Convergent Privacy Framework for Multi-layer GNNs through Contractive Message Passing

Yu Zheng*, Chenang Li*, Zhou Li* and Qingsong Wang†
* University of California, Irvine, Email: {yu.zheng, chenangl, zhou.li}@uci.edu
† University of California, San Diego, Email: qiw072@ucsd.edu

Abstract

Differential privacy (DP) has been integrated into graph neural networks (GNNs) to protect sensitive structural information, e.g., edges, nodes, and associated features across various applications. A prominent approach is to perturb the message-passing process, which forms the core of most GNN architectures. However, existing methods typically incur a privacy cost that grows linearly with the number of layers (e.g., GAP published in Usenix Security'23), ultimately requiring excessive noise to maintain a reasonable privacy level. This limitation becomes particularly problematic when multi-layer GNNs, which have shown better performance than one-layer GNN, are used to process graph data with sensitive information.

In this paper, we theoretically establish that the privacy budget converges with respect to the number of layers by applying privacy amplification techniques to the message-passing process, exploiting the contractive properties inherent to standard GNN operations. Motivated by this analysis, we propose a simple yet effective *Contractive Graph Layer (CGL)* that ensures the contractiveness required for theoretical guarantees while preserving model utility. Our framework, CARIBOU, supports both training and inference, equipped with a contractive aggregation module, a privacy allocation module, and a privacy auditing module. Experimental evaluations demonstrate that CARIBOU significantly improves the privacy-utility trade-off and achieves superior performance in privacy auditing tasks.

Acknowledgment

This work [1] has been accepted for publication in the Proceedings of the 33rd Network and Distributed System Security Symposium (NDSS) 2026. The official version is available via DOI: https://dx.doi.org/10.14722/ndss.2026.240255. For extended proofs and technical details, please refer to the full version on https://arxiv.org/abs/2506.22727.

References

[1] Y. Zheng, C. Li, Z. Li, and Q. Wang, "Convergent privacy framework for multi-layer gnns through contractive message passing," in *Proceedings of the 33rd Network and Distributed System Security Symposium (NDSS)*, 2026.

Convergent Privacy Framework for Multi-layer GNNs through Contractive Message Passing

Yu Zheng*, Chenang Li*, Zhou Li*, and Qingsong Wang#

* University of California, Irvine # University of California, San Diego

Summary of Contribution

✦ A novel privacy analysis for GNNs that leverages the **contractiveness** of message passing to achieve the ***convergent*** privacy costs.

✦ New design of *perturbed graph contractive layer* and a practical private GNN framework **CARIBOU**.

Full-version Paper

Code

Artifact Evaluated NDSS

Available

Functional

Reproduced

Scenarios & Motivation

Perturbed Message Passing with Differential Privacy (DP)

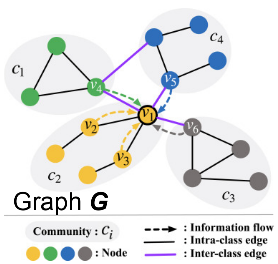

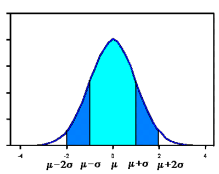

G,G': Edge/node-level neighboring.

Formulation: $X^{(k+1)} = \Pi(\mathrm{MP}_G(X^{(k)}) + Z^{(k)})$

Node representation/embeddings from message passing (MP); $X^{(0)}$: input feature matrix.

Gaussian noise $\sim N(0, \sigma^2)$.

Framework	Mechanism	Noise (σ^2)	Utility
PertGraph [1,2]	*G* Pert.	$\propto 1$	Fair
DPDGC [3]	Decoupled *G* w/ Pert.	$\propto K$	Good
GAP [4]	Pert. MP	$\propto K$	Good

✦ Share a **critical limitation**: privacy loss grows linearly with K.

✦ Require large σ to maintain a reasonable privacy level for multi-layer GNNs, degrading utility.

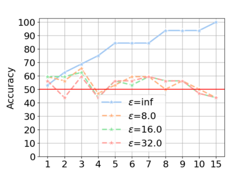

Convergent Privacy Budget

✦ **Insight**: leverage the inherent **privacy amplification** that occurs in multi-layer GNNs through contractiveness.

❖ When perturbed MP is contractive, the distance between GNNs trained on neighboring ***G,G'*** shrinks at each step.

❖ Consequently, the influence of individual data points diminishes, leading to the amplified privacy rooted from "over-smoothing."

✓ Remove the over-estimated privacy loss.

✦ Derive a much **tighter bound** for the finally released GNN model.

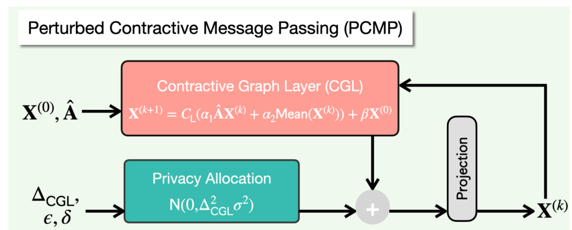

CARIBOU Framework

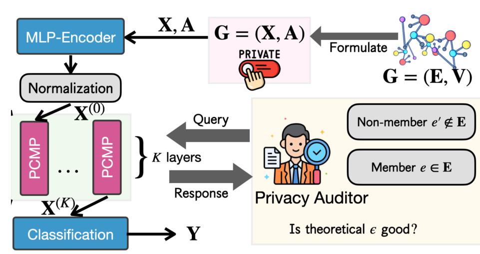

✦ **Theorem 1 [DP guarantee for CGL layers]**. Let *G* be a graph and *K* be the number of contractive graph layers in CARIBOU. Let $C_L < 1$ be Lipschitz constant. Then, the *K*-hop message passing of CARIBOU satisfies: $(\frac{\alpha}{2}\frac{\Delta^2}{\sigma^2}\min\{K, \frac{1-C_L^K}{1+C_L^K}\frac{1+C_L}{1-C_L}\} + \frac{\log(1/\delta)}{\alpha-1}, \delta)$-DP, where σ is the noise scale, Δ is the sensitivity of MP, and α, δ are DP parameters.

Experiments

Accuracy: EDP over the Cora dataset.

ϵ	1	2	4	8	16	32
CARIBOU	85%	87%	87%	88%	89%	89%
GAP	76%	78%	75%	76%	78%	80%

Accuracy: NDP over the Cora dataset.

ϵ	1	2	4	8	16	32
CARIBOU	81%	83%	86%	87%	88%	88%
GAP	34%	32%	32%	44%	56%	64%

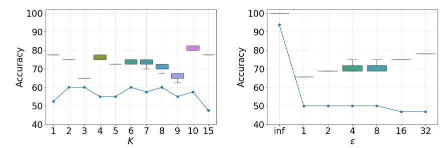

(a) Different K when $\epsilon = 4$ (b) Different ϵ when $K = 10$

Fig. 5: Comparison between CARIBOU (colored boxes) and GAP (blue lines) for ablation study.

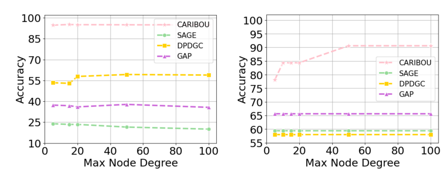

(a) $\epsilon = 2$ (Photo) (b) $\epsilon = 2$ (Chain-S)

Fig. 6: NDP Accuracy with Varying Max Node Degree.

Acknowledge & Reference

[1] A. Kolluri, T. Baluta, B. Hooi, and P. Saxena, "Lpgnet: Link private graph networks for node classification," in ACM SIGSAC Conference on Computer and Communications Security, CCS, 2022, pp. 1813–1827.

[2] F. Wu, Y. Long, C. Zhang, and B. Li, "LINKTELLER: recovering private edges from graph neural networks via influence analysis," in IEEE Symposium on Security and Privacy, SP, 2022, pp. 2005–2024.

[3] E. Chien, W.-N. Chen, C. Pan, P. Li, A. Ozgur, and O. Milenkovic, "Differentially private decoupled graph convolutions for multigranular topology protection," in Advances in Neural Information Processing Systems, vol. 36, 2023.

[4] S. Sajadmanesh, A. S. Shamsabadi, A. Bellet, and D. Gatica-Perez, "Gap: Differentially private graph neural networks with aggregation perturbation," in USENIX Security 2023-32nd USENIX Security Symposium, 2023.