POST-CALIBRATION TOPOLOGY DRIFT REVEALS OPERATOR-DEPENDENT NUMERICAL SENSITIVITY IN QUANTIZED GRAPH NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Post-training quantization (PTQ) is attractive for graph neural networks (GNNs), yet it is commonly evaluated on the same topology used for calibration. We study whether post-calibration topology drift changes the numerical discrepancy between a quantized GNN and its full-precision counterpart beyond the error already present on the calibration graph. We introduce Component-Preserving Degree-Preserving Rewiring (CP-DPR), a controlled intervention that changes neighborhood composition while preserving the node set, features, labels, model weights, edge count, exact node-wise degree sequence, and connected-component partition. We measure the baseline-adjusted excess logit divergence $\mathrm{EQD}_Z = \mathrm{QD}_Z(A_\delta) - \mathrm{QD}_Z(A_0)$. Experiments cover Cora, CiteSeer, PubMed, and Roman-empire; GCN and GraphSAGE; W8A8 and W4A4 simulated PTQ; five training seeds; and ten perturbation seeds per severity. At the strongest perturbation level, 10 of 12 external dataset–operator–precision settings exhibit positive excess divergence, with substantially larger absolute effects under W4A4. Roman-empire with GCN provides a consistent negative case that persists after validation-only model refinement. Activation diagnostics show negligible additional clipping in the dominant positive regimes and a stronger association with rounding mismatch. GraphSAGE further exposes an observer-blindness failure mode in which min–max calibration parameters remain unchanged despite internal activation-distribution motion. These results provide evidence that topology drift induces operator-dependent numerical sensitivity not captured by calibration-graph accuracy alone, while the evidence for operator dependence currently rests on one heterophilous benchmark.

1 INTRODUCTION

Graph neural networks (GNNs) are deployed in settings where relational structure can evolve: citation links appear, transactions create interactions, and communication patterns change over time. Because topology is part of the inference operator, changing the edge set changes which representations are aggregated and therefore the hidden activations produced by a fixed trained model.

Low-precision inference is an important route to reducing the computational and memory cost of GNN deployment. Existing GNN quantization methods already exploit graph structure. Degree-Quant uses degree-aware quantization-aware training (Tailor et al., 2021); $\mathrm{A}^2\mathrm{Q}$ uses aggregation behavior to learn mixed precision (Zhu et al., 2023); MixQ-GNN broadens mixed-precision design across GNN computations (Moustafa et al., 2025); and TopGQ targets topology-aware PTQ through TopPIN-based grouping (Kwon et al., 2026). These works establish that graph structure matters for quantization.

A separate literature studies predictive robustness under structural perturbations and graph distribution shift (Zügner & Günnemann, 2020; Fan et al., 2024; Zhang et al., 2026a). Our question is narrower. Suppose a model is trained in FP32 and calibrated for PTQ on a reference graph G_0, after which the quantization parameters are frozen. If deployment topology changes to G_δ, does the quantized model remain as close to its FP32 counterpart as it was on G_0? We do not ask whether topology drift changes the GNN output—that is expected. We ask whether it changes the *additional numerical discrepancy introduced by quantization*.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Arbitrary edge perturbations confound neighborhood composition with degree, normalization, component structure, and other graph statistics. We therefore introduce Component-Preserving Degree-Preserving Rewiring (CP-DPR), which changes connections while preserving every node degree, the total edge count, and the exact connected-component partition. For GCN, degree preservation also fixes the degree-derived normalization factors, leaving the identities of aggregated neighbors as the principal manipulated quantity.

Let $\mathrm{QD}_Z(A)$ denote relative FP32–PTQ logit divergence under adjacency A. Our primary outcome is

$$\mathrm{EQD}_Z(\delta) = \mathrm{QD}_Z(A_\delta) - \mathrm{QD}_Z(A_0). \tag{1}$$

A positive value means that topology drift increases FP32–PTQ separation beyond the quantizer's baseline error; a negative value means that the separation decreases.

Across the controlled benchmark, Cora shows a severity-dependent increase in EQD_Z, and 10 of 12 external dataset–architecture–precision settings are positive at the strongest perturbation level. Effects are substantially larger under W4A4 than W8A8. The pattern is not universal: Roman-empire with GCN exhibits a stable sign reversal. Mechanism diagnostics further indicate that the dominant positive regime is associated more strongly with rounding mismatch than clipping, and GraphSAGE reveals that min–max calibration can remain unchanged while the activation distribution moves internally.

Our contributions are: (i) a formulation of post-calibration topology-induced quantization drift and the baseline-adjusted EQD_Z metric; (ii) CP-DPR, a controlled protocol preserving exact degree sequence, edge count, and component partition; (iii) cross-dataset, cross-architecture, and cross-precision evidence that topology drift often amplifies low-bit numerical divergence, together with a reproducible negative boundary case that provides evidence of architecture dependence; and (iv) deployment-relevant diagnostics showing rounding-dominated mismatch and min–max observer blindness without reliable correction from recalibration.

2 Background and Related Work

Topology and message passing. A broad class of GNNs aggregates messages from a node's neighborhood (Gilmer et al., 2017); consequently, changing adjacency changes the values entering the aggregation operator even when features and learned parameters remain fixed. In GCN (Kipf & Welling, 2017), propagation uses $\widetilde{D}^{-1/2}\widetilde{A}\widetilde{D}^{-1/2}$, so naive edge perturbations can change both neighborhood membership and degree-dependent normalization. GraphSAGE with mean aggregation (Hamilton et al., 2017) uses a different update construction, making architecture-specific activation responses to the same structural intervention plausible. Detailed equations are given in Appendix A.

PTQ and calibration. Affine quantization maps floating-point values to a finite grid through a scale and zero point (Jacob et al., 2018). In PTQ, weights are quantized after training and activation quantizers are determined from calibration statistics (Nagel et al., 2021). Our setting therefore creates a frozen numerical interface: calibration observes activations on G_0, whereas deployment may later evaluate the same weights and quantizer on G_δ. A topology shift can move activations within a fixed quantization range, changing rounding error even when clipping does not increase. Appendix A gives the quantizer and clipping/rounding decomposition used in our diagnostics.

GNN quantization. SGQuant introduced graph-specific multi-granularity quantization (Feng et al., 2020); Degree-Quant linked quantization sensitivity to node degree (Tailor et al., 2021); $\mathrm{A}^2\mathrm{Q}$ uses aggregation-aware mixed precision and studies transfer to unseen graphs (Zhu et al., 2023); MEGA combines degree-aware mixed precision with accelerator co-design (Zhu et al., 2024); and MixQ-GNN expands mixed precision across message-passing computations (Moustafa et al., 2025). TopGQ is the closest recent work: it explicitly targets GNN PTQ with topology-aware grouping (Kwon et al., 2026). Thus, topology-aware quantization itself is not our novelty. We study what happens *after* calibration when topology changes.

Structural robustness and graph OOD. Structural-robustness work treats edge changes as perturbations to the predictive model (Zügner & Günnemann, 2020). Graph OOD work studies broader development-to-deployment shifts and invariant, adaptive, or extrapolative strategies (Fan et al., 2024;

Li et al., 2024a;b; Zhang et al., 2026a). CauVQ uses vector quantization as a representation-learning mechanism for graph OOD generalization (Zhang et al., 2026b), whereas we study finite-precision weights and activations at inference. These works primarily ask whether predictions remain reliable under shift; we instead compare PTQ and FP32 *on the same shifted graph* to isolate the quantization-specific component.

Controlled rewiring and gap. Degree-preserving edge switching modifies structure while holding the degree sequence fixed (Carstens & Horadam, 2017), but it can still merge or split connected components. CP-DPR strengthens this control by preserving the exact component partition as well. The resulting gap is narrow but distinct: whether a low-bit PTQ model develops additional numerical sensitivity when topology changes after calibration, including when degree-derived normalization and connected-component membership are held fixed.

3 METHODOLOGY

3.1 POST-CALIBRATION TOPOLOGY DRIFT

Let $G_0 = (V, E_0, X, Y)$ be the graph used for training and calibration and let f_θ be the trained FP32 GNN. PTQ calibration on G_0 produces quantization parameters ϕ_0. Deployment changes only adjacency, $A_0 \to A_\delta$, while V, X, Y, θ and, under the primary static condition, ϕ_0 remain fixed. Define

$$Z_{\mathrm{FP}}(A) = f_\theta(X, A), \qquad Z_Q(A; \phi_0) = f_\theta^Q(X, A; \phi_0), \tag{2}$$

$$\mathrm{QD}_Z(A) = \frac{\|Z_{\mathrm{FP}}(A) - Z_Q(A; \phi_0)\|_F}{\|Z_{\mathrm{FP}}(A)\|_F + \varepsilon}. \tag{3}$$

Our primary endpoint is the baseline-adjusted excess divergence

$$\boxed{\mathrm{EQD}_Z(\delta) = \mathrm{QD}_Z(A_\delta) - \mathrm{QD}_Z(A_0)}. \tag{4}$$

For task performance, we analogously define

$$\mathrm{EQD}_M(\delta) = \big[M_{\mathrm{FP}}(A_\delta) - M_Q(A_\delta)\big] - \big[M_{\mathrm{FP}}(A_0) - M_Q(A_0)\big]. \tag{5}$$

3.2 COMPONENT-PRESERVING DEGREE-PRESERVING REWIRING

A degree-preserving double-edge swap changes adjacency without changing the degree of affected nodes. CP-DPR adds the hard requirement that the exact connected-component partition also remain unchanged. Every accepted graph satisfies

$$V_\delta = V_0, \quad X_\delta = X_0, \quad Y_\delta = Y_0, \quad |E_\delta| = |E_0|, \quad d_\delta(v) = d_0(v)\ \forall v, \quad \mathcal{C}(G_\delta) = \mathcal{C}(G_0). \tag{6}$$

Rewiring is performed within the original largest connected component; edges outside it remain unchanged. Candidate graphs are accepted only after exact verification of degree equality, edge-count equality, component-partition equality, graph simplicity, and target severity. Severity is edge-set Jaccard distance,

$$\delta_E = 1 - \frac{|E_0 \cap E_\delta|}{|E_0 \cup E_\delta|}, \tag{7}$$

with $\delta_E \in \{0.10, 0.20, 0.40\}$ in the main controlled benchmark.

3.3 QUANTIZATION AND DIAGNOSTICS

We use simulated PTQ with W8A8 and W4A4. Weights are quantized per output channel with symmetric quantization; activations use per-tensor affine quantization. Biases and final logits remain FP32, and adjacency is not quantized. Activation fake quantization is applied to the model input and to the hidden representation after the first message-passing layer and ReLU. Calibration uses G_0 and all available node activations without labels; deployment freezes those activation quantization parameters under G_δ.

Table 1: Cora under CP-DPR. Median EQD_Z with crossed-bootstrap 95% confidence intervals.

Condition	$\delta_E = 0.10$	$\delta_E = 0.20$	$\delta_E = 0.40$
W8A8	.000164 [.000136,.000193]	.000379 [.000295,.000411]	.000898 [.000687,.000964]
W4A4	.002362 [.002003,.002960]	.005696 [.004816,.006788]	.014040 [.011850,.015777]

Target-topology recalibration recomputes hidden-activation quantization parameters on G_δ while holding weights fixed. Its shift-specific benefit is

$$\mathrm{Gain}_{\mathrm{Recal}}(\delta) = \mathrm{EQD}_Z^{\mathrm{Static}}(\delta) - \mathrm{EQD}_Z^{\mathrm{Recal}}(\delta). \tag{8}$$

We record clipping fraction, clipping error, and rounding MSE, and compare min–max calibration with a percentile observer using $[P_{0.1}, P_{99.9}]$ bounds. The percentile observer is diagnostic rather than a proposed quantization method.

3.4 Datasets, Architectures, and Statistical Design

We use Cora, CiteSeer, PubMed, and Roman-empire. The first three are standard citation-network benchmarks associated with Planetoid-style evaluation (Yang et al., 2016); Roman-empire is a heterophilous benchmark introduced by Platonov et al. (2023). We evaluate two-layer GCN and mean-aggregation GraphSAGE. The default configuration uses hidden dimension 64, ReLU, dropout 0.5, Adam with learning rate 0.01 and weight decay 5×10^{-4}, and validation-loss early stopping.

For each dataset–architecture setting, five independently trained models are crossed with ten independently generated perturbation graphs per severity, giving 50 model–graph units per condition. Training seeds are 1101, 2202, 3303, 4404, and 5505; perturbation seeds are 101, 202, 303, 404, 505, 606, 707, 808, 909, and 1010. Roman-empire+GCN receives an additional validation-only refinement because all original runs reached the 300-epoch ceiling: a predefined four-configuration grid is trained for up to 1000 epochs with patience 150, with the test split excluded from configuration selection.

The primary summary is median EQD_Z. Because training seed and perturbation realization are crossed sources of variation, uncertainty is estimated with a crossed bootstrap: each replicate resamples training seeds and perturbation seeds independently with replacement, forms their Cartesian product, and recomputes the median. We report percentile 95% confidence intervals. Mechanism correlations are Spearman associations and are interpreted descriptively, not causally. We omit the previously used pair-level Wilcoxon tests from the main inferential analysis because the 50 crossed model–graph pairs are not independent observational units.

4 Results

4.1 Cora Survives Degree and Component Controls

All 30 Cora CP-DPR graphs preserve the exact degree sequence and connected-component partition. Under W8A8, median EQD_Z rises from 0.000164 at $\delta_E = 0.10$ to 0.000898 at 0.40; W4A4 follows the same monotonic pattern at larger magnitude, from 0.002362 to 0.014040. Crossed-bootstrap intervals remain entirely above zero at every severity (Table 1).

The corresponding comparison with unconstrained-component degree-preserving rewiring shows only small differences in effect magnitude, indicating that component preservation removes a confound rather than the phenomenon. On the calibration graph, mean baseline divergence across the five Cora/GCN models is $\mathrm{QD}_Z(A_0) = 0.00471 \pm 0.00049$ for W8A8 and 0.0885 ± 0.0080 for W4A4. Thus, the excess at $\delta_E = 0.40$ is approximately 19% and 16% of the corresponding baseline error, respectively. This relative-scale interpretation is specific to Cora because external absolute baselines were not recorded.

Task-level check. On the separate Cora degree-preserving rewiring bank without the component constraint, W8A8 EQD_M remains small and its crossed-bootstrap interval includes zero at every severity. W4A4 medians are positive and increase with severity, but the bootstrap interval excludes

Table 2: External validation at $\delta_E = 0.40$. Entries are median EQD_Z [crossed-bootstrap 95% CI].

Dataset	Architecture	W8A8	W4A4
CiteSeer	GCN	.000920 [.000856,.000984]	.015979 [.012569,.017743]
CiteSeer	GraphSAGE	.000368 [.000309,.000467]	.007454 [.006328,.008576]
PubMed	GCN	.001254 [.001154,.001315]	.013779 [.011890,.016633]
PubMed	GraphSAGE	.000822 [.000769,.000899]	.010800 [.009194,.011713]
Roman-empire	GCN	-.000432 [-.000495,-.000325]	-.005031 [-.006804,-.003038]
Roman-empire	GraphSAGE	.000480 [.000426,.000516]	.004918 [.004434,.005861]

zero only at $\delta_E = 0.05$; at larger severities the uncertainty intervals include zero. We therefore treat the task-level result as suggestive rather than as uniform statistical confirmation. Full values are reported in Appendix C.

4.2 External Validation and Precision Dependence

At $\delta_E = 0.40$, 10 of 12 external dataset–architecture–precision settings have positive EQD_Z with crossed-bootstrap intervals above zero (Table 2). The two exceptions are Roman-empire with GCN, for which both precision settings are robustly negative. The positive pattern therefore extends beyond one citation graph and beyond GCN normalization, while the sign reversal prevents a universal degradation claim.

In every positive setting, W4A4 yields much larger excess divergence than W8A8. At $\delta_E = 0.40$, the W4A4/W8A8 ratio is approximately 17.4 for CiteSeer-GCN, 20.3 for CiteSeer-GraphSAGE, 11.0 for PubMed-GCN, 13.1 for PubMed-GraphSAGE, and 10.2 for Roman-empire-GraphSAGE. Because EQD_Z is baseline adjusted, this compares the topology-induced *change* in FP32–PTQ discrepancy rather than only the absolute quantization error.

4.3 Roman-Empire GCN Is a Stable Negative Boundary Case

Roman-empire with GCN is the only systematic sign reversal. We test whether this is an artifact of weak training by applying validation-selected refinement. Test accuracy and Macro-F1 improve for all five training seeds, yet the sign remains negative. Refined W8A8 median EQD_Z values are -0.000102, -0.000223, and -0.000414 across increasing severity; W4A4 gives -0.001197, -0.003765, and -0.006579. At $\delta_E = 0.40$, the W4A4 crossed-bootstrap interval is $[-0.008585, -0.004939]$. Several best checkpoints remain near the 1000-epoch ceiling, so we do not claim formal convergence; the supported conclusion is only that the negative regime is not explained solely by the original 300-epoch limit.

4.4 Mechanism Diagnostics: Rounding, Clipping, and Observer Blindness

Across the external benchmark, median excess clipping fraction is zero in every dataset–architecture–precision setting. Excess rounding MSE is instead strongly associated with EQD_Z in most positive regimes: pooled Spearman correlations are approximately 0.885/0.888 for CiteSeer-GCN (W8A8/W4A4), 0.892/0.832 for CiteSeer-GraphSAGE, and 0.868/0.779 for Roman-empire-GraphSAGE. This pattern is consistent with grid mismatch rather than a clipping-dominated explanation, but the association is diagnostic rather than causal.

GraphSAGE reveals a sharper observer-blindness case. On CiteSeer, min–max hidden quantization parameters remain exactly unchanged for every shifted graph in both W8A8 and W4A4 despite increased EQD_Z, so target min–max recalibration provides exactly zero median mitigation. At $\delta_E = 0.40$, the mean exact-equality fraction of GraphSAGE target qparams across the external datasets is approximately 0.647 for min–max calibration versus 0.037 for percentile calibration. Percentile recalibration detects distribution movement more often, but robust positive mitigation is supported in only one of six GraphSAGE dataset–precision settings; CiteSeer-W4A4 improves, whereas CiteSeer-W8A8 and PubMed-W8A8 consistently worsen. Detecting distribution motion and reducing FP32–PTQ divergence are therefore distinct objectives.

5 Discussion and Limitations

The main result is not that topology drift universally "breaks" quantized GNNs. Rather, it changes the numerical relationship between a low-bit model and its FP32 reference even after obvious structural confounds are removed. Exact degree preservation fixes connectivity magnitude and, for GCN, degree-derived normalization; exact component preservation prevents the effect from being driven by reachability changes between disconnected regions. The persistence of EQD_Z under these controls shows that the response cannot be reduced to degree drift alone.

The larger W4A4 response is consistent with a coarser quantization grid making structurally induced activation motion more consequential for rounding. The diagnostics support this interpretation: clipping barely changes in the dominant positive regimes, whereas rounding-related error tracks EQD_Z. GraphSAGE further shows that range stability is not distribution stability: min–max qparams can remain unchanged while internal occupancy moves and numerical divergence grows. More responsive percentile observers detect that motion, but recalibration does not reliably correct it.

Roman-empire is equally important because it prevents over-generalization. On the same graph, GraphSAGE has positive excess divergence while GCN has a robust negative effect. This provides evidence that the sign depends on an interaction among graph structure, message-passing operator, precision, and calibration. However, architecture dependence is currently evidenced by one heterophilous benchmark and two architectures; it should not yet be interpreted as a general taxonomy of GNN operators.

The study also uses simulated rather than hardware-specific PTQ, covers node classification with only GCN and GraphSAGE, leaves adjacency in full precision, and perturbs topology synthetically through rewiring. CP-DPR preserves degrees and components but not higher-order properties such as homophily, clustering, communities, motifs, or spectra. Mechanism results are associational, task-level EQD_M is evaluated only on Cora and under the unconstrained-component rewiring bank, and absolute baseline $\mathrm{QD}_Z(A_0)$ is available only for Cora/GCN. These constraints delimit the present claims; Appendix D provides the full limitation list.

From a deployment perspective, the results motivate three checks: PTQ should not be validated only on the calibration graph when connectivity is expected to evolve; low-bit settings deserve particular scrutiny; and degree statistics or min–max activation ranges alone are insufficient monitoring signals. Robustness of the quantizer to plausible structural alternatives should be evaluated separately from both topology-aware quantizer design and conventional predictive robustness.

6 Conclusion

We isolate a deployment effect that standard PTQ evaluation can miss: a quantizer is calibrated on one graph while the message-passing topology later changes. Using EQD_Z and CP-DPR, we separate topology-induced change in FP32–PTQ discrepancy from baseline quantization error while holding node features, labels, model weights, edge count, degree sequence, and connected-component membership fixed. Across four datasets, two architectures, and two low-bit regimes, topology drift frequently increases divergence and does so much more strongly at 4 bits. The effect is not universal: Roman-empire+GCN exhibits a stable sign reversal, providing evidence of architecture-dependent numerical sensitivity rather than monotonic degradation. The activation analysis is more consistent with rounding mismatch than clipping in the dominant positive regime and shows that min–max observers can miss meaningful internal distribution motion. These findings motivate explicit post-calibration structural-stability evaluation whenever graph connectivity may evolve.

AI use statement

Generative AI tools were used during manuscript preparation to provide feedback on statistical presentation and experimental methodology, verify bibliographic metadata, improve wording and readability, and assist with LaTeX/template formatting. In this editing process, the tools did not generate or alter the reported experimental measurements. All AI-assisted changes were reviewed by the author, who takes responsibility for the final text, claims, citations, and artifacts.

Reproducibility statement

The main paper specifies the datasets, architectures, quantization configurations, perturbation severities, training and perturbation seeds, primary metric, and crossed-bootstrap design. Appendix B consolidates implementation-level details available in the manuscript source, and Appendix C reports the task-level Cora results. Any accompanying code or supplementary implementation details should be anonymized before submission.

References

C. J. Carstens and K. J. Horadam. Switching edges to randomize networks: What goes wrong and how to fix it. *Journal of Complex Networks*, 5(3):337–351, 2017. doi: 10.1093/comnet/cnw027.

Shaohua Fan, Xiao Wang, Chuan Shi, Peng Cui, and Bai Wang. Generalizing graph neural networks on out-of-distribution graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(1):322–337, 2024. doi: 10.1109/TPAMI.2023.3321097.

Boyuan Feng, Yuke Wang, Xu Li, Shu Yang, Xueqiao Peng, and Yufei Ding. SGQuant: Squeezing the last bit on graph neural networks with specialized quantization. In *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 1044–1052. IEEE, 2020. doi: 10.1109/ICTAI50040.2020.00198.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017.

William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, volume 30, pp. 1024–1034, 2017.

Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2704–2713. IEEE, 2018. doi: 10.1109/CVPR.2018.00286.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.

Dain Kwon, Kanghyun Choi, Hyeyoon Lee, Sunjong Park, Seoyong Lee, Sukjin Kim, and Jinho Lee. TopGQ: Fast GNN post-training quantization leveraging topology information. *arXiv preprint arXiv:2608.30394*, 2026.

Haoyang Li, Xin Wang, Zeyang Zhang, Haibo Chen, Ziwei Zhang, and Wenwu Zhu. Disentangled graph self-supervised learning for out-of-distribution generalization. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 28890–28904. PMLR, 2024a.

Xiner Li, Shurui Gui, Youzhi Luo, and Shuiwang Ji. Graph structure extrapolation for out-of-distribution generalization. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 27846–27874. PMLR, 2024b.

Samir Moustafa, Nils M. Kriege, and Wilfried N. Gansterer. Efficient mixed precision quantization in graph neural networks. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, pp. 4038–4052. IEEE, 2025. doi: 10.1109/ICDE65448.2025.00301.

Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021.

Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *International Conference on Learning Representations*, 2023.

Shyam A. Tailor, Javier Fernandez-Marques, and Nicholas D. Lane. Degree-Quant: Quantization-aware training for graph neural networks. In *International Conference on Learning Representations*, 2021.

Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 40–48. PMLR, 2016.

Kexin Zhang, Shuhan Liu, Song Wang, Weili Shi, Chen Chen, Pan Li, Sheng Li, Jundong Li, and Kaize Ding. A survey of deep graph learning under distribution shifts: From graph out-of-distribution generalization to adaptation. *ACM Transactions on Knowledge Discovery from Data*, 20(2):1–38, 2026a. doi: 10.1145/3785475.

Weihong Zhang, Liang Bai, Hangyuan Du, and Xian Yang. CauVQ: Causal vector quantization for graph OOD generalization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(33): 28409–28417, 2026b. doi: 10.1609/aaai.v40i33.40070.

Zeyu Zhu, Fanrong Li, Zitao Mo, Qinghao Hu, Gang Li, Zejian Liu, Xiaoyao Liang, and Jian Cheng. A2Q: Aggregation-aware quantization for graph neural networks. In *International Conference on Learning Representations*, 2023.

Zeyu Zhu, Fanrong Li, Gang Li, Zejian Liu, Zitao Mo, Qinghao Hu, Xiaoyao Liang, and Jian Cheng. MEGA: A memory-efficient GNN accelerator exploiting degree-aware mixed-precision quantization. In *2024 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 124–138. IEEE, 2024. doi: 10.1109/HPCA57654.2024.00020.

Daniel Zügner and Stephan Günnemann. Certifiable robustness of graph convolutional networks under structure perturbations. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1656–1665. ACM, 2020. doi: 10.1145/3394486.3403217.

A Background Details

A generic message-passing layer can be written as

$$m_v^{(\ell)} = \mathrm{AGG}^{(\ell)}\Big(\Big\{M^{(\ell)}(h_v^{(\ell)}, h_u^{(\ell)}) : u \in N(v)\Big\}\Big), \tag{9}$$

$$h_v^{(\ell+1)} = U^{(\ell)}\Big(h_v^{(\ell)}, m_v^{(\ell)}\Big), \tag{10}$$

following Gilmer et al. (2017). For GCN (Kipf & Welling, 2017),

$$H^{(\ell+1)} = \sigma\Big(\widetilde{D}^{-1/2}\widetilde{A}\widetilde{D}^{-1/2}H^{(\ell)}W^{(\ell)}\Big), \tag{11}$$

with $\widetilde{A} = A + I$ and degree matrix $\widetilde{D}$.

For an affine quantizer (Jacob et al., 2018),

$$q(x) = \mathrm{clip}\Big(\mathrm{round}\Big(\frac{x}{s}\Big) + z, q_{\min}, q_{\max}\Big), \qquad \hat{x} = s(q(x) - z), \tag{12}$$

where s is the scale and z the zero point. Under min–max activation calibration,

$$s = \frac{x_{\max} - x_{\min}}{q_{\max} - q_{\min}}. \tag{13}$$

If $x_c = \mathrm{clip}(x, l, u)$ denotes range clipping, then quantization error can be decomposed conceptually as

$$x - \hat{x} = \underbrace{(x - x_c)}_{\text{clipping}} + \underbrace{(x_c - \hat{x})}_{\text{rounding}}. \tag{14}$$

A min–max observer can remain unchanged when $\min(H_\delta) \approx \min(H_0)$ and $\max(H_\delta) \approx \max(H_0)$ even though probability mass moves within that range. We call this specific failure mode *observer blindness*.

B Experimental and Statistical Details

Datasets and splits. Cora, CiteSeer, and PubMed are loaded with the PyTorch Geometric `Planetoid` wrapper and use the dataset-provided train/validation/test masks. The corresponding split sizes are 140/500/1000 nodes for Cora, 120/500/1000 for CiteSeer, and 60/500/1000 for PubMed. Roman-empire is loaded with PyTorch Geometric's `HeterophilousGraphDataset`; because it provides multiple predefined masks, we deterministically use split index 0, containing 11,331/5,665/5,666 train/validation/test nodes. Dataset-provided features and labels are retained without an additional feature-normalization transform. For all datasets, topology is canonicalized to a simple undirected graph by removing self-loops and duplicate edges while preserving node identities, features, labels, and masks.

Training and model selection. The default model configuration uses two message-passing layers, hidden dimension 64, ReLU, dropout 0.5, and Adam with learning rate 0.01 and weight decay 5×10^{-4}. Models are trained for at most 300 epochs with early-stopping patience 50. The checkpoint with the lowest validation cross-entropy loss is retained; the test split is never used for checkpoint selection. For Roman-empire+GCN refinement, a predefined four-configuration grid is trained for up to 1000 epochs with patience 150, again using validation loss for early stopping and excluding the test split from candidate selection. The globally selected refined configuration uses hidden dimension 128, learning rate 0.01, weight decay 5×10^{-4}, and dropout 0.5, and the original CP-DPR graph bank is reused unchanged.

Quantization and evaluation scope. Weights use symmetric per-output-channel fake quantization with integer ranges $[-127, 127]$ for 8 bits and $[-7, 7]$ for 4 bits. For each output channel, the scale is the channel maximum absolute weight divided by the positive range endpoint. Activations use per-tensor affine min–max quantization with unsigned ranges $[0, 255]$ for 8 bits and $[0, 15]$ for 4 bits. The affine scale is $(x_{\max} - x_{\min})/(q_{\max} - q_{\min})$, and the zero point is obtained by rounding $q_{\min} - x_{\min}/s$ and clipping it to the integer range. Fake quantization uses `torch.round` followed by clipping and dequantization. Activation calibration on G_0 uses all node activations without labels at the model input and at the hidden representation after the first message-passing layer and ReLU. Biases and final logits remain in FP32, and the adjacency is not quantized. Although calibration uses all nodes, QD$_Z and task metrics reported in the paper are computed on the test - node mask.$

Seeds, perturbations, and uncertainty. Training seeds are 1101, 2202, 3303, 4404, and 5505. Perturbation seeds are 101, 202, 303, 404, 505, 606, 707, 808, 909, and 1010. The main CP-DPR benchmark uses $\delta_E \in \{0.10, 0.20, 0.40\}$, producing ten perturbed graphs per severity. The separate task-level Cora bank uses the same degree-preserving double-edge-swap procedure without the exact component-partition constraint and additionally includes $\delta_E = 0.05$.

All crossed-bootstrap confidence intervals reported in the present manuscript use 5,000 replicates with base seed 20260918. Each replicate independently resamples the five training seeds and ten perturbation seeds with replacement, forms their Cartesian product, and recomputes the median; percentile 95% confidence intervals are reported. The original analysis also contained pair-level Wilcoxon signed-rank tests across the 50 model–graph pairs. Because those pairs share training and perturbation seeds and therefore are not independent observational units, the submission version does not use those p-values for its main inferential claims.

C Task-Level Cora Analysis

Table 3 reports EQD_M using Macro-F1 on the separate Cora degree-preserving rewiring bank without component preservation.

Under W8A8, the intervals include zero at every severity. Under W4A4, the medians are positive and increase with severity, but the interval excludes zero only at $\delta_E = 0.05$. Accordingly, the submission version does not describe the higher-severity task-level effect as statistically established by the crossed-bootstrap analysis.

Table 3: Cora task-level analysis. Median EQD_M (Macro-F1) with crossed-bootstrap 95% confidence intervals.

Condition	δ_E	Median EQD_M [95% CI]
W8A8	0.05	0.000003 [0.000000, 0.000417]
W8A8	0.10	0.000021 [0.000000, 0.000654]
W8A8	0.20	0.000026 [-0.000006, 0.000600]
W8A8	0.40	0.000251 [-0.000485, 0.000672]
W4A4	0.05	0.001163 [0.000119, 0.002982]
W4A4	0.10	0.002030 [-0.000651, 0.004092]
W4A4	0.20	0.003527 [-0.000165, 0.005842]
W4A4	0.40	0.004719 [-0.000994, 0.009554]

D Extended Limitations

First, the experiments use simulated PTQ rather than hardware-specific integer kernels, so they do not support claims about latency, energy, memory bandwidth, or device-specific arithmetic. Second, the benchmark covers node classification, two message-passing architectures, and W8A8/W4A4; attention-based or transformer-style graph models, link prediction, graph classification, temporal graphs, and other formats remain open. Third, adjacency itself is not quantized.

Fourth, CP-DPR preserves degrees and connected components but not every higher-order graph statistic. Homophily, clustering, assortativity, motifs, communities, spectral properties, and multi-hop organization may still change; the present experiments do not identify which of these best predicts numerical response. Fifth, rewiring is confined to the original largest connected component while smaller components remain unchanged. Sixth, the mechanism analysis is associational: the evidence is consistent with rounding mismatch and inconsistent with a clipping-dominated explanation in dominant positive regimes, but it does not identify a unique causal mediator.

Seventh, task-level EQD_M is confined to Cora and is evaluated on the unconstrained-component degree-preserving rewiring bank rather than the main CP-DPR bank. Eighth, absolute pre-perturbation $\mathrm{QD}_Z(A_0)$ is reported only for Cora/GCN, so the relative-to-baseline interpretation should not be generalized to external settings. Ninth, the architecture-dependence evidence comes from a single heterophilous dataset with two architectures; additional heterophilous benchmarks and operators are required to distinguish a GCN-specific effect from properties specific to Roman-empire. Finally, the crossed design contains five trained models and ten perturbation realizations per severity; broader replication would further strengthen uncertainty quantification.