck Identification View paper
- [20] PIORF: Physics-Informed Ollivier-Ricci Flow for Long-Range Interactions in Mesh Graph Neural Networks View paper
- [21] Discrete Curvature Graph Information Bottleneck View paper
- [22] On the Trade-off between Over-smoothing and Over-squashing in Deep Graph Neural Networks View paper
- [23] Effective Structural Encodings via Local Curvature Profiles View paper
- [24] A Rewiring Contrastive Patch PerformerMixer Framework for Graph Representation Learning View paper
- [25] Curvature and over-squashing in Graph Neural Networks View paper
- [26] Discrete Curvature and Applications in Graph Machine Learning View paper
- [27] Are GNNs fundamentally bottlenecked? View paper
- [28] Robust Explanations of Graph Neural Networks via Graph Curvatures View paper
- [29] Fractal-Inspired Message Passing Neural Networks with Fractal Nodes View paper
- [30] Curvature MPNNs: Improving Message Passing with Local Structural Properties View paper
- [31] FoSR: First-order spectral rewiring for addressing oversquashing in GNNs View paper
- [32] DeepRicci: Self-supervised Graph Structure-Feature Co-Refinement for Alleviating Over-squashing View paper