Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Oversmoothing as Disagreement Contraction inMessage-Passing Graph Neural Networks

Fatih Uludağ

fatihuludag@yyu.edu.tr

Van Yüz

H. Eray Celik

Van Yüz

Research Article

Keywords: Graph Neural Networks, Oversmoothing, Message Passing, Random-Walk Diffusion, Stationary Projection, Representation Collapse

Posted Date: April 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9446792/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Oversmoothing as Disagreement Contraction in Message-Passing Graph Neural Networks

Fatih Uludağ[a], H. Eray Çelik[a]

[a] *Van Yuzuncu Yil University, Department of Economics, Van, Turkey*

Abstract

Oversmoothing is a central limitation of deep message-passing Graph Neural Networks (GNNs), wherein node representations become increasingly indistinguishable as depth grows and downstream discrimination deteriorates. This paper studies oversmoothing for a fixed-diffusion subclass of message-passing architectures, with random-walk diffusion $P = D^{-1}A$ on undirected graphs as the primary case; data-dependent propagators such as attention and dynamic rewiring are outside the scope of the present theory. We develop a diffusion-aligned, component-aware formulation that makes oversmoothing explicit as disagreement contraction under graph diffusion. To handle general irregular graphs, we introduce a stationary (degree-weighted) projector onto component-wise constants, $\Pi_D = C(C^\top DC)^{-1}C^\top D$, which preserves the invariance identities $P\Pi = \Pi$ and $\Pi P = \Pi$. This yields a principled decomposition of representations into consensus and disagreement components, with oversmoothing quantified by the decay of $\|(I - \Pi)H\|_{F,D}$ in the natural stationary geometry. Defining the disagreement-restricted diffusion operator $T = (I - \Pi)P(I - \Pi)$ and the mixing coefficient $\mu = \|T\|_{2,D}$, we obtain non-expansive or contractive disagreement bounds for linear propagation and layerwise contraction proxies $\rho_\ell = L_\sigma \mu \|W_\ell\|_2$ for nonlinear layers with row-wise Lipschitz activations. For residual mixing $H_{\ell+1} = (1 - \alpha_\ell)H_\ell + \alpha_\ell \sigma(PH_\ell W_\ell)$, the bound $(1 - \alpha_\ell) + \alpha_\ell \rho_\ell$ makes the role of residual strength explicit in the contractive regime. The framework yields operator-consistent diagnostics based on relative disagreement and collapse depth, linking diffusion mixing to oversmoothing behavior across datasets, depths, and propagation choices.

Email addresses: fatihuludag@yyu.edu.tr (Fatih Uludağ), [eray.celik@yyu.edu.tr] (H. Eray Çelik)

Keywords: Graph Neural Networks, Oversmoothing, Message Passing, Random-Walk Diffusion, Stationary Projection, Representation Collapse

1. Introduction

Message-passing Graph Neural Networks (MPNNs) learn node representations by repeatedly aggregating neighborhood information and applying learnable transformations and nonlinearities (Kipf and Welling, 2017; Hamilton et al., 2017; Gilmer et al., 2017). The combination of depth with receptive field expansion can improve denoising and capture multi-hop structure, but deep stacks often show deterioration in node-level discrimination. This depth barrier is commonly attributed to *oversmoothing*, where node embeddings become increasingly indistinguishable as depth grows (Li et al., 2018; Oono and Suzuki, 2019).

In this paper, the theoretical results are derived for *fixed graph-diffusion operators* that preserve a stationary distribution, with random-walk diffusion $P = D^{-1}A$ as the primary case. This setting covers a broad fixed-diffusion subclass of message-passing architectures, but it does not include data-dependent propagators such as attention or dynamic rewiring. Within this scope, we develop a diffusion-aligned, component-aware formulation that makes oversmoothing explicit as a contraction mechanism.

The core idea is to decompose any layer representation H into: (i) a consensus component that is constant on each connected component in the sense preserved by the diffusion, and (ii) a within-component disagreement component capturing deviations from that consensus. Oversmoothing is then defined as the progressive decay of disagreement toward zero as depth increases.

The random-walk diffusion $P = D^{-1}A$ preserves degree-weighted (stationary) averages on irregular graphs rather than uniform averages. For example, on the three-node path 1–2–3, the preserved component average is proportional to weights $(1, 2, 1)$, not the uniform mean. Using uniform component means therefore breaks key invariance identities unless the graph is regular. To ensure mathematical correctness on general undirected graphs, we define consensus via a stationary (degree-weighted) projector compatible with P, and we derive contraction bounds in the corresponding natural D-weighted geometry. The novelty is therefore not a new generic Markov-chain contraction theorem per se, but an operator-consistent formulation of oversmoothing measurement on irregular and disconnected graphs.

Contributions.

1. **Operator-aligned oversmoothing definition.** We formalize oversmoothing as decay of the disagreement component $(I - \Pi)H$ under message passing, where Π is the stationary (degree-weighted) projector compatible with random-walk diffusion.
2. **Stationary-consistent theory on irregular graphs.** By adopting $\Pi_D = C(C^\top DC)^{-1}C^\top D$ and the D-weighted geometry, we ensure invariance identities $P\Pi = \Pi$ and $\Pi P = \Pi$ hold on general undirected graphs (not only regular graphs).
3. **Mixing-controlled disagreement bounds.** Defining $T = (I - \Pi)P(I - \Pi)$ yields a disagreement-restricted operator whose induced norm $\mu = \|T\|_{2,D}$ governs worst-case non-expansive or contractive behavior for linear propagation and yields layerwise bounds for nonlinear layers via a disagreement–Lipschitz lemma.
4. **Residual mixing as a stability dial.** For residual-mixing propagation $H_{\ell+1} = (1 - \alpha_\ell)H_\ell + \alpha_\ell \sigma(PH_\ell W_\ell)$, we obtain an explicit bound $(1 - \alpha_\ell) + \alpha_\ell \rho_\ell$, clarifying how residual strength modulates disagreement contraction in the contractive regime.
5. **Operator-consistent diagnostics.** The theory motivates depth-resolved diagnostics including relative disagreement and collapse depth, enabling failure-aware empirical analysis across datasets, backbones, and diffusion choices.

2. Related Work

2.1. Oversmoothing: theory, dynamics, and operator viewpoints

Oversmoothing is widely understood as the depth-driven homogenization of node representations in message-passing GNNs, which reduces class-discriminative variation. Early analyses interpret the phenomenon through repeated low-pass filtering or Laplacian smoothing (Li et al., 2018; Oono and Suzuki, 2019). A complementary operator-theoretic view connects repeated message passing to Markov-chain-like diffusion toward stationary structure (Zhao et al., 2022).

Energy- and spectrum-based diagnostics quantify variation decay and connect oversmoothing to contraction of graph Dirichlet energy across layers.

For example, Dirichlet Energy Constrained Learning introduces energy constraints that enable deeper stacks by controlling information loss (Zhou et al., 2021). Extensions consider irregular graphs and generalized diffusions, including fractional Laplacians, which alter the speed and pattern of smoothing (Maskey et al., 2023).

Oversmoothing is not limited to linear GCN-style diffusions. Analyses of attention-based models treat attention as a nonlinear, data-dependent operator and study when homogenization persists and to what extent attention can resist depth-related collapse (Wu et al., 2023). These findings motivate diagnostics that preserve the natural symmetries of propagation operators and separate preserved consensus structure from shrinking disagreement components.

2.2. Making GNNs deep: stabilization, diffusion control, and propagation design

Deep GNN design uses residual or dense connectivity and controlled propagation to stabilize training and improve expressivity (Li et al., 2019). Diffusion-based architectures also provide routes to mitigate oversmoothing while retaining multi-hop information flow, for example through source-term diffusion (Thorpe et al., 2022) or personalized propagation schemes (Gasteiger et al., 2019). These lines of work motivate measuring whether an intervention genuinely changes contraction behavior, rather than merely shifting optimization outcomes.

2.3. Oversquashing and related depth pathologies

Depth limitations also arise from over-squashing, where long-range signals are compressed through structural bottlenecks (Alon and Yahav, 2021; Topping et al., 2021). Recent surveys emphasize that deep GNN behavior is best understood through multiple interacting failure modes rather than through oversmoothing alone (Akansha, 2025). This makes it important to distinguish between preserved consensus structure, disappearing disagreement, and bottleneck-induced information loss.

2.4. Positioning of this paper

We target a gap that remains under-specified in much oversmoothing literature: operator-consistent, stationarity-aligned measurement on irregular and disconnected graphs. In contrast to uniform-mean or component-agnostic baselines, we define consensus using a stationary (degree-weighted) projector

compatible with random-walk diffusion, enabling a principled decomposition into consensus and disagreement. Oversmoothing is quantified by disagreement contraction in the natural stationary geometry, linking diffusion mixing to depth behavior and yielding experiment-ready diagnostics.

3. Preliminaries & Notation

We consider an undirected graph $G = (V, E)$ with $|V| = n$. Let $A \in \mathbb{R}^{n\times n}$ be the adjacency matrix and $D = \mathrm{diag}(d_1, \ldots, d_n)$ the degree matrix with $d_i = \sum_j A_{ij}$ and $d_i > 0$.

3.1. Random-walk diffusion and stationarity

The random-walk diffusion is

$$P = D^{-1}A, \tag{1}$$

which is row-stochastic and can be interpreted as a Markov transition matrix. We also consider lazy diffusion

$$P_{\mathrm{lazy}} = (1-\lambda)I + \lambda P, \qquad \lambda \in (0, 1], \tag{2}$$

which improves aperiodicity while preserving stationary behavior. For an undirected graph, the stationary distribution satisfies

$$\pi_i = \frac{d_i}{\sum_{j=1}^{n} d_j}, \qquad \pi^\top P = \pi^\top. \tag{3}$$

Although the main results are stated for $P = D^{-1}A$, the same projector-based viewpoint extends to any fixed row-stochastic operator with stationary distribution π after replacing D by a diagonal matrix proportional to π.

3.2. Message passing layers

Let $H_\ell \in \mathbb{R}^{n\times d}$ denote node representations at layer ℓ (rows are node embeddings). The base update is

$$H_{\ell+1} = \sigma(PH_\ell W_\ell), \tag{4}$$

and the *residual-mixing* (skip) variant is

$$H_{\ell+1} = (1-\alpha_\ell)H_\ell + \alpha_\ell\, \sigma(PH_\ell W_\ell), \qquad \alpha_\ell \in [0, 1]. \tag{5}$$

We also define the corresponding *pre-activation* linear message:

$$Z_\ell := PH_\ell W_\ell. \tag{6}$$

3.3. D-weighted geometry

For vectors $x, y \in \mathbb{R}^n$, define

$$\langle x, y \rangle_D = x^\top D y, \qquad \|x\|_{2,D} = \sqrt{x^\top D x}. \tag{7}$$

For matrices $H \in \mathbb{R}^{n \times d}$, we use the D-weighted Frobenius norm

$$\|H\|_{F,D} := \sqrt{\mathrm{tr}(H^\top D H)} = \left(\sum_{i=1}^{n} d_i \|H_{i:}\|_2^2 \right)^{1/2}, \tag{8}$$

and, when convenient, its squared form

$$\|H\|_{F,D}^2 = \mathrm{tr}(H^\top D H). \tag{9}$$

The induced operator norm is

$$\|M\|_{2,D} = \sup_{x \neq 0} \frac{\|Mx\|_{2,D}}{\|x\|_{2,D}}. \tag{10}$$

3.4. Stationary (degree-weighted) consensus projector

Let m be the number of connected components and $C \in \mathbb{R}^{n \times m}$ the component-indicator matrix. The stationary projector is

$$\Pi \equiv \Pi_D := C(C^\top D C)^{-1} C^\top D, \tag{11}$$

the D-orthogonal projector onto the component-wise constant (consensus) subspace. For the path graph 1–2–3, this projector averages with weights $(1, 2, 1)$ on the connected component; the uniform mean is not diffusion-invariant unless all degrees are equal. For matrix features, ΠH is applied columnwise. We decompose

$$H = \Pi H + (I - \Pi) H, \qquad H^\perp := (I - \Pi) H. \tag{12}$$

3.5. Layerwise disagreement metrics and collapse depth

To directly track oversmoothing as disagreement contraction, we define the layerwise *relative disagreement* on pre-activations:

$$S_{\mathrm{rel,pre}}(\ell) := \frac{\|(I - \Pi) Z_\ell\|_{F,D}}{\|(I - \Pi) Z_1\|_{F,D}}, \tag{13}$$

and analogously (optional) post-activation relative disagreement

$$S_{\mathrm{rel,post}}(\ell) := \frac{\|(I-\Pi)H_\ell\|_{F,D}}{\|(I-\Pi)H_1\|_{F,D}}. \tag{14}$$

For a threshold $\tau \in (0,1)$, the *collapse depth* is

$$L_c(\tau) := \min\{\ell \geq 1 : S_{\mathrm{rel,pre}}(\ell) \leq \tau\}, \tag{15}$$

and we record $L_c(\tau) = \text{None}$ if the threshold is not reached within the explored depth.

4. Theory: Oversmoothing as Disagreement Contraction

4.1. Invariance of stationary component-wise constants

Proposition 1 (Stationary invariance)**.** *For an undirected graph and $P = D^{-1}A$, the stationary projector Π satisfies*

$$P\Pi = \Pi, \qquad \Pi P = \Pi. \tag{16}$$

Proposition 1 implies that the consensus subspace is invariant and that disagreement evolves autonomously under the projected operator. In particular,

$$(I-\Pi)P = (I-\Pi)P(I-\Pi) = P(I-\Pi). \tag{17}$$

Thus the dynamics can be restricted to the disagreement subspace without loss.

4.2. Disagreement-restricted diffusion and the mixing coefficient

Define the disagreement-restricted diffusion operator

$$T := (I-\Pi)P(I-\Pi), \tag{18}$$

and the mixing coefficient

$$\mu := \|T\|_{2,D}. \tag{19}$$

4.3. Linear disagreement contraction

Proposition 2 (Linear contraction on disagreement)**.** *If $H_{\ell+1} = PH_\ell$, then*

$$\|(I-\Pi)H_{\ell+1}\|_{F,D} \leq \mu\|(I-\Pi)H_\ell\|_{F,D}, \tag{20}$$

and hence $\|(I-\Pi)H_\ell\|_{F,D} \leq \mu^\ell\|(I-\Pi)H_0\|_{F,D}$.

Interpretation. Since P is non-expansive in the D-geometry and $I - \Pi$ is a D-orthogonal projector, one has $\mu \leq 1$. Strict exponential contraction requires $\mu < 1$. For plain random walks on periodic components (e.g., bipartite graphs), μ can equal 1, in which case Proposition 2 yields a non-expansive bound rather than a strict contraction. This motivates lazy diffusion or self-looped propagation in practice.

4.4. Nonlinearity in D-geometry: a Lipschitz route

Assumption 1 (Row-wise Euclidean Lipschitz)**.** *The activation σ acts row-wise via $\varphi : \mathbb{R}^d \to \mathbb{R}^d$ and is L_σ-Lipschitz: $\|\varphi(u) - \varphi(v)\|_2 \leq L_\sigma \|u - v\|_2$.*

Lemma 1 (Lipschitz in $\|\cdot\|_{F,D}$)**.** *For any $U, V \in \mathbb{R}^{n\times d}$,*

$$\|\sigma(U) - \sigma(V)\|_{F,D} \leq L_\sigma \|U - V\|_{F,D}. \tag{21}$$

Lemma 2 (Disagreement Lipschitzness)**.** *For any $Z \in \mathbb{R}^{n\times d}$,*

$$\|(I - \Pi)\sigma(Z)\|_{F,D} \leq L_\sigma \|(I - \Pi)Z\|_{F,D}. \tag{22}$$

4.5. Base-layer disagreement contraction for fixed diffusion

Define the layer-wise contraction proxy

$$\rho_\ell := L_\sigma \, \mu \, \|W_\ell\|_2. \tag{23}$$

Proposition 3 (Base-layer disagreement bound)**.** *For the fixed-diffusion update $H_{\ell+1} = \sigma(PH_\ell W_\ell)$,*

$$\|(I - \Pi)H_{\ell+1}\|_{F,D} \leq \rho_\ell \|(I - \Pi)H_\ell\|_{F,D}. \tag{24}$$

4.6. Residual layer: a contraction dial

Consider residual mixing

$$H_{\ell+1} = (1 - \alpha_\ell)H_\ell + \alpha_\ell \sigma(PH_\ell W_\ell). \tag{25}$$

Proposition 4 (Residual disagreement bound)**.** *Under residual mixing,*

$$\|(I - \Pi)H_{\ell+1}\|_{F,D} \leq \Big((1 - \alpha_\ell) + \alpha_\ell \rho_\ell\Big) \|(I - \Pi)H_\ell\|_{F,D}. \tag{26}$$

Interpretation. The factor $(1 - \alpha_\ell) + \alpha_\ell \rho_\ell$ makes the role of residual mixing explicit. In the contractive regime $\rho_\ell < 1$, smaller α_ℓ moves the bound closer to 1 and therefore slows disagreement contraction, delaying oversmoothing onset. When $\rho_\ell \geq 1$, the same expression is only an upper bound and does not by itself imply delay.

5. Experimental Design

We investigate oversmoothing through the lens of disagreement contraction in deep graph neural networks. All oversmoothing measurements are computed on pre-activation linear messages evaluated in evaluation mode. Let $\widetilde{P}$ denote the fixed propagation operator used in a given run (e.g., rw, lazy, or ppr). We define the pre-activation representation by $Z_\ell := \widetilde{P} H_\ell W_\ell$ and quantify disagreement as $(I - \Pi)Z_\ell$.

Relative disagreement and collapse depth. We measure contraction using $S_{\mathrm{rel,pre}}(\ell) = \|(I - \Pi)Z_\ell\|_{F,D} / \|(I - \Pi)Z_1\|_{F,D}$ and summarize onset of oversmoothing by $L_c(\tau)$, the smallest layer such that $S_{\mathrm{rel,pre}}(\ell) \le \tau$. We report $\tau = 0.1$ in the main experiments, and perform a post-hoc sensitivity analysis for $\tau \in \{0.05, 0.1, 0.2\}$ (Table A.12).

Failure-aware learning-regime diagnostics. Oversmoothing metrics must be interpreted within a learning regime. We therefore report (i) validation accuracy at the best epoch, (ii) training accuracy at that checkpoint, and (iii) explicit stability labels. We count a run as *near-collapse* when `train_acc_at_best` $<$ 0.35. We additionally define a *hard-collapse* detector that flags trivial predictors when predicted labels concentrate and predictions are confident and near-chance.

Hard-collapse detector (operational). At the best-validation checkpoint, let $p_{i,c}$ be the predicted probability of class c at node i, and $\hat{y}_i = \arg\max_c p_{i,c}$. Let $q_c = \frac{1}{n}\sum_i \mathbf{1}[\hat{y}_i = c]$ be the predicted-class histogram and $H_{\mathrm{hist}} = -\frac{1}{\log C}\sum_c q_c \log q_c$ its normalized entropy. Let $H_{\mathrm{pred}} = \frac{1}{n}\sum_i \Big(-\frac{1}{\log C}\sum_c p_{i,c}\log p_{i,c}\Big)$ be mean normalized predictive entropy. We flag *hard-collapse* if ($\max_c q_c \ge 0.95$ or $H_{\mathrm{hist}} \le 0.05$) and $H_{\mathrm{pred}} \le 0.20$ and $\mathrm{Acc}_{\mathrm{train}} \le 1/C + 0.05$.

Intervention knobs. We study three intervention knobs: (i) depth L, (ii) residual mixing α, and (iii) diffusion family (rw/lazy/ppr) and mixing strength (lazy λ). The theory predicts smaller α slows contraction and delays oversmoothing, while overly weak mixing (small α or λ) can induce under-propagation.

Datasets and seeds. Main sweeps use Cora (real-world), a synthetic SBM (mechanistic), and ogbn-arxiv (large-scale). Unless noted otherwise, we report mean±std over five standard seeds $\{1, 7, 13, 42, 123\}$.

Additional breadth matrix. As a robustness check, we also evaluate a small cross-dataset/backbone matrix across Cora, Chameleon, and ogbn-arxiv using FD-MPNN, GCN, and GraphSAGE at depths $L \in \{2, 4, 8\}$. These results are reported in Table A.9 and should be interpreted as a breadth check rather than a direct theorem-validation experiment.

Baseline oversmoothing metrics. Alongside our operator-aligned diagnostics, we compute commonly used baseline metrics from layerwise representations: Dirichlet energy, mean cosine similarity on edges, and feature variance. We assess *predictive validity* by correlating diagnostics with test accuracy and by classifying collapse outcomes (Tables A.10–A.11).

6. Results

6.1. Cora: depth boundary and intervention effects

Table A.1 shows a sharp stability boundary under strong mixing ($\alpha = 1.0$) for lazy diffusion. Shallow models ($L = 2, 4$) remain in a learning regime with high accuracy and no failures, and $L_c(0.1)$ is not reached within the explored depth. In contrast, at $L = 8$ and $L = 16$ the model enters a near-degenerate regime (near-collapse rate 100%) with early collapse depth (median $L_c(0.1) = 2$).

Residual mixing provides an effective stability dial at depth $L = 8$ (Table A.2). Decreasing α from 1.0 to 0.5 restores stable learning (no near-collapse) and yields strong performance, and further decreasing to $\alpha = 0.2$ improves accuracy while remaining stable. These results align with Proposition 4: smaller α slows the effective contraction factor $(1 - \alpha) + \alpha\rho_\ell$.

Diffusion-family effects are modest once the system operates in a stable regime (Table A.3): rw/lazy/ppr achieve comparable accuracy at ($L = 8, \alpha = 0.5$) with zero failures and no observed $L_c(0.1)$ within depth. In contrast, mixing-strength λ in lazy diffusion exhibits a clear stability–propagation trade-off (Table A.4): very weak mixing ($\lambda = 0.1$) degrades accuracy (under-propagation), while moderate mixing (best here at $\lambda = 0.5$) maximizes accuracy.

6.2. SBM: mechanistic contraction under near-degeneracy

SBM provides a mechanism-transparent environment. Under lazy diffusion with $\alpha = 1.0$, shallow depth ($L = 4$) achieves near-perfect performance with no collapse, while deeper models ($L \geq 8$) rapidly enter near-collapse with

early collapse depth (Table A.5). Notably, the hard-collapse detector is not triggered in these SBM runs; degradation appears mainly as near-collapse, motivating explicit failure-aware reporting.

At $L = 32$, residual mixing modulates contraction dynamics (Table A.6): decreasing α substantially delays collapse depth (median $L_c(0.1)$ increases from 2 to 12), confirming the directional effect on contraction speed. However, predictive accuracy remains at a near-chance floor in this extremely deep regime, illustrating that contraction diagnostics must be interpreted jointly with learning-regime indicators. Diffusion-family changes at $(L = 32, \alpha = 0.5)$ do not rescue accuracy and yield similar collapse depths (Table A.7).

6.3. ogbn-arxiv: large-scale depth sensitivity

On ogbn-arxiv under lazy diffusion with $\alpha = 0.5$, accuracy decreases with depth (Table A.8), and a subset of runs at $L = 8$ enters near-collapse. In these depths, $L_c(0.1)$ is typically not reached within the explored range, so we rely on layerwise disagreement summaries and predictive-validity analyses below to interpret contraction.

6.4. Cross-dataset/backbone breadth check

Table A.9 reports a small breadth matrix (seeds $\{1, 7\}$) across homophily (Cora), heterophily (Chameleon), and large-scale (ogbn-arxiv). Across backbones, depth increases reduce performance most sharply on heterophily and large-scale graphs, and near-collapse incidence can be backbone-dependent (e.g., GraphSAGE exhibits higher variance on Chameleon at $L = 8$). This breadth check supports that depth-induced degradation is not limited to a single dataset.

6.5. Predictive validity and threshold sensitivity

Across 207 unique runs spanning datasets, depths, backbones, and interventions, our disagreement-contraction diagnostics show strong predictive validity. Table A.10 reports Spearman correlations between diagnostics and test accuracy: the minimum relative disagreement ($S_{\text{rel,pre}}$ and $S_{\text{rel,post}}$ minima) exhibits strong monotone association with accuracy, comparable to strong baselines such as feature variance. For collapse classification (hard or near), Table A.11 shows that disagreement summaries nearly perfectly separate degenerate runs (ROC-AUC ≥ 0.998), while threshold-based collapse depth $L_c(\tau)$ is sensitive to τ and performs best at $\tau = 0.2$ in this study. Table A.12 summarizes the sensitivity of $L_c(\tau)$ to τ, motivating reporting

S_{rel} summaries alongside $L_c(\tau)$, especially when $L_c(\tau)$ is not reached within measured depth.

7. Discussion

This work advances a metric-centric account of oversmoothing by treating it as disagreement contraction toward a degree-weighted stationary, component-wise constant subspace. The main contribution is not a new generic Markov-chain convergence theorem, but an operator-consistent formulation of oversmoothing measurement on irregular and disconnected graphs. The stationary projector Π_D yields a consensus/disagreement decomposition aligned with random-walk diffusion invariances and thereby supports contraction bounds in the natural stationary geometry.

Empirically, depth and mixing jointly define a stability envelope. On Cora, strong mixing ($\alpha = 1.0$) induces an early transition to near-degeneracy at moderate depth, while reducing residual mixing ($\alpha \in \{0.5, 0.2\}$) restores stable learning and delays contraction. On SBM, extreme depth regimes exhibit near-degenerate predictors where contraction metrics change under interventions (e.g., smaller α increases L_c) without recovering predictive performance, emphasizing the necessity of failure-aware interpretation.

Across a broad set of runs, disagreement summaries show strong predictive validity. At the same time, they should be viewed as complementary diagnostics rather than universally dominant baselines: feature variance is also very strong in our study, whereas disagreement measures add stationary-aligned interpretability and a direct connection to the analyzed propagation operator. Taken together, the results support disagreement contraction as a practical and theoretically grounded diagnostic lens for oversmoothing in fixed-diffusion message passing.

8. Limitations and Future Work

This work has several limitations. First, the current theory covers fixed, π-stationary diffusion operators, with random-walk diffusion $P = D^{-1}A$ as the main case. Extending the disagreement framework to data-dependent propagators such as attention, adaptive graph rewiring, or sample-conditioned normalization remains an important open direction.

Second, our empirical failure thresholds are operational rather than universal. Near-collapse uses `train_acc_at_best` < 0.35, and collapse depth

$L_c(\tau)$ depends on the chosen threshold τ. Although these diagnostics are informative in our sweeps, threshold calibration may vary across datasets, class imbalance, and task difficulty.

Third, the cross-dataset/backbone breadth matrix in this revision is intentionally small and should be read as a robustness check rather than a one-to-one validation of the fixed-diffusion theory. It uses two seeds and includes backbones whose effective propagation operators are related to, but not identical with, the operator analyzed in Sections 3–4.

Finally, diffusion-family effects can be hyperparameter-sensitive, especially for PPR and residual mixing. A fuller study should report standardized protocols, larger hyperparameter grids, and direct links between empirical contraction and operator quantities such as spectral gap, periodicity, and mixing time.

Declarations

Funding The author received no specific funding for this work.

Competing interests The author declares no competing interests.

Data availability Public datasets analyzed in this study are available from their original public sources, including Cora, Chameleon, and OGBN-Arxiv. The synthetic SBM data used in this study are generated from the parameters reported in Appendix A.2. The experiment configurations, scripts, and aggregated outputs supporting the findings of this study are available from the corresponding author upon reasonable request.

References

[1] S. Akansha, "Over-squashing in Graph Neural Networks: A comprehensive survey," *Neurocomputing*, 2025.

[2] U. Alon and E. Yahav, "On the Bottleneck of Graph Neural Networks and its Practical Implications," ICLR 2021.

[3] J. Gasteiger et al., "Predict then Propagate: Graph Neural Networks meet Personalized PageRank," ICLR 2019.

[4] J. Gilmer et al., "Neural Message Passing for Quantum Chemistry," ICML 2017.

[5] W. L. Hamilton et al., "Inductive Representation Learning on Large Graphs," NeurIPS 2017.

[6] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," ICLR 2017.

[7] Q. Li et al., "Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning," AAAI 2018.

[8] G. Li et al., "DeepGCNs: Can GCNs Go As Deep As CNNs?," ICCV 2019.

[9] S. Maskey et al., "A Fractional Graph Laplacian Approach to Oversmoothing," NeurIPS 2023.

[10] K. Oono and T. Suzuki, "Graph Neural Networks Exponentially Lose Expressive Power for Node Classification," arXiv:1905.10947, 2019.

[11] M. Thorpe et al., "GRAND++: Graph Neural Diffusion with a Source Term," ICLR 2022.

[12] J. Topping et al., "Understanding over-squashing and bottlenecks on graphs via curvature," arXiv:2111.14522, 2021.

[13] X. Wu et al., "Demystifying Oversmoothing in Attention-Based Graph Neural Networks," NeurIPS 2023.

[14] W. Zhao et al., "Comprehensive Analysis of Oversmoothing in Graph Neural Networks from Markov Chains Perspective," arXiv:2211.06605, 2022.

[15] K. Zhou et al., "Dirichlet Energy Constrained Learning for Deep Graph Neural Networks," NeurIPS 2021.

Appendix A. Supplementary Material

Appendix A.1. Proofs for Section 4

Useful norm facts. For any matrix $X \in \mathbb{R}^{n\times d}$ and any feature-space operator $W \in \mathbb{R}^{d\times d'}$,

$$\|XW\|_{F,D} = \|D^{1/2}XW\|_F \le \|W\|_2 \, \|D^{1/2}X\|_F = \|W\|_2 \, \|X\|_{F,D}. \tag{A.1}$$

Likewise, for any node-space operator $M \in \mathbb{R}^{n\times n}$ and any matrix $X = [x_1, \ldots, x_d]$,

$$\|MX\|_{F,D}^2 = \sum_{j=1}^{d} \|Mx_j\|_{2,D}^2 \le \|M\|_{2,D}^2 \sum_{j=1}^{d} \|x_j\|_{2,D}^2 = \|M\|_{2,D}^2 \|X\|_{F,D}^2. \tag{A.2}$$

Finally,

$$S := D^{1/2}PD^{-1/2} = D^{-1/2}AD^{-1/2} \tag{A.3}$$

is symmetric. Hence $\|P\|_{2,D} = \|S\|_2$. Since $I - S$ is the normalized Laplacian, the spectrum of S lies in $[-1, 1]$, so $\|P\|_{2,D} \le 1$. Therefore P is non-expansive in the D-geometry.

Proof of Proposition 1. Let c_k denote the k-th column of C, i.e., the indicator vector of connected component k. Because each connected component is closed under the random walk, every node in component k sends all of its probability mass to nodes in the same component. Therefore

$$Pc_k = c_k \qquad \text{for each } k, \tag{A.4}$$

and hence $PC = C$. It follows that

$$P\Pi = PC(C^\top DC)^{-1}C^\top D = C(C^\top DC)^{-1}C^\top D = \Pi. \tag{A.5}$$

For the second identity, use $DP = A$:

$$\Pi P = C(C^\top DC)^{-1}C^\top DP = C(C^\top DC)^{-1}C^\top A. \tag{A.6}$$

Now $C^\top A = C^\top D$. Indeed, if node j belongs to component k, then

$$(C^\top A)_{kj} = \sum_{i\in V_k} A_{ij} = d_j = (C^\top D)_{kj}, \tag{A.7}$$

while for any $r \neq k$ one has

$$(C^\top A)_{rj} = 0 = (C^\top D)_{rj}. \tag{A.8}$$

Therefore

$$\Pi P = C(C^\top DC)^{-1}C^\top D = \Pi. \tag{A.9}$$

As an immediate corollary,

$$(I - \Pi)P = P - \Pi = P(I - \Pi), \tag{A.10}$$

and also

$$(I - \Pi)P(I - \Pi) = (I - \Pi)P, \tag{A.11}$$

because $P\Pi = \Pi$.

Proof of Proposition 2. Define the disagreement component

$$E_\ell := (I - \Pi)H_\ell. \tag{A.12}$$

Using Proposition 1,

$$E_{\ell+1} = (I - \Pi)H_{\ell+1} = (I - \Pi)PH_\ell = (I - \Pi)P(I - \Pi)H_\ell = TE_\ell. \tag{A.13}$$

Applying the induced operator norm columnwise gives

$$\|E_{\ell+1}\|_{F,D} = \|TE_\ell\|_{F,D} \leq \|T\|_{2,D}\, \|E_\ell\|_{F,D} = \mu\, \|E_\ell\|_{F,D}. \tag{A.14}$$

Iterating this bound yields

$$\|E_\ell\|_{F,D} \leq \mu^\ell \|E_0\|_{F,D}. \tag{A.15}$$

Proof of Lemma 1. Let u_i and v_i denote the i-th rows of U and V. Since σ acts row-wise via an L_σ-Lipschitz map φ, we have

$$\|\sigma(U) - \sigma(V)\|_{F,D}^2 = \sum_{i=1}^n d_i\, \|\varphi(u_i) - \varphi(v_i)\|_2^2 \leq L_\sigma^2 \sum_{i=1}^n d_i\, \|u_i - v_i\|_2^2 = L_\sigma^2 \|U - V\|_{F,D}^2. \tag{A.16}$$

Taking square roots proves the claim.

Proof of Lemma 2. Because ΠZ is component-wise constant, all rows of ΠZ are identical within each connected component. Since σ acts row-wise using the same map on every node, $\sigma(\Pi Z)$ is again component-wise constant. Therefore

$$(I - \Pi)\sigma(\Pi Z) = 0. \tag{A.17}$$

Using this identity,

$$(I - \Pi)\sigma(Z) = (I - \Pi)\Big(\sigma(Z) - \sigma(\Pi Z)\Big). \tag{A.18}$$

Because $I - \Pi$ is a D-orthogonal projector, it is non-expansive in $\|\cdot\|_{F,D}$, so

$$\|(I-\Pi)\sigma(Z)\|_{F,D} \le \|\sigma(Z)-\sigma(\Pi Z)\|_{F,D} \le L_\sigma \|Z-\Pi Z\|_{F,D} = L_\sigma \|(I-\Pi)Z\|_{F,D}, \tag{A.19}$$

where the second inequality uses Lemma 1.

Proof of Proposition 3. Recall

$$Z_\ell = P H_\ell W_\ell. \tag{A.20}$$

By Lemma 2,

$$\|(I - \Pi)H_{\ell+1}\|_{F,D} = \|(I - \Pi)\sigma(Z_\ell)\|_{F,D} \le L_\sigma \|(I - \Pi)Z_\ell\|_{F,D}. \tag{A.21}$$

Next, using Proposition 1 and the norm facts above,

$$\begin{aligned}
\|(I - \Pi)Z_\ell\|_{F,D} &= \|(I - \Pi)P H_\ell W_\ell\|_{F,D} && \text{(A.22)}\\
&= \|T H_\ell W_\ell\|_{F,D} && \text{(A.23)}\\
&= \|T(I - \Pi)H_\ell W_\ell\|_{F,D} && \text{(A.24)}\\
&\le \|T\|_{2,D}\, \|(I - \Pi)H_\ell W_\ell\|_{F,D} && \text{(A.25)}\\
&\le \mu\, \|W_\ell\|_2\, \|(I - \Pi)H_\ell\|_{F,D}. && \text{(A.26)}
\end{aligned}$$

Combining the two displays gives

$$\|(I-\Pi)H_{\ell+1}\|_{F,D} \le L_\sigma \mu \|W_\ell\|_2\, \|(I-\Pi)H_\ell\|_{F,D} = \rho_\ell\, \|(I-\Pi)H_\ell\|_{F,D}. \tag{A.27}$$

Proof of Proposition 4. Applying $I - \Pi$ to the residual update yields

$$(I - \Pi)H_{\ell+1} = (1 - \alpha_\ell)(I - \Pi)H_\ell + \alpha_\ell (I - \Pi)\sigma(P H_\ell W_\ell). \tag{A.28}$$

Taking $\|\cdot\|_{F,D}$ and using the triangle inequality together with Proposition 3,

$$\|(I-\Pi)H_{\ell+1}\|_{F,D} \le (1-\alpha_\ell)\|(I-\Pi)H_\ell\|_{F,D} + \alpha_\ell\|(I-\Pi)\sigma(PH_\ell W_\ell)\|_{F,D} \tag{A.29}$$

$$\le (1-\alpha_\ell)\|(I-\Pi)H_\ell\|_{F,D} + \alpha_\ell\rho_\ell\|(I-\Pi)H_\ell\|_{F,D} \tag{A.30}$$

$$= \Big((1-\alpha_\ell)+\alpha_\ell\rho_\ell\Big)\|(I-\Pi)H_\ell\|_{F,D}. \tag{A.31}$$

This proves the stated residual disagreement bound.

Appendix A.2. Reproducibility details

Task, checkpoint selection, and reporting. All experiments are node-classification runs. Model selection is based on validation accuracy, and the reported test accuracy for a run is taken from the best-validation checkpoint. Unless otherwise noted, results are reported as mean±std over the seed set $\{1, 7, 13, 42, 123\}$. The breadth matrix in Table A.9 uses seeds $\{1, 7\}$ only and should be read as a robustness check rather than as a directly comparable five-seed sweep.

Propagation operators and model families. The paper-aligned fixed-diffusion model (FD-MPNN) uses a fixed propagation operator $\widetilde{P}$ inside each layer:

$$H_{\ell+1} = (1-\alpha_\ell)H_\ell + \alpha_\ell\,\sigma(\widetilde{P}H_\ell W_\ell), \tag{A.32}$$

with the base case recovered by $\alpha_\ell = 1$. For random-walk and lazy diffusion we use

$$P_{\text{rw}} = D^{-1}A, \qquad P_{\text{lazy}} = (1-\lambda)I + \lambda D^{-1}A. \tag{A.33}$$

The exposed fixed-diffusion base configs enable self-loops. PPR-based sweeps use the `ppr` diffusion type; in the exposed Cora and ogbn-arxiv base configs, the associated parameters are $\alpha_{\text{ppr}} = 0.15$ and $K = 10$.

Measured representations and collapse diagnostics. All oversmoothing diagnostics are computed on pre-activation linear messages in evaluation mode:

$$Z_\ell := \widetilde{P}H_\ell W_\ell, \qquad S_{\text{rel,pre}}(\ell) := \frac{\|(I-\Pi)Z_\ell\|_{F,D}}{\|(I-\Pi)Z_1\|_{F,D}}. \tag{A.34}$$

The main experiments use $\tau = 0.1$ and `norm_layer=1`. Collapse depth is defined as

$$L_c(\tau) := \min\{\ell \ge 1 : S_{\text{rel,pre}}(\ell) \le \tau\}. \tag{A.35}$$

A run is labeled near-collapse when

$$\texttt{train_acc_at_best} < 0.35. \tag{A.36}$$

In the exposed Cora and ogbn-arxiv base configs, the hard-collapse detector uses

$$q_{\max} = 0.95, \qquad H_{\text{hist}} \leq 0.05, \qquad H_{\text{pred}} \leq 0.20, \qquad \text{chance margin} = 0.05. \tag{A.37}$$

Threshold sensitivity is evaluated for $\tau \in \{0.05, 0.1, 0.2\}$.

Run aggregation and predictive-validity pipeline. Run-level summaries are collected from `metrics.json`. Canonical metadata are read from `config_resolved.yaml`, and layerwise summaries are extracted from `layerwise.csv`. For the flat aggregation used in Tables A.10–A.12, duplicate $(\texttt{run}, \texttt{seed}, \tau)$ tuples are removed by keeping the latest timestamped run directory. Table A.10 reports Spearman correlation with `test_acc_at_best`. Table A.11 labels a run as positive when either `hard_collapse` or `near_collapse` is true. When $L_c(\tau)$ is not reached, the filled-depth predictor imputes $L + 1$.

Sweep grids reproduced from the study. The exact sweep grids used for the reported tables are as follows.

- **Cora depth sweep (Table A.1):** lazy diffusion, $\alpha = 1.0$, depths $L \in \{2, 4, 8, 16\}$.
- **Cora residual sweep (Table A.2):** lazy diffusion, depth $L = 8$, $\alpha \in \{0.2, 0.5, 1.0\}$.
- **Cora diffusion-family sweep (Table A.3):** $\{\text{rw}, \text{lazy}, \text{ppr}\}$ at $(L, \alpha) = (8, 0.5)$.
- **Cora lazy-strength sweep (Table A.4):** lazy diffusion, $(L, \alpha) = (8, 0.5)$, $\lambda \in \{0.1, 0.3, 0.5, 1.0\}$.
- **SBM depth sweep (Table A.5):** lazy diffusion, $\alpha = 1.0$, depths $L \in \{4, 8, 16, 32\}$.
- **SBM residual sweep (Table A.6):** lazy diffusion, depth $L = 32$, $\alpha \in \{0.2, 0.5, 1.0\}$.

- **SBM diffusion-family sweep (Table A.7):** $\{\text{rw}, \text{lazy}, \text{ppr}\}$ at $(L, \alpha) = (32, 0.5)$.
- **ogbn-arxiv depth sweep (Table A.8):** lazy diffusion, $\alpha = 0.5$, depths $L \in \{2, 4, 8\}$.
- **Breadth matrix (Table A.9):** datasets $\{\text{Cora}, \text{Chameleon}, \text{ogbn-arxiv}\}$, backbones $\{\text{FD-MPNN}, \text{GCN}, \text{GraphSAGE}\}$, depths $L \in \{2, 4, 8\}$, seeds $\{1, 7\}$.
- **Predictive-validity analysis (Tables A.10–A.12):** 207 unique runs after deduplication across datasets, depths, backbones, and intervention settings.

Dataset protocol. The main sweeps use Cora, a synthetic SBM, and ogbn-arxiv; the breadth matrix additionally includes Chameleon. The exposed fixed-diffusion configs use `data.dataset=cora` and `data.dataset=ogbn-arxiv` with `root=data`. The exposed SBM generator uses `sizes=[800,800,800]`, `p_in=0.02`, `p_out=0.002`, `feat_dim=128`, and split `{train:0.6, val:0.2, test:0.2}`.

Training hyperparameters. Across the exposed fixed-diffusion base configs, all runs use `hidden_dim=128`, ReLU activation, learning rate 0.01, weight decay 5×10^{-4}, patience 50, `residual_apply_last=false`, `add_input_proj=true`, and `device=auto`. Maximum epochs are 300 for Cora and SBM, and 200 for ogbn-arxiv. Dropout is dataset-specific: 0.5 for Cora, 0.0 for SBM, and 0.2 for ogbn-arxiv.

Breadth-matrix comparability note. Table A.9 combines the paper-aligned fixed-diffusion model with GCN and GraphSAGE backbones whose effective propagation operators are related to, but not identical with, the operator analyzed in Sections 3–4. In the matrix-generation script, the non-FD PyG backbones are instantiated with the `lazy` tag and $\alpha = 0.5$ for configuration uniformity, but these values are ignored at the model level. For this reason, Table A.9 should be interpreted as a robustness check rather than as a one-to-one theorem-validation table.

L	Test acc.	Near	$L_c(0.1)$ defined	median $L_c(0.1)$
2	0.811 ± 0.003	0%	0%	–
4	0.783 ± 0.008	0%	0%	–
8	0.251 ± 0.093	100%	100%	2
16	0.319 ± 0.000	100%	100%	2

Table A.1: Cora depth sweep under lazy diffusion with $\alpha = 1.0$ (mean±std over 5 seeds). Near-collapse uses `train_acc_at_best` < 0.35.

α	Test acc.	Near	$L_c(0.1)$ defined	median $L_c(0.1)$
0.2	0.805 ± 0.006	0%	0%	–
0.5	0.764 ± 0.021	0%	0%	–
1.0	0.251 ± 0.093	100%	100%	2

Table A.2: Cora residual-mixing sweep at depth $L = 8$ under lazy diffusion (mean±std over 5 seeds). Smaller α improves stability and delays contraction.

Appendix A.3. Core sweep tables

Diffusion	Test acc.	Near
lazy	0.760 ± 0.027	0%
ppr	0.768 ± 0.010	0%
rw	0.770 ± 0.021	0%

Table A.3: Cora diffusion-family sweep at $L = 8$ with $\alpha = 0.5$ (mean±std over 5 seeds).

λ (lazy)	Test acc.
0.1	0.628 ± 0.046
0.3	0.751 ± 0.019
0.5	0.784 ± 0.016
1.0	0.758 ± 0.021

Table A.4: Cora lazy mixing-strength sweep at $L = 8$ with $\alpha = 0.5$ (mean±std over 5 seeds).

L	Test acc.	Near	median $L_c(0.1)$
4	1.000 ± 0.001	0%	–
8	0.329 ± 0.015	100%	2
16	0.333 ± 0.007	100%	2
32	0.336 ± 0.006	100%	2

Table A.5: SBM depth sweep under lazy diffusion with $\alpha = 1.0$ (mean±std over 5 seeds).

α	Test acc.	median $L_c(0.1)$
0.2	0.336 ± 0.006	12
0.5	0.336 ± 0.006	5
1.0	0.336 ± 0.006	2

Table A.6: SBM residual-mixing sweep at $L = 32$ under lazy diffusion (mean±std over 5 seeds). Smaller α delays contraction (larger L_c) but does not recover accuracy in this near-degenerate regime.

Diffusion	Test acc.	median $L_c(0.1)$
lazy	0.336 ± 0.006	5
ppr	0.336 ± 0.006	5
rw	0.336 ± 0.006	5

Table A.7: SBM diffusion-family sweep at $L = 32$ with $\alpha = 0.5$ (mean±std over 5 seeds).

L	Test acc.	Near
2	0.541 ± 0.001	0%
4	0.522 ± 0.002	0%
8	0.411 ± 0.110	20%

Table A.8: ogbn-arxiv depth sweep under lazy diffusion with $\alpha = 0.5$ (mean±std over 5 seeds).

Dataset	Backbone	L=2	L=4	L=8	Near@8
Chameleon	FD-MPNN	0.450 ± 0.040	0.384 ± 0.037	0.344 ± 0.006	0%
Chameleon	GCN	0.359 ± 0.002	0.320 ± 0.012	0.272 ± 0.006	0%
Chameleon	GraphSAGE	0.527 ± 0.011	0.492 ± 0.005	0.364 ± 0.143	50%
Cora	FD-MPNN	0.804 ± 0.011	0.790 ± 0.040	0.745 ± 0.040	0%
Cora	GCN	0.807 ± 0.004	0.801 ± 0.002	0.741 ± 0.040	0%
Cora	GraphSAGE	0.806 ± 0.006	0.800 ± 0.004	0.736 ± 0.021	0%
ogbn-arxiv	FD-MPNN	0.541 ± 0.002	0.525 ± 0.002	0.216 ± 0.000	100%
ogbn-arxiv	GCN	0.502 ± 0.001	0.486 ± 0.005	0.275 ± 0.084	50%
ogbn-arxiv	GraphSAGE	0.505 ± 0.001	0.493 ± 0.002	0.303 ± 0.024	0%

Table A.9: Cross-dataset/backbone matrix (mean±std over seeds $\{1, 7\}$). Fixed-diffusion runs use lazy diffusion with $\alpha = 0.5$ and $\tau = 0.1$.

Appendix A.4. Breadth matrix and predictive validity

Metric	Spearman ρ	p-value
Var(min)	0.810	2.54×10^{-49}
Var(final)	0.767	2.49×10^{-41}
$S_{\text{rel,pre}}$ (min)	0.710	4.58×10^{-33}
$S_{\text{rel,post}}$ (min)	0.696	2.71×10^{-31}
$S_{\text{rel,pre}}$ (final)	0.577	9.72×10^{-20}
$S_{\text{rel,post}}$ (final)	0.555	3.92×10^{-18}
Dirichlet(min)	0.498	2.21×10^{-14}
Dirichlet(final)	0.493	4.38×10^{-14}
L_c fill ($\tau = 0.2$)	0.490	6.46×10^{-14}
L_c fill ($\tau = 0.1$)	0.457	4.53×10^{-11}
Cosine(final)	-0.483	1.72×10^{-13}

Table A.10: Predictive validity across 207 unique runs (duplicates removed): Spearman correlation between diagnostics and test accuracy at the best-validation checkpoint. L_c fill imputes $L + 1$ when $L_c(\tau)$ is not reached.

Metric	ROC-AUC	AUPRC
Var(min)	1.000	1.000
Var(final)	1.000	1.000
$S_{\text{rel,post}}$ (min)	1.000	1.000
$S_{\text{rel,pre}}$ (min)	0.998	0.998
Dirichlet(final)	0.987	0.988
$S_{\text{rel,post}}$ (final)	0.986	0.991
Dirichlet(min)	0.986	0.989
$S_{\text{rel,pre}}$ (final)	0.984	0.989
Cosine(final)	0.979	0.987
L_c fill (τ=0.2)	0.844	0.811
L_c fill (τ=0.1)	0.786	0.763

Table A.11: Binary collapse prediction (hard or near). Scores are oriented so that larger values indicate higher collapse risk (e.g., $-S_{\text{rel}}$ and $-L_c$ fill). Results over 207 unique runs.

τ	L_c defined	median L_c	Spearman ρ	ROC-AUC	AUPRC
0.05	37.7%	3.0	0.305	0.697	0.616
0.10	38.2%	2.0	0.422	0.786	0.763
0.20	38.6%	2.5	0.490	0.844	0.811

Table A.12: Sensitivity of $L_c(\tau)$ to the threshold τ (207 unique runs). Predictive validity uses filled depth (L_c when reached, else $L + 1$).