

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 JACOBIGAD: JACOBI POLYNOMIAL-POWERED HETEROGENEOUS GRAPH-LEVEL ANOMALY DETECTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Heterogeneous graph-level anomaly detection is vital for applications such as fraud detection and drug discovery, yet remains challenging due to mixed features, complex structures, and severe class imbalance. This paper introduces JacobiGAD, a unified framework that addresses these challenges through three key innovations. First, learnable multiscale filters based on Jacobi Polynomials adapt to different node and edge types, fusing multiple graph views to enhance anomaly signals. Second, these polynomials enable efficient approximation of targeted functions and naturally encode diverse geometries. Third, a Ricci Flow-inspired loss amplifies gradients for rare anomalies, mitigating class imbalance without distorting graph embeddings, ensuring stable convergence. Extensive experiments on real-world benchmarks show JacobiGAD outperforms the best baseline by up to 2.79% (AUROC), 7.78% (AUPRC), 7.11% (Recall@k), and 5.96% (F1-score) on average.

1 INTRODUCTION

Graph-level anomaly detection (GAD) identifies entire graphs that exhibit structural or attributively deviations from norm ones, a critical task for applications (Ma et al., 2023; Lin et al., 2024), such as financial fraud detection, drug toxicity screening, and infrastructure monitoring. These graphs are often heterogeneous, containing multiple node and edge types, presenting three core challenges: Mixed feature spaces (Xu et al., 2024). Heterogeneous graphs combine diverse attributes with varying dimensions, making it hard to design unified filters that capture relevant anomalous patterns across all types. (2) Structural complexity (Zhang et al., 2022). Multiple edge types and intricate structures create difficulties in detecting anomalous substructures within graphs, which can easily mislead the detector. (3) Imbalanced label distribution (Dong et al., 2024). Genuine anomalies are exceedingly rare, leading to highly skewed training sets that bias models toward normal graphs. Existing methods for graph classification or GAD struggle to surface anomalous signals in such imbalanced heterogeneous data, as illustrated in Section 2 and demonstrated in Section 5.

To tackle these challenges, this paper presents JacobiGAD, an end-to-end framework that unifies adaptive multi-scale spectral filters with imbalance-sensitive loss for heterogeneous GAD. Specifically, our learnable Jacobi Polynomial filters adapt to diverse node and edge types across multiple graph views, enhancing true anomaly signals while suppressing noise. A complementary Ricci Flow-inspired loss dynamically amplifies gradients for rare anomalies, effectively combating class imbalance. Theoretically, we prove that these filters enable fast, stable approximation while preserving feature distances, and that the loss ensures reliable convergence. Empirically, JacobiGAD consistently outperforms all baselines across AUROC, AUPRC, Recall@k, and F1-score on 15 real-world benchmarks. In summary, our contributions are threefold:

- We propose JacobiGAD, a novel framework for heterogeneous GAD that integrates adaptive Jacobi Polynomial filters with a Ricci Flow-inspired loss function.
- We provide theoretical guarantees on filter stability, information preservation, and loss convergence, ensuring principled and efficient learning.
- We comprehensively validate JacobiGAD on diverse real-world datasets, showcasing its superior ability to detect rare anomalies that existing methods fail to identify.

054 **2 RELATED WORK**

055

056 **Homogeneous Graph Classification.** Early successes in graph classification on homogeneous
 057 networks include GCN (Kipf & Welling, 2017), which approximates spectral graph convolutions,
 058 SAGE (Hamilton et al., 2017), which samples neighborhoods, GAT (Velickovic et al., 2018), which
 059 applies attention to neighbor messages, and GIN (Xu et al., 2019), which demonstrated that sum-
 060 aggregation matches the Weisfeiler–Leman test’s expressivity. Recent extensions such as LRGNN
 061 (Wei et al., 2023) stack GNNs for long-range dependencies, GRDL (Wang & Fan, 2024) treats node
 062 embeddings as discrete distributions for direct classification, UQGNN (Wu et al., 2025) introduces
 063 uncertainty-aware objectives for robustness, and UIL (Sui et al., 2025) offers a unified view on
 064 invariant graph learning. While effective on balanced, homogeneous benchmarks, these models
 065 struggle with integrating multiple node/edge types and detecting rare anomalies in complex structures.

066 **Heterogeneous Graph Classification.** Methods such as HMGNN (Yu & Gao, 2022) and muxGNN
 067 (Melton & Krishnan, 2023) capture heterogeneity using motifs or multiplex networks, while HeGCL
 068 (Shi et al., 2024) employs contrastive learning on multiple views. Subsequent approaches, such as
 069 RFAGNN (Wu et al., 2024) and SHGLNN (Hayat et al., 2024), use relational attention or hypergraphs
 070 to model complex interactions. Although these methods perform heterogeneous graph classification,
 071 they rely on fixed filters or heuristic fusion strategies, assume balanced data, and lack principled
 072 mechanisms for anomaly detection.

073 **Graph-level Anomaly Detection.** Current anomaly detection literature includes iGAD (Zhang et al.,
 074 2022), which learns anomalous substructures in graphs, GmapAD (Ma et al., 2023), which maps
 075 graphs into feature spaces based on similarity to representative nodes, RumorMixer (Xu et al., 2024),
 076 focusing on the echo chamber effect and platform heterogeneity; RQGNN (Dong et al., 2024), which
 077 uses the Rayleigh Quotient to uncover sample properties, and UniGAD (Lin et al., 2024), which
 078 tackles multi-level tasks for diverse information. While these methods perform well in GAD, they
 079 struggle to generalize to heterogeneous scenarios due to their inability to adapt filters to diverse
 080 feature domains, fuse multiple graph views, and incorporate theoretically guaranteed loss for handling
 081 imbalanced data in complex structures.

082 In contrast, JacobiGAD is an innovative end-to-end framework specifically designed for heteroge-
 083 neous GAD. It introduces learnable, multi-scale spectral filters that adaptively fuse signals across
 084 diverse node and edge types, and a Ricci Flow–inspired loss that counteracts class imbalance by
 085 dynamically emphasizing rare anomalies. Unlike homogeneous methods, it natively handles hetero-
 086 geneous complexity; unlike existing heterogeneous classifiers, it uses learned, geometry-aware filters
 087 instead of fixed bases; and unlike all prior approaches, it addresses severe imbalance in a theoretically
 088 grounded manner, enabling the detection of subtle anomalies that are missed by other methods.

089

090 **3 PRELIMINARIES**

091

092 **Heterogeneous Graph.** A heterogeneous graph is defined as $G = (\mathcal{V}, \mathcal{A}, \mathcal{X}, T_V, R_E)$, where the
 093 node set $\mathcal{V} = \bigcup_{t=1}^{|T_V|} \{V_t\}$ comprises $|T_V|$ distinct types of nodes, each endowed with an attribute
 094 matrix $\mathbf{X}_t \in \mathbb{R}^{|V_t| \times d_t}$ in \mathcal{X} . The set of adjacency matrices $\mathcal{A} = \{\mathbf{A}_r\}_{r=1}^{|R_E|}$ with each $\mathbf{A}_r \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$,
 095 encodes the $|R_E|$ relation types by setting $(\mathbf{A}_r)_{ij} = 1$ if nodes $i, j \in \mathcal{V}$ are linked under relation
 096 r , otherwise $(\mathbf{A}_r)_{ij} = 0$. The schema is completed by T_V , the set of node types, and R_E , the set
 097 of relation types, which together satisfy $|T_V| + |R_E| > 2$. In practice, heterogeneous graphs often
 098 exhibit heterogeneity in attribute dimensions, i.e., $d_{T_i} \neq d_{T_j}$ for $T_i \neq T_j$, $T_i, T_j \in T_V$.

099

100 **Task Definition.** Given a heterogeneous graph set $\mathcal{G} = \{G^{(i)} = (\mathcal{V}^{(i)}, \mathcal{A}^{(i)}, \mathcal{X}^{(i)}, T_V, R_E)\}_{i=1}^N$,
 101 we partition \mathcal{G} into two disjoint subsets, anomalous graphs \mathcal{G}^{an} and normal graphs \mathcal{G}^{no} , with
 102 $\mathcal{G}^{an} \cap \mathcal{G}^{no} = \emptyset$. The GAD task then seeks to assign each $G^{(i)} \in \mathcal{G}$ to one of these classes, based on
 103 atypical structural or attribute patterns that distinguish anomalous instances. Beyond the difficulties of
 104 complex feature and structure caused by heterogeneity, heterogeneous GAD also exhibits severe class
 105 imbalance, i.e., $|\mathcal{G}^{an}| \ll |\mathcal{G}^{no}|$, which compounds the difficulty of reliable anomaly discrimination.
 106 Building on this formulation. Our study proposes a novel spectral GNN based on Jacobi Polynomials
 107 under the guidance of Ricci Flow-inspired loss, specifically designed for heterogeneous GAD to
 address the challenges mentioned in Section 1.

108 **Spectral Graph Neural Network.** The key ideas of spectral GNNs are to conduct graph convolutional
 109 operations in the Fourier domain, which can be defined as $g \star \mathbf{X} = g(\mathbf{L})\mathbf{X}$, where $g(\cdot)$ is the graph
 110 filter, \mathbf{X} is the feature matrix of the graph, and \mathbf{L} is the normalized Laplacian matrix, which can be
 111 defined as $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}}$, given the adjacency matrix \mathbf{A} , corresponding degree matrix \mathbf{D} ,
 112 and an identity matrix \mathbf{I} . The successful choices of $g(\cdot)$ from prior work (Defferrard et al., 2016), are
 113 polynomials, inspiring our exploration of the optimal basis of the graph filter in heterogeneous GAD.
 114

115 **Jacobi Polynomials.** Jacobi Polynomials $\{P_n^{\alpha,\beta}(x)\}_{n=0}^{\infty}$ are a family of orthogonal polynomials on
 116 the interval $x \in [-1, 1]$ with weight function $w(x) = (1-x)^\alpha(1+x)^\beta$ for parameters $\alpha, \beta > -1$:
 117

$$\begin{aligned} P_0^{\alpha,\beta}(x) &= 1, \\ P_1^{\alpha,\beta}(x) &= (\alpha+1) + \frac{\alpha+\beta+2}{2}(x-1), \\ P_k^{\alpha,\beta}(x) &= (\theta_k^{(1)}x + \theta_k^{(2)})P_{k-1}^{\alpha,\beta}(x) - \theta_k^{(3)}P_{k-2}^{\alpha,\beta}, \end{aligned}$$

122 where

$$\begin{aligned} \theta_k^{(1)} &= \frac{(2k+\alpha+\beta-1)(2k+\alpha+\beta)}{2k(k+\alpha+\beta)}, \\ \theta_k^{(2)} &= \frac{(2k+\alpha+\beta-1)(\alpha^2-\beta^2)}{2k(k+\alpha+\beta)(2k+\alpha+\beta-2)}, \\ \theta_k^{(3)} &= \frac{(k+\alpha-1)(k+\beta-1)(2k+\alpha+\beta)}{k(k+\alpha+\beta)(2k+\alpha+\beta-2)} \end{aligned}$$

130 Jacobi Polynomials provide a general solution for graph signal filtering. In more detail, increasing α
 131 decreases contributions near the upper end of the spectrum, i.e., high-frequency or rapidly varying
 132 components, while increasing β down-weights contributions near the lower end, i.e., low-frequency
 133 or smooth components. In practice, this parametrization yields an efficient, k -hop localized graph
 134 convolution operator whose passband can be finely tuned by selecting α and β to match the topology
 135 and signal characteristics of diverse graph domains. Special cases include classical polynomials, such
 136 as Legendre Polynomials ($\alpha = \beta = 0$), Chebyshev Polynomials ($\alpha = \beta = -\frac{1}{2}$), and Gegenbauer
 137 Polynomials ($\alpha = \beta = \lambda - \frac{1}{2}$, where $\lambda > -\frac{1}{2}$).

138 **Ricci Flow.** Ricci Flow is a geometric process that deforms a Riemannian metric $g(t)$ according to:
 139

$$\frac{\partial g(t)}{\partial t} = -\gamma \text{Ric}(g(t)),$$

142 where $\text{Ric}(g(t))$ denotes the Ricci curvature tensor and $\gamma \in \mathbb{R}^+$. Under this evolution, regions of high
 143 curvature flatten out, leading to a more uniform geometry. In graphs, edgewise curvature measures
 144 are defined via optimal transport between local neighborhood distributions. A discrete Ricci Flow
 145 then updates edge weights to equalize the curvature across the graph. This curvature-guided objective
 146 counteracts extreme class imbalance without globally distorting the graph representation, ensuring
 147 that rare but structurally distinctive anomalies receive proportionally larger gradient updates.
 148

4 METHOD

4.1 OVERVIEW

153 In this section, we present an overview of JacobiGAD in Figure 1. First, we unify heterogeneous
 154 features via Gaussian projection and construct a multi-view topology in Section 4.2. Additionally,
 155 we demonstrate the distance preservation property of our alignment, as shown in Theorem 1. Next,
 156 we propose JPGNN, a spectral filter based on Jacobi Polynomials, which fuses multiple views
 157 while provably preserving feature and structural information in Section 4.3. Furthermore, we
 158 prove the optimal basis, information preservation, target amplification, multiple spaces, extensive
 159 approximation, and converged approximation properties of JPGNN, as shown in Theorems 2, 3, 4, 5,
 160 6, and 7, respectively. Finally, we introduce RFACE, a Ricci Flow-inspired loss, that intrinsically
 161 adapts to imbalanced distributions in Section 4.4. Moreover, we verify the weight balance and
 convergence guarantee properties of RFACE, as shown in Theorems 8 and 9.

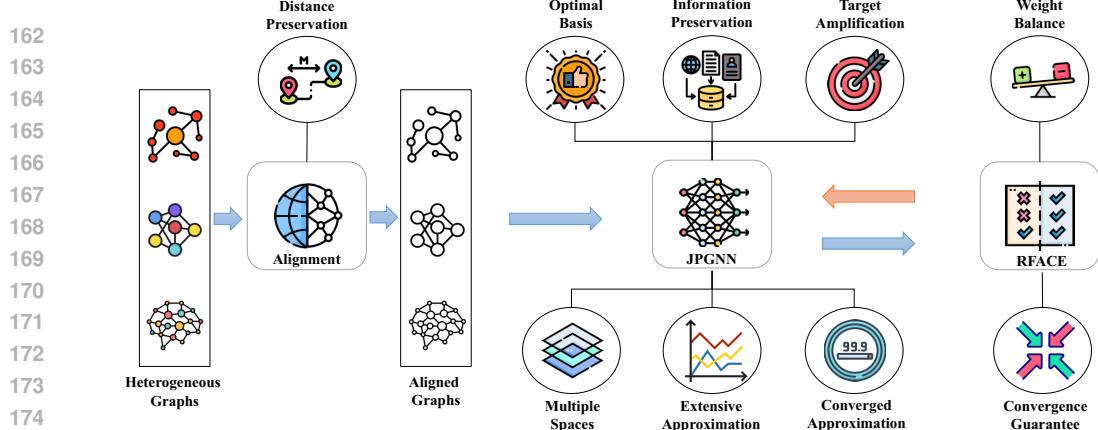


Figure 1: Overview of JacobiGAD.

4.2 HETEROGENEOUS FEATURE PROJECTION AND MULTI-VIEW TOPOLOGY

The initial processing of heterogeneous data is a critical determinant of model performance. Conventional approaches often fail to adequately address two fundamental challenges: the misalignment of features across node types and the synergistic integration of multiple relational contexts. Our method addresses these challenges through a principled framework.

Heterogeneous feature projection. Current methods for handling heterogeneous node features, such as decomposition (Ren & Du, 2020), concatenation (Gao et al., 2023), and independent learning (Yang et al., 2023), each face significant limitations. Decomposition reduces dimensionality at the cost of information loss. Concatenation increases feature size, leading to overfitting and high computational cost. Independent learning ignores semantic alignment between views, introducing noise and raising training expense. All fail to adequately align semantics across node types.

To address the above drawbacks, we propose a theoretically guaranteed approach that can align features without severe information loss and requires no parameter tuning. Given a set of heterogeneous graphs with in total $|T_V|$ types of node features, we first apply zero padding to the end of each type so that all types have the same dimensionality d_{max} , i.e., $\mathbf{X}'_t = \mathbf{X}_t \oplus \mathbf{0}$, where $\mathbf{X}'_t \in \mathbb{R}^{|V_t| \times d_{max}}$ and $\mathbf{0} \in \mathbb{R}^{|V_t| \times (d_{max} - d_t)}$, $\forall t \in T_V$. Then, we use a shared Gaussian matrix to project them into a lower latent space with dimension d , i.e., $(\mathbf{X}'_t)^T = \mathbf{P}(\mathbf{X}'_t)^T$, where $\mathbf{P} \in \mathbb{R}^{d \times d_{max}}$ and $\mathbf{P}_{i,j} \sim \mathcal{N}(0, \frac{1}{d})$. The result feature matrix $\mathbf{X}^{proj} = \mathbf{X}_1^{proj} \oplus \mathbf{X}_2^{proj} \oplus \dots \oplus \mathbf{X}_{|T_V|}^{proj}$. The validity of this approach is rigorously guaranteed by the following theorem, whose proof is in Appendix A.

Theorem 1. *Given any finite set of vectors with different dimensions, zero-padding at any position can equivalently preserve their original information. For any zero-padded vector $\mathbf{x} \in \mathbb{R}^D$, a data-independent Gaussian projection $f(\mathbf{x}) = \mathbf{P}\mathbf{x}$, where $\mathbf{P} \in \mathbb{R}^{d \times D}$ and $\mathbf{P}_{i,j} \sim \mathcal{N}(0, \frac{1}{d})$, can preserve pairwise Euclidean distance for M pairs up to a factor ϵ with high probability, $1 - 2Me^{-\frac{\epsilon^2 d}{4}}$.*

Theorem 1 shows that high-dimensional vectors can be projected into a lower-dimensional space while preserving their pairwise distances with high probability. This ensures the semantic relationships between nodes are maintained isotropically in the latent space, without additional cost during training.

Multi-view topology. Additionally, to address the multifaceted topology of heterogeneous graphs, we move beyond the naive summation of adjacency matrices, which assumes all relation types are equally important. Instead, we employ learnable weights for each relation: $\tilde{\mathbf{A}} = \sum_{r=1}^{|R_E|} \omega_r \mathbf{A}_r$. This allows the model to dynamically discern the hierarchical importance of different relational contexts. However, prior heuristic weighting schemes often fail to leverage inherent structural patterns. Our method, detailed in Section 4.3, provides a theoretically grounded approach for optimal multi-view fusion, ensuring convergence and information preservation to leverage the formulation effectively.

216 4.3 JACOBI POLYNOMIAL-BASED GRAPH NEURAL NETWORK
217

218 Our spectral GNN takes as input a graph with aligned features and fused topology. Selecting an
219 appropriate spectral filter basis is critical for heterogeneous GAD, as different bases offer distinct
220 expressivity. We posit that Jacobi Polynomials are the optimal basis, a claim supported by the
221 following analysis and theorems. First, we consider homogeneous graph classification, a closely
222 related case of heterogeneous GAD. The core of our argument rests on a theorem, proved in Appendix
223 A, establishing the direct and general optimality of Jacobi Polynomials.

224 **Theorem 2.** *Consider the optimization process of a spectral GNN in graph classification:
225 $\operatorname{argmin}_{\theta_k, W} \mathcal{L}(\mathbf{y}, \operatorname{Pooling}(\sum_{k=1}^K \theta_k g_k(\mathbf{L}) \mathbf{X} \mathbf{W}))$. Assuming that it can reach a global minimum
226 by tuning learnable parameters θ_k, W , then the optimal choice of the basis of the graph filter $g_k(\cdot)$
227 can be the Jacobi Polynomials, according to its convergence speed to the minimum area.*

228 Theorem 2 provides the foundational justification for our architecture, demonstrating that Jacobi
229 Polynomials are optimal for the case of homogeneous graph classification. This inspires their use for
230 more complex heterogeneous graph-level tasks. As established in Section 4.2, a heterogeneous graph’s
231 multi-view topology is a linear combination of homogeneous adjacency matrices. Consequently, the
232 task can be represented as a combination of its homogeneous variants. Therefore, the expressivity of
233 Jacobi Polynomials for heterogeneous GAD hinges on their ability to filter and fuse this multi-view
234 information, a capability demonstrated by the following theorems, whose proofs are in Appendix A.

235 **Theorem 3.** *Given different views of a graph, the combination of Jacobi Polynomial-based graph
236 filter can preserve the full information from the original graph due to injectivity.*

237 **Theorem 4.** *Combining information from V views using the combination of Jacobi Polynomial-based
238 graph filter will amplify targeted patterns (the enhancement factor grows as $\Theta(V)$) while suppressing
239 noise (the signal-to-noise ratio grows as $\Theta(\sqrt{V})$).*

241 Theorem 3 and 4 demonstrate that a Jacobi Polynomial-based filter comprehensively preserves
242 information while selectively amplifying targeted patterns and reducing noise. This is vital for
243 heterogeneous GAD, where anomalies are often subtle inconsistencies across relational views. Unlike
244 filters that may smooth over these faint cues, our Jacobi basis can be tuned to amplify cross-view
245 discrepancies while dampening common normal signals.

246 However, a filter constrained to a Euclidean prior is insufficient, as anomalies can exhibit complex
247 structures such as hierarchical or cyclical patterns (Dong et al., 2025). Effective heterogeneous
248 GAD thus requires a filter capable of leveraging multi-geometric information from Hyperbolic
249 (for hierarchical data) and Spherical (for cyclical data) spaces. The following theorem, proved in
250 Appendix A, establishes that Jacobi Polynomials possess this essential capability.

251 **Theorem 5.** *After appropriate coordinate transformations, Jacobi Polynomials can serve as eigen-
252 functions of the Laplace-Beltrami operator in the κ -stereographic model (Bachmann et al., 2020).
253 The connections for each geometry are as follows:*

- 254 • *Spherical geometry ($\kappa > 0$): The Laplace-Beltrami operator in stereographic coordinates has
255 eigenfunctions with radial and angular parts. The angular part is handled by spherical harmonics,
256 while the radial part satisfies a differential equation solvable by Jacobi Polynomials.*
- 257 • *Hyperbolic geometry ($\kappa < 0$): The spectrum of the Laplace-Beltrami operator is continuous, and
258 the radial eigenfunctions are not polynomials but can be expressed as Jacobi Polynomials.*
- 259 • *Euclidean geometry ($\kappa = 0$): The Laplace-Beltrami operator reduces to the standard Laplacian,
260 and the radial eigenfunctions are Bessel functions, which arise as a limit of Jacobi Polynomials.*

263 Theorem 5 elevates our model beyond Euclidean-centric approaches. By adjusting its parameters
264 (α, β), the Jacobi filter performs a soft selection of the optimal geometric domain for the fused graph’s
265 structure. This enables a single model to detect anomalies manifesting in any of these paradigms, a
266 critical capability for complex real-world heterogeneous graphs.

267 In summary, Jacobi Polynomials offer key advantages for our task: performance guarantee, effective
268 multi-view fusion, and adaptability to complex structural patterns. This naturally raises the question
269 of whether a Jacobi Polynomial-based GNN can converge efficiently during training. We address this
with the following theorems, which are demonstrated in Appendix A.

270 **Theorem 6.** Assuming using Jacobi Polynomials as graph filter $g(\cdot)$, and the eigenvalues of the
 271 shifted Laplacian matrix \mathbf{L} fall in $[-1, 1]$, then $g(\mathbf{L})$ can approximate any continuous function lying in
 272 the space $C[-1, 1]$ (contains continuous functions on $[-1, 1]$). Moreover, it can also approximate any
 273 function in the $L_w^2[-1, 1]$ space (contains measurable functions satisfying $\int_{-1}^1 |f(x)|^2 w(x) dx < \infty$,
 274 where $w(x) = (1-x)^\alpha (1+x)^\beta$, and $\alpha, \beta > -1$).
 275

276 **Theorem 7.** Jacobi Polynomials satisfy sharp approximation bounds. In particular, if the function
 277 $f(x)$ has r continuous derivatives, then there exists a constant C depending on r, α, β such that
 278 the Jacobi Polynomials $g(x)$ obeys $\min_{\deg(g(x)) \leq N} \|f(x) - g(x)\|_\infty \leq \frac{C}{N} \|f^{(r)}(x)\|_{L_w^1}$, which
 279 guarantees that a low-degree Jacobi filter will approximate $f(x)$ well.

280 Theorems 6 and 7 guarantee our model’s high expressiveness and computational efficiency. A low-
 281 order polynomial suffices to capture complex patterns, enabling a shallow architecture that avoids the
 282 over-smoothing typical of deep GNNs, a critical advantage for preserving the fine-grained differences
 283 between normal and anomalous graphs in heterogeneous GAD.

284 Based on Theorems 2–7, Jacobi Polynomials are theoretically justified as an optimal basis for
 285 heterogeneous GAD. We therefore operationalize this framework into a neural network layer, adhering
 286 to the parameter constraints specified in Theorem 6. We first transform the input adjacency matrix
 287 $\tilde{\mathbf{A}} = \sum_{r=1}^{|R_E|} \omega_r \mathbf{A}_r$ to the normalized Laplacian matrix $\tilde{\mathbf{L}} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \tilde{\mathbf{A}} \mathbf{D}^{-\frac{1}{2}}$, and rescale the
 288 normalized Laplacian matrix $\hat{\mathbf{L}} = \frac{2}{\lambda_{\max}} \tilde{\mathbf{L}} - \mathbf{I}$, where λ_{\max} is the largest eigenvalue of $\tilde{\mathbf{L}}$. Then the
 289 k -th layer of the Jacobi Polynomial-based Graph Neural Network (JPGNN) can be defined as:
 290

$$291 \mathbf{H}^{(k)} = \sigma\left(\left(\sum_{t=0}^T \theta_t^{(k)} P_t^{(\alpha^{(k)}, \beta^{(k)})}(\hat{\mathbf{L}})\right) \mathbf{H}^{(k-1)} \mathbf{W}^{(k)}\right),$$

294 where $\mathbf{H}^{(0)} = \sigma(\mathbf{X}^{proj} \mathbf{W}^{(0)})$, σ is a activation function, and $\omega_r, \alpha^{(k)}, \beta^{(k)}, \theta_t^{(k)}, \mathbf{W}^{(k)}$ are learn-
 295 able parameters. Then, the graph embedding \mathbf{z} can be obtained by:
 296

$$297 \mathbf{H}^{stack} = \mathbf{H}^{(0)} \oplus \mathbf{H}^{(1)} \oplus \dots \oplus \mathbf{H}^{(K)},$$

$$298 \mathbf{H} = \sigma(\mathbf{H}^{stack} \mathbf{W}),$$

$$299 \mathbf{z} = \text{Pooling}(\mathbf{H})$$

300 where \mathbf{W} is a learnable parameter. This design yields a fully co-adaptive model: the multi-view
 301 fusion, governed by view weights ω_r , and the spectral processing via JPGNN are jointly optimized to
 302 excel at heterogeneous GAD.
 303

304 4.4 RICCI FLOW-INSPIRED LOSS FUNCTION

306 The above design addresses the first two challenges outlined in Section 1, while the final component
 307 of our framework tackles the severe class imbalance in heterogeneous GAD, where normal graphs
 308 significantly outnumber anomalies. A standard Cross-Entropy loss is ill-suited for this scenario, as
 309 it can become dominated by the majority class. To counteract this, we introduce the Ricci Flow
 310 Adjusted Cross-Entropy Loss (RFACE), which dynamically reshapes the learning landscape based on
 311 the model’s output geometry.
 312

313 For a graph-level classification task with C classes, given predicted probability of i -th sample
 $\mathbf{p}_i = \text{Sigmoid}(\mathbf{z}_i)$, the Cross-Entropy loss is:
 314

$$315 \mathcal{L}_{CE} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C \mathbf{y}_{i,c} \log(\mathbf{p}_{i,c})$$

318 In highly imbalanced settings (e.g., $C = 2$), the standard Cross-Entropy loss, \mathcal{L}_{CE} , produces much
 319 larger gradients for the frequent class. This biases model updates toward the majority class, often
 320 harming minority class performance. To counteract this, we adapt principles from differential
 321 geometry, mimicking the Ricci Flow, which homogenizes a manifold’s curvature. We apply this
 322 concept to the loss landscape’s curvature per class, defined for a class c as:
 323

$$\kappa_c = \log\left(\frac{f_c}{\max_{c'} f_{c'} + \epsilon}\right),$$

324 Table 1: Average performance with multiple runs (homogeneous graph classification models).
325

| Datasets | Metrics | GCN | SAGE | GAT | GIN | LRGNN | GRDL | UQGNN | UIL | JacobiGAD |
|-------------|----------|--------|--------|--------|--------|--------|--------|--------|--------|---------------|
| SF-295 | AUROC | 0.6687 | 0.7178 | 0.7409 | 0.6914 | 0.7578 | 0.6389 | 0.5248 | 0.7334 | 0.7729 |
| | AUPRC | 0.0856 | 0.1600 | 0.1645 | 0.0961 | 0.1962 | 0.0871 | 0.0525 | 0.1598 | 0.2623 |
| | Recall@k | 0.1078 | 0.2362 | 0.2099 | 0.1029 | 0.2494 | 0.1342 | 0.0519 | 0.1975 | 0.3210 |
| | F1-score | 0.4870 | 0.5685 | 0.5560 | 0.5255 | 0.4952 | 0.4871 | 0.4871 | 0.5223 | 0.6356 |
| SN12C | AUROC | 0.7034 | 0.7440 | 0.7475 | 0.7194 | 0.7747 | 0.5639 | 0.4922 | 0.7604 | 0.7797 |
| | AUPRC | 0.1032 | 0.1598 | 0.1772 | 0.1122 | 0.2038 | 0.0792 | 0.0486 | 0.1789 | 0.2666 |
| | Recall@k | 0.1279 | 0.2430 | 0.2234 | 0.1364 | 0.2515 | 0.1168 | 0.0537 | 0.2702 | 0.3240 |
| | F1-score | 0.4905 | 0.5846 | 0.5791 | 0.5242 | 0.5001 | 0.4875 | 0.4875 | 0.5105 | 0.6329 |
| UACC257 | AUROC | 0.6654 | 0.7324 | 0.6998 | 0.6848 | 0.7220 | 0.5711 | 0.4611 | 0.7124 | 0.7613 |
| | AUPRC | 0.0726 | 0.1588 | 0.1228 | 0.0903 | 0.1626 | 0.0916 | 0.0378 | 0.1134 | 0.1995 |
| | Recall@k | 0.0832 | 0.2404 | 0.1633 | 0.1136 | 0.2312 | 0.1471 | 0.0355 | 0.1562 | 0.2819 |
| | F1-score | 0.4921 | 0.5722 | 0.5214 | 0.5096 | 0.4905 | 0.4895 | 0.4895 | 0.4955 | 0.6246 |
| DBLP | AUROC | 0.9816 | 0.9746 | 0.9671 | 0.9730 | 0.7475 | 0.9800 | 0.9352 | 0.9698 | 0.9830 |
| | AUPRC | 0.9829 | 0.9758 | 0.9515 | 0.9740 | 0.6285 | 0.9761 | 0.9497 | 0.9701 | 0.9842 |
| | Recall@k | 0.9418 | 0.9441 | 0.9172 | 0.9284 | 0.6197 | 0.9374 | 0.8926 | 0.9172 | 0.9575 |
| | F1-score | 0.9598 | 0.9594 | 0.9346 | 0.9445 | 0.4863 | 0.9492 | 0.9268 | 0.3959 | 0.9651 |
| IMDB | AUROC | 0.6601 | 0.6771 | 0.6707 | 0.6601 | 0.6643 | 0.6677 | 0.6575 | 0.6487 | 0.7263 |
| | AUPRC | 0.7007 | 0.7260 | 0.7161 | 0.6948 | 0.7063 | 0.6878 | 0.6982 | 0.6764 | 0.7619 |
| | Recall@k | 0.7056 | 0.6982 | 0.6824 | 0.6982 | 0.6772 | 0.6909 | 0.6993 | 0.7003 | 0.7192 |
| | F1-score | 0.6387 | 0.6225 | 0.6045 | 0.6363 | 0.6242 | 0.6258 | 0.6305 | 0.6100 | 0.6585 |
| PDNS | AUROC | 0.7773 | 0.8577 | 0.6735 | 0.6249 | 0.8159 | 0.6935 | 0.4377 | 0.5683 | 0.8728 |
| | AUPRC | 0.4434 | 0.6110 | 0.3263 | 0.2349 | 0.5188 | 0.3578 | 0.1565 | 0.2224 | 0.6871 |
| | Recall@k | 0.4788 | 0.5900 | 0.3444 | 0.2766 | 0.5206 | 0.3299 | 0.1429 | 0.2205 | 0.6283 |
| | F1-score | 0.6743 | 0.7561 | 0.4917 | 0.4526 | 0.5718 | 0.4553 | 0.4526 | 0.4526 | 0.7760 |
| RCDD | AUROC | 0.9581 | 0.9811 | 0.9602 | 0.9658 | 0.9805 | 0.9609 | 0.8033 | 0.9593 | 0.9826 |
| | AUPRC | 0.8619 | 0.9291 | 0.8871 | 0.8823 | 0.9267 | 0.8605 | 0.3491 | 0.8645 | 0.9332 |
| | Recall@k | 0.8006 | 0.8695 | 0.8261 | 0.8249 | 0.8743 | 0.8111 | 0.4106 | 0.8081 | 0.8747 |
| | F1-score | 0.8782 | 0.9230 | 0.8995 | 0.8981 | 0.9280 | 0.8271 | 0.4616 | 0.8694 | 0.9280 |
| Transaction | AUROC | 0.9085 | 0.9437 | 0.9216 | 0.9202 | 0.9461 | 0.8773 | 0.7162 | 0.9245 | 0.9543 |
| | AUPRC | 0.3722 | 0.4811 | 0.4520 | 0.3961 | 0.5063 | 0.3688 | 0.0927 | 0.4376 | 0.5642 |
| | Recall@k | 0.3410 | 0.5172 | 0.4462 | 0.3730 | 0.5217 | 0.3730 | 0.0915 | 0.4577 | 0.5835 |
| | F1-score | 0.6502 | 0.7469 | 0.7088 | 0.6552 | 0.7762 | 0.6138 | 0.4105 | 0.6785 | 0.7944 |

349 where f_c is the frequency of class c in training set, and ϵ is for numerical stability. To adjust the
350 gradients based on Ricci Flow, we further define the Ricci Flow adjustment term for the i -th sample:
351

$$\Delta \mathbf{p}_{i,c} = -\gamma \kappa_c \nabla_{\mathbf{p}_{i,c}} \mathcal{L}_{CE},$$

353 where γ is a hyperparameter, and the RFACE can be defined as:
354

$$\mathcal{L}_{RFACE} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C \mathbf{y}_{i,c} \log(\tilde{\mathbf{p}}_{i,c}),$$

355 where $\tilde{\mathbf{p}}_{i,c} = \text{Sigmoid}(\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c})$. The following theorems proves the benefits of utilizing RFACE
356 as the training objective for heterogeneous GAD, demonstrated in Appendix A:
357

358 **Theorem 8.** For a rare class c , $|\nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{RFACE}| > |\nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{CE}|$ with amplifying factor proportional
359 to $\gamma |\kappa_c|$ and the amplification follows $(1 + \gamma |\kappa_c|)$ -Lipschitz continuous, preserving the topology of
360 the latent graph embedding space.
361

362 **Theorem 9.** When the adjusted predictions are perfect, i.e., $\tilde{\mathbf{p}}_{i,c} = \mathbf{y}_{i,c}, \forall i, c$, the adjustment term
363 vanishes, i.e., $\Delta \mathbf{p}_{i,c} \rightarrow 0$, and the raw predictions also converge to the true labels, i.e., $\mathbf{p}_{i,c} \rightarrow \mathbf{y}_{i,c}$
364 for all classes, including rare ones.
365

366 In summary, the RFACE is a dynamic system that actively recalibrates the learning focus based on
367 per-class performance, not a simple weighting scheme. This ensures our JPGNN is optimized for
368 detecting rare anomalies, making the entire pipeline, from feature projection and topology fusion
369 to spectral filtering, coherent and optimal for the task. Beyond its theoretical foundation, extensive
370 experiments in Section 5 confirm the practical superiority of JacobiGAD.
371

5 EXPERIMENTS

5.1 EXPERIMENTAL SETUP

375 **Datasets and baselines.** We evaluate JacobiGAD on 14 public and 1 private real-world datasets,
376 divided into 20%/20%/60% for train/validation/test, and compare our JacobiGAD with 18 baselines
377

378 Table 2: Average performance with multiple runs (heterogeneous graph classification models).
379

| Datasets | Metrics | HMGN | muxGNN | HeGCL | RFAGNN | SHGLNN | JacobiGAD |
|-------------|----------|--------|--------|--------|--------|--------|---------------|
| SF-295 | AUROC | 0.4112 | 0.4348 | 0.6584 | 0.6799 | 0.5088 | 0.7729 |
| | AUPRC | 0.0421 | 0.0417 | 0.1129 | 0.1090 | 0.0471 | 0.2623 |
| | Recall@k | 0.0477 | 0.0230 | 0.1745 | 0.1802 | 0.0198 | 0.3210 |
| | F1-score | 0.4871 | 0.4871 | 0.4876 | 0.4871 | 0.4871 | 0.6356 |
| SN12C | AUROC | 0.4309 | 0.5169 | 0.6373 | 0.6958 | 0.5053 | 0.7797 |
| | AUPRC | 0.0410 | 0.0784 | 0.0894 | 0.1077 | 0.0458 | 0.2666 |
| | Recall@k | 0.0315 | 0.1347 | 0.1287 | 0.1577 | 0.0188 | 0.3240 |
| | F1-score | 0.4892 | 0.5383 | 0.4874 | 0.4920 | 0.4875 | 0.6329 |
| UACC257 | AUROC | 0.5512 | 0.4207 | 0.6835 | 0.7129 | 0.5022 | 0.7613 |
| | AUPRC | 0.0698 | 0.0332 | 0.1305 | 0.1297 | 0.0381 | 0.1995 |
| | Recall@k | 0.1014 | 0.0142 | 0.2049 | 0.1755 | 0.0132 | 0.2819 |
| | F1-score | 0.4895 | 0.4895 | 0.4962 | 0.4953 | 0.4895 | 0.6246 |
| DBLP | AUROC | 0.4849 | 0.9697 | 0.9696 | 0.9814 | 0.7684 | 0.9830 |
| | AUPRC | 0.3786 | 0.9697 | 0.9698 | 0.9826 | 0.5949 | 0.9842 |
| | Recall@k | 0.3602 | 0.9172 | 0.9306 | 0.9530 | 0.5996 | 0.9575 |
| | F1-score | 0.5079 | 0.9388 | 0.9522 | 0.9595 | 0.6763 | 0.9651 |
| IMDB | AUROC | 0.5256 | 0.6176 | 0.6512 | 0.6594 | 0.5220 | 0.7263 |
| | AUPRC | 0.6063 | 0.6716 | 0.7033 | 0.7119 | 0.5978 | 0.7619 |
| | Recall@k | 0.5983 | 0.6572 | 0.6709 | 0.6909 | 0.5889 | 0.7192 |
| | F1-score | 0.3682 | 0.5171 | 0.6089 | 0.6214 | 0.3682 | 0.6585 |
| PDNS | AUROC | 0.5563 | 0.6250 | 0.7796 | 0.7359 | 0.5173 | 0.8728 |
| | AUPRC | 0.2115 | 0.2528 | 0.4190 | 0.3977 | 0.1912 | 0.6871 |
| | Recall@k | 0.2420 | 0.2855 | 0.4583 | 0.3980 | 0.2470 | 0.6283 |
| | F1-score | 0.4525 | 0.5289 | 0.5432 | 0.6194 | 0.4526 | 0.7760 |
| RCDD | AUROC | 0.7105 | 0.9523 | 0.9390 | 0.9809 | 0.5819 | 0.9826 |
| | AUPRC | 0.2848 | 0.8470 | 0.8031 | 0.9219 | 0.1531 | 0.9332 |
| | Recall@k | 0.2911 | 0.7870 | 0.7366 | 0.8645 | 0.0739 | 0.8747 |
| | F1-score | 0.4625 | 0.8610 | 0.8404 | 0.9182 | 0.4614 | 0.9280 |
| Transaction | AUROC | 0.6409 | 0.8415 | 0.8853 | 0.9338 | 0.5745 | 0.9543 |
| | AUPRC | 0.0773 | 0.3302 | 0.3781 | 0.4331 | 0.0533 | 0.5642 |
| | Recall@k | 0.0572 | 0.3021 | 0.3753 | 0.4348 | 0.0984 | 0.5835 |
| | F1-score | 0.4118 | 0.6241 | 0.6447 | 0.6690 | 0.3504 | 0.7944 |

402 Table 3: Average performance with multiple runs (GAD models).
403

| Datasets | Metrics | iGAD | GmapAD | RumorMixer | RQGNN | UniGAD | JacobiGAD |
|-------------|----------|--------|--------|------------|--------|--------|---------------|
| SF-295 | AUROC | 0.6768 | 0.6190 | 0.4092 | 0.7657 | 0.5947 | 0.7729 |
| | AUPRC | 0.1040 | 0.0670 | 0.0414 | 0.1938 | 0.0724 | 0.2623 |
| | Recall@k | 0.1531 | 0.0724 | 0.0280 | 0.2683 | 0.1095 | 0.3210 |
| | F1-score | 0.5427 | 0.4095 | 0.4871 | 0.6154 | 0.4971 | 0.6356 |
| SN12C | AUROC | 0.7416 | 0.5957 | 0.3549 | 0.7695 | 0.6281 | 0.7797 |
| | AUPRC | 0.1581 | 0.0605 | 0.0353 | 0.1973 | 0.0769 | 0.2666 |
| | Recall@k | 0.2242 | 0.0733 | 0.0290 | 0.2558 | 0.1151 | 0.3240 |
| | F1-score | 0.5476 | 0.3477 | 0.4875 | 0.5844 | 0.4756 | 0.6329 |
| UACC257 | AUROC | 0.7404 | 0.5936 | 0.4997 | 0.7599 | 0.5973 | 0.7613 |
| | AUPRC | 0.1323 | 0.0507 | 0.0411 | 0.1894 | 0.0672 | 0.1995 |
| | Recall@k | 0.2140 | 0.0527 | 0.0456 | 0.2465 | 0.1176 | 0.2819 |
| | F1-score | 0.5429 | 0.3461 | 0.4895 | 0.6064 | 0.5058 | 0.6246 |
| DBLP | AUROC | 0.9791 | 0.5551 | 0.5000 | 0.9804 | 0.9644 | 0.9830 |
| | AUPRC | 0.9803 | 0.4131 | 0.3835 | 0.9829 | 0.9541 | 0.9842 |
| | Recall@k | 0.9396 | 0.4452 | 0.3792 | 0.9463 | 0.8881 | 0.9575 |
| | F1-score | 0.9598 | 0.5479 | 0.2772 | 0.9509 | 0.9122 | 0.9651 |
| IMDB | AUROC | 0.6530 | 0.5079 | 0.4989 | 0.6707 | 0.6528 | 0.7263 |
| | AUPRC | 0.6971 | 0.5866 | 0.5822 | 0.7254 | 0.6902 | 0.7619 |
| | Recall@k | 0.6909 | 0.5794 | 0.5783 | 0.6845 | 0.6982 | 0.7192 |
| | F1-score | 0.6313 | 0.5073 | 0.3681 | 0.6294 | 0.6332 | 0.6585 |
| PDNS | AUROC | 0.8502 | 0.5173 | 0.5928 | 0.7550 | 0.7310 | 0.8728 |
| | AUPRC | 0.6399 | 0.1810 | 0.2606 | 0.4109 | 0.4293 | 0.6871 |
| | Recall@k | 0.5870 | 0.1999 | 0.2732 | 0.4309 | 0.4146 | 0.6283 |
| | F1-score | 0.7308 | 0.5058 | 0.4526 | 0.5420 | 0.6240 | 0.7760 |
| RCDD | AUROC | 0.9794 | 0.7895 | 0.7335 | 0.9624 | 0.9561 | 0.9826 |
| | AUPRC | 0.9225 | 0.3065 | 0.3134 | 0.8740 | 0.8636 | 0.9332 |
| | Recall@k | 0.8611 | 0.3288 | 0.4513 | 0.8073 | 0.7960 | 0.8747 |
| | F1-score | 0.9190 | 0.6463 | 0.4614 | 0.8874 | 0.8541 | 0.9280 |
| Transaction | AUROC | 0.9431 | 0.7384 | 0.6422 | 0.9353 | 0.8862 | 0.9543 |
| | AUPRC | 0.4626 | 0.0873 | 0.0992 | 0.4939 | 0.3505 | 0.5642 |
| | Recall@k | 0.4897 | 0.0709 | 0.1030 | 0.4966 | 0.3501 | 0.5835 |
| | F1-score | 0.7152 | 0.3875 | 0.4213 | 0.7422 | 0.6307 | 0.7944 |

426
427 in the related area. Details can be found in Appendix B. Due to the limited space, we present results
428 of 7 public and 1 private datasets in Section 5, and those of the other 7 public datasets in Appendix E.
429430 **Experimental Settings.** We ensure a fair evaluation by standardizing our approach: baseline models
431 use code from GitHub and their authors' recommended hyperparameters. Note that, since the most
432 commonly used three GNNs, GCN, SAGE, and GAT, are designed for node classification tasks, we

thus implement them with Pytorch_Geometric package and the weighted Cross-Entropy Loss, using the default hyperparameters. JacobiGAD’s hyperparameters are rigorously tuned via grid search to maximize validation performance (summed AUROC/AUPRC/Recall@k/F1-score). Configurations are listed in Appendix D.

5.2 EXPERIMENTAL RESULTS

We conduct a comprehensive comparison of JacobiGAD against three major groups of competing methods: 8 widely used homogeneous graph classification models, 5 representative heterogeneous graph classification approaches, and 5 novel graph-level anomaly detection methods. The results across 8 datasets are summarized in Tables 1, 2, and 3. We elaborate on our findings in detail next.

To begin with, Table 1 demonstrates that JacobiGAD consistently surpasses classical homogeneous GNN architectures, including GCN, SAGE, GAT, and GIN. These baselines, although foundational, remain surprisingly competitive compared with several more advanced techniques. Remarkably, newer homogeneous GNNs, such as LRGNN, GRDL, UQGNN, and UIL, do not perform as well, frequently falling behind even the simpler models. Their limited performance on heterogeneous graph anomaly detection can be attributed to two main issues: they cannot adaptively integrate information across multiple semantic views well, and they lack mechanisms to properly address the severe class imbalance inherent in GAD tasks.

We then compare JacobiGAD with contemporary heterogeneous graph classification methods, including HMGNN, muxGNN, HeGCL, RFAGNN, and SHGLNN. As shown in Table 2, JacobiGAD consistently yields better detection accuracy across all datasets. Although these models are designed specifically for heterogeneous graphs, their representation learning pipelines often rely on fixed or suboptimal strategies for combining heterogeneous modalities, limiting their expressiveness. Thus, Such drawbacks may distort the graph information, especially when running on complex real-world heterogeneous graphs, leading to sometimes inferior performance, even compared to state-of-the-art homogeneous models. Moreover, most of them do not explicitly mitigate data imbalance, which is especially detrimental in anomaly detection scenarios where abnormal samples are extremely scarce.

Finally, we benchmark against the dedicated graph anomaly detection methods iGAD, GmapAD, RumorMixer, RQGNN, and UniGAD. Their comparative performance, reported in Table 3, indicates that JacobiGAD achieves substantially stronger detection capability. These GAD models are tailored for specific anomaly settings, primarily in homogeneous graphs, and therefore struggle with our target task. Their architectures generally lack the capacity to jointly capture multi-view semantic signals and the high-order structural irregularities that characterize anomalies in heterogeneous graphs. Consequently, even though they are specialized for anomaly detection, their design inherently limits their applicability in the heterogeneous graph setting considered in this work.

5.3 ABLATION STUDY

We further examine the influence of key components in JacboGAD, i.e., \mathcal{L}_{RFACE} , tunable Jacobi Polynomial parameters α, β , and learnable view weights ω_r . As shown in Table 4, the ablation study demonstrates the critical contribution of each proposed component to the overall performance of the JacobiGAD. Using \mathcal{L}_{CE} to replace \mathcal{L}_{RFACE} results in significant and consistent performance degradation across all datasets, underscoring its vital role in effectively tackling imbalanced issues in heterogeneous GAD tasks. The learnable parameters (α, β) of Jacobi Polynomials also prove essential, as the fix of them leads to a clear decline in performance, which shows that, without a flexible enough graph filter, the model can not handle the complex information within heterogeneous graphs. Similarly, the learnable relation weight ω_r contributes positively, with its fixing causing noticeable dips, demonstrating the importance of adaptive weights for different relations in the heterogeneous graphs. To sum up, the full model consistently outperforms all ablated variants, confirming that all three components work in concert to achieve state-of-the-art anomaly detection performance across diverse heterogeneous graph datasets.

5.4 HYPERPARAMETER ANALYSIS

Figure 2 reports the AUROC, AUPRC, Recall@k, and F1-score of JacobiGAD on the RCDD dataset as we vary $\eta, h_{\text{dim}}, K, T, \epsilon, \gamma$, where η is the learning rate, h_{dim} is the hidden dimension of

Table 4: Ablation study for component deactivation.

| Datasets | Metrics | JacobiGAD w/o \mathcal{L}_{RFACE} | w/o learnable (α, β) | w/o learnable ω_r |
|----------|----------|-------------------------------------|---------------------------------|--------------------------|
| SF-295 | AUROC | 0.7729 | 0.7591 | 0.7727 |
| | AUPRC | 0.2623 | 0.1953 | 0.2172 |
| | Recall@k | 0.3210 | 0.2815 | 0.2905 |
| | F1-score | 0.6356 | 0.5984 | 0.6161 |
| SN12C | AUROC | 0.7797 | 0.7463 | 0.7505 |
| | AUPRC | 0.2666 | 0.2345 | 0.2119 |
| | Recall@k | 0.3240 | 0.3035 | 0.2933 |
| | F1-score | 0.6329 | 0.6129 | 0.5733 |
| UACC257 | AUROC | 0.7613 | 0.7282 | 0.7365 |
| | AUPRC | 0.1995 | 0.1617 | 0.1713 |
| | Recall@k | 0.2819 | 0.2241 | 0.2475 |
| | F1-score | 0.6246 | 0.5634 | 0.5829 |
| DBLP | AUROC | 0.9830 | 0.9820 | 0.9745 |
| | AUPRC | 0.9842 | 0.9831 | 0.9762 |
| | Recall@k | 0.9575 | 0.9530 | 0.9441 |
| | F1-score | 0.9651 | 0.9641 | 0.9576 |
| IMDB | AUROC | 0.7263 | 0.6962 | 0.7116 |
| | AUPRC | 0.7619 | 0.7313 | 0.7514 |
| | Recall@k | 0.7192 | 0.7045 | 0.7108 |
| | F1-score | 0.6585 | 0.6439 | 0.6325 |
| PDNS | AUROC | 0.8728 | 0.8691 | 0.8728 |
| | AUPRC | 0.6871 | 0.6749 | 0.6812 |
| | Recall@k | 0.6283 | 0.6149 | 0.6239 |
| | F1-score | 0.7760 | 0.7695 | 0.7685 |
| RCDD | AUROC | 0.9826 | 0.9808 | 0.9814 |
| | AUPRC | 0.9332 | 0.9325 | 0.9318 |
| | Recall@k | 0.8747 | 0.8749 | 0.8739 |
| | F1-score | 0.9280 | 0.9274 | 0.9264 |

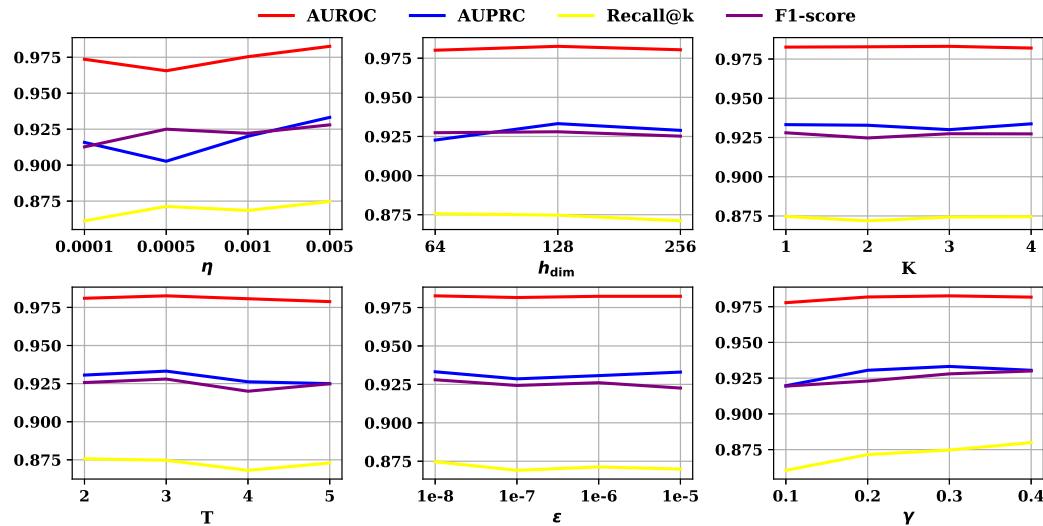


Figure 2: The change of performance on RCDD when varying different hyperparameters.

JacobiGAD, K , T are the width and depth of JacobiGAD respectively, ϵ is the small value to keep κ_c in RFACE valid, and γ is the adjusted hyperparameter in RFACE. As shown in Figure 2, JacobiGAD remains stable when varying the hyperparameters, demonstrating its stability.

6 CONCLUSION

This paper proposed JacobiGAD, a novel framework for heterogeneous GAD. Our approach integrates a theoretically grounded random projection for feature alignment, a Jacobi Polynomial-based spectral GNN for superior multi-view fusion and cross-geometric representation learning, and a Ricci Flow-inspired loss that dynamically counteracts class imbalance. Supported by strong theoretical guarantees and extensive experimental validation, JacobiGAD establishes a new state-of-the-art, providing a powerful and principled methodology for GAD on complex heterogeneous graphs.

540 REFERENCES
541

542 Gregor Bachmann, Gary Bécigneul, and Octavian Ganea. Constant curvature graph convolutional
543 networks. In *ICML*, pp. 486–496, 2020.

544 Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on
545 graphs with fast localized spectral filtering. In Daniel D. Lee, Masashi Sugiyama, Ulrike von
546 Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *NeurIPS*, pp. 3837–3845, 2016.

547 Xiangyu Dong, Xingyi Zhang, and Sibo Wang. Rayleigh quotient graph neural networks for graph-
548 level anomaly detection. In *ICLR*, 2024.

549 Xiangyu Dong, Xingyi Zhang, Lei Chen, Mingxuan Yuan, and Sibo Wang. Spacegnn: Multi-space
550 graph neural network for node anomaly detection with extremely limited labels. In *ICLR*, 2025.

551 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. In
552 *ICLR Workshop*, 2019.

553 Shen Gao, Haotong Zhang, Xiuying Chen, Chongyang Tao, Dongyan Zhao, and Rui Yan. A trend of
554 AI conference convergence in similarity: An empirical study through trans-temporal heterogeneous
555 graph. *IEEE Trans. Knowl. Data Eng.*, 35:9642–9655, 2023.

556 William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large
557 graphs. In *NeurIPS*, pp. 1024–1034, 2017.

558 Malik Khizar Hayat, Shan Xue, and Jian Yang. Self-supervised heterogeneous hypergraph learning
559 with context-aware pooling for graph-level classification. In *ICDM*, pp. 140–149, 2024.

560 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.
561 In *ICLR*, 2017.

562 Yiqing Lin, Jianheng Tang, Chenyi Zi, H. Vicky Zhao, Yuan Yao, and Jia Li. Unigad: Unifying
563 multi-level graph anomaly detection. In *NeurIPS*, pp. 136120–136148, 2024.

564 Xiaoxiao Ma, Jia Wu, Jian Yang, and Quan Z. Sheng. Towards graph-level anomaly detection via
565 deep evolutionary mapping. In *KDD*, pp. 1631–1642, 2023.

566 Joshua Melton and Siddharth Krishnan. muxgnn: Multiplex graph neural network for heterogeneous
567 graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45:11067–11078, 2023.

568 Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion
569 Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. In *ICML
570 Workshop*, 2020.

571 Yuanhang Ren and Ye Du. Specializing word vectors by spectral decomposition on heterogeneously
572 twisted graphs. In *COLING*, pp. 3599–3609, 2020.

573 Walter Rudin. *Real and complex analysis*. McGraw-Hill, Inc., 1987.

574 Gen Shi, Yifan Zhu, Jian K. Liu, and Xuesong Li. Hegcl: Advance self-supervised learning in
575 heterogeneous graph-level representation. *IEEE Trans. Neural Networks Learn. Syst.*, 35:13914–
576 13925, 2024.

577 Yongduo Sui, Jie Sun, Shuyao Wang, Zemin Liu, Qing Cui, Longfei Li, and Xiang Wang. A unified
578 invariant learning framework for graph classification. In *KDD*, pp. 1301–1312, 2025.

579 Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
580 Bengio. Graph attention networks. In *ICLR*, 2018.

581 Xiyuan Wang and Muhan Zhang. How powerful are spectral graph neural networks. In *ICML*, pp.
582 23341–23362, 2022.

583 Zixiao Wang and Jicong Fan. Graph classification via reference distribution learning: Theory and
584 practice. In *NeurIPS*, pp. 137698–137740, 2024.

594 Lanning Wei, Zhiqiang He, Huan Zhao, and Quanming Yao. Search to capture long-range dependency
595 with stacking gnns for graph classification. In *WWW*, pp. 588–598, 2023.
596

597 Lirong Wu, Haitao Lin, Bozhen Hu, Cheng Tan, Zhangyang Gao, Zicheng Liu, and Stan Z. Li.
598 Beyond homophily and homogeneity assumption: Relation-based frequency adaptive graph neural
599 networks. *IEEE Trans. Neural Networks Learn. Syst.*, 35:8497–8509, 2024.

600 Yujia Wu, Bo Yang, Elynn Y. Chen, Yuzhou Chen, and Zheshi Zheng. Conditional prediction ROC
601 bands for graph classification. In *AISTATS*, pp. 2458–2466, 2025.
602

603 Huawei Xu, Chao Gao, Xianghua Li, and Zhen Wang. Rumormixer: Exploring echo chamber effect
604 and platform heterogeneity for rumor detection. In *PKDD*, pp. 21–37, 2024.

605 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
606 networks? In *ICLR*, 2019.
607

608 Xiaocheng Yang, Mingyu Yan, Shirui Pan, Xiaochun Ye, and Dongrui Fan. Simple and efficient
609 heterogeneous graph neural network. In *AAAI*, pp. 10816–10824, 2023.

610 Zhaoning Yu and Hongyang Gao. Molecular representation learning via heterogeneous motif graph
611 neural networks. In *ICML*, pp. 25581–25594, 2022.
612

613 Anru R Zhang and Yuchen Zhou. On the non-asymptotic and sharp lower tail bounds of random
614 variables. *Stat*, 9:e314, 2020.

615 Ge Zhang, Zhenyu Yang, Jia Wu, Jian Yang, Shan Xue, Hao Peng, Jianlin Su, Chuan Zhou, Quan Z.
616 Sheng, Leman Akoglu, and Charu C. Aggarwal. Dual-discriminative graph neural network for
617 imbalanced graph-level anomaly detection. In *NeurIPS*, pp. 24144–24157, 2022.
618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A PROOF**
 649

650 **Proof of Theorem 1.** Suppose we have N vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N$, where $\mathbf{v}_i \in \mathbb{R}^{d_i}$. We first prove
 651 the equivalence of different zero-padding methods, namely end padding, front padding, and scatter
 652 padding. Let $D = \max_i d_i$, then we have:

653

- 654 • End padding: $g(\mathbf{v}_i) = (\mathbf{v}_{i1}, \mathbf{v}_{i2}, \dots, \mathbf{v}_{id_i}, 0, 0, \dots, 0)$,
- 655 • Front padding: $g(\mathbf{v}_i) = (0, 0, \dots, 0, \mathbf{v}_{i1}, \mathbf{v}_{i2}, \dots, \mathbf{v}_{id_i})$,
- 656 • Scatter padding: Randomly generate a subset I_i of $\{1, 2, \dots, D\}$, where $|I_i| = d_i$, then

657
$$(g(\mathbf{v}_i))_j = \begin{cases} \mathbf{v}_k, & j = \text{the } k\text{-th element of } I_i, \\ 0, & j \notin I_i \end{cases}$$

660 All of these $g(\cdot)$ are linear isometries, i.e., $\|g(\mathbf{u}) - g(\mathbf{v})\| = \|\mathbf{u} - \mathbf{v}\|$, which means the zero padding
 661 ways are equivalent for preserving distance. \square

662 Denote $\mathbf{x}_i = g(\mathbf{v}_i) \in \mathbb{R}^D$ and draw $\mathbf{P} \in \mathbb{R}^{d \times D}$ with $\mathbf{P}_{ij} \sim \mathcal{N}(0, \frac{1}{d})$, then we have:

663
$$f : \mathbb{R}^D \rightarrow \mathbb{R}^d, f(\mathbf{x}) = \mathbf{P}\mathbf{x}$$

664 Consider the random variable

665
$$\mathbf{X} = \|f(\mathbf{u})\|_2^2 = \sum_{k=1}^d \langle \mathbf{P}_{k,*}, \mathbf{u} \rangle^2,$$

666 where $\mathbf{u} = \mathbf{x}_i - \mathbf{x}_j$, and $\mathbf{P}_{k,*}$ is the k -th row of \mathbf{P} , we have each inner product $\langle \mathbf{P}_{k,*}, \mathbf{u} \rangle$ is Gaussian
 667 with mean 0 and variance

668
$$\text{Var}(\langle \mathbf{P}_{k,*}, \mathbf{u} \rangle) = \sum_{l=1}^D \text{Var}(\mathbf{P}_{k,l}) \mathbf{u}_l^2 = \frac{1}{d} \sum_{l=1}^D \mathbf{u}_l^2 = \frac{\|\mathbf{u}\|_2^2}{d}$$

669 Hence we have:

670
$$\mathbf{Y}_k = \sqrt{\frac{d}{\|\mathbf{u}\|_2^2}} \langle \mathbf{P}_{k,*}, \mathbf{u} \rangle \sim \mathcal{N}(0, 1),$$

 671

$$\mathbf{X} = \sum_{k=1}^d \langle \mathbf{P}_{k,*}, \mathbf{u} \rangle^2 = \frac{\|\mathbf{u}\|_2^2}{d} \sum_{k=1}^d \mathbf{Y}_k^2$$

672 Therefore, $\mathbf{Z} = \sum_{k=1}^d \mathbf{Y}_k^2$ is χ^2 with d degrees of freedom, and $\mathbf{X} = \frac{\|\mathbf{u}\|_2^2}{d} \mathbf{Z}$.

673 A standard inequality (Zhang & Zhou, 2020) for tail bound of χ^2 random variable demonstrates:

674
$$\Pr[\|\mathbf{X} - \mathbf{u}\|_2^2 \geq \epsilon \mathbf{u}\|_2^2] \leq 2e^{-\frac{\epsilon^2 d}{4}}, \forall 0 < \epsilon < 1$$

675 Then for any \mathbf{u} , we can have:

676
$$\Pr[(1 - \epsilon)\mathbf{u}\|_2^2 \leq \|\mathbf{P}\mathbf{u}\|_2^2 \leq (1 + \epsilon)\mathbf{u}\|_2^2] \geq 1 - 2e^{-\frac{\epsilon^2 d}{4}}$$

677 We care about M pairs of vectors, then by the union bound, the probability that all M pairs of
 678 distances are preserved is at least $1 - 2Me^{-\frac{\epsilon^2 d}{4}}$. \square

679 **Proof of Theorem 2.** The optimization process of a spectral GNN in graph classification can be
 680 defined as:

681
$$\text{argmin}_{\theta_k, \mathbf{W}} \mathcal{L}(y, \text{Pooling}(\sum_{k=1}^K \theta_k g_k(\mathbf{L}) \mathbf{X} \mathbf{W}))$$

682 For simplicity, let \mathcal{L} be the MSE loss function, and the pooling function be the mean pooling function.
 683 Then we can reformulate the process for i -th sample with n_i nodes as:

684
$$\text{argmin}_{\theta_k, \mathbf{W}} \frac{1}{2} (\mathbf{p}^T \sum_{k=1}^K \theta_k g_k(\mathbf{L}) \mathbf{X} \mathbf{W} - y)^2,$$

702 where $\mathbf{p}^T \in \mathbb{R}^{1 \times n_i}$ represents the mean pooling function vector with each entry as $\frac{1}{n_i}$. Then, over a
 703 dataset of N graphs, we define the targeted process after reordering as:
 704

$$705 \quad \operatorname{argmin}_{\theta_k, \mathbf{W}} \frac{1}{2N} \sum_{i=1}^N \left(\sum_{k=1}^K \theta_k a_k^{(i)} - y^{(i)} \right)^2,$$

708 where $a_k^{(i)} = (\mathbf{p}^{(i)})^T g_k(\mathbf{L}^{(i)}) \mathbf{X}^{(i)} \mathbf{W}$.
 709

710 According to previous work (Wang & Zhang, 2022), the learned filter function is nearly identical
 711 across different bases since they share the same expressive power and can all reach the global
 712 minimum. Therefore, the optimization of \mathbf{W} is largely independent of the basis selection near the
 713 global minimum. In contrast, the optimization of θ_k is significantly influenced by the choice of basis.
 714 To emphasize the impact of basis selection, we will focus exclusively on the optimization of θ_k .
 715

716 To analyze the convergence speed near the global minimum, we then derive the Hessian matrix \mathbf{H} of
 717 the process with respect to θ_k :

$$717 \quad \mathbf{H}_{jk} = \frac{1}{N} \sum_{i=1}^N a_j^{(i)} a_k^{(i)}$$

719 Diagonalize each Laplacian matrix $\mathbf{L}^{(i)} = \mathbf{U}^{(i)} \mathbf{\Lambda}^{(i)} (\mathbf{U}^{(i)})^T$, we can get:
 720

$$721 \quad a_k^{(i)} = \sum_{l=1}^{n_i} g_k(\lambda_l^{(i)}) \phi_l^{(i)} \psi_l^{(i)},$$

724 where $\lambda_l^{(i)}$ is the l -th eigenvalue of $\mathbf{L}^{(i)}$, $\phi_l^{(i)} = [(\mathbf{U}^{(i)})^T \mathbf{X}^{(i)} \mathbf{W}]_l$, and $\psi_l^{(i)} = \mathbf{p}^{(i)} \mathbf{u}_l^{(i)}$.
 725

726 Since the Laplacian matrix is the normalized Laplacian matrix, the eigenvalue distribution of all
 727 graphs converges to a density $\rho(\lambda)$ on $[0, 2]$. Assume the random coefficients $\phi_l^{(i)}, \psi_l^{(i)}$ decorrelate
 728 between different l , and $\mathbb{E}[(\phi\psi)^2 | \lambda]$ depends only on λ . We have:

$$729 \quad \frac{1}{N} \sum_{i=1}^N a_j^{(i)} a_k^{(i)} \xrightarrow{N \rightarrow \infty} \int_0^2 g_j(\lambda) g_k(\lambda) \omega(\lambda) d\lambda,$$

732 where $\omega(\lambda) = \rho(\lambda) \mathbb{E}[(\phi\psi)^2 | \lambda]$. In other words, we have:
 733

$$734 \quad \mathbf{H} \propto [\langle g_j, g_k \rangle_\omega]_{j,k=0}^K,$$

735 where $\langle g_j, g_k \rangle_\omega = \int_0^2 g_j(\lambda) g_k(\lambda) \omega(\lambda) d\lambda$. Reaching the global minimum means \mathbf{H} is a diagonal
 736 matrix ($\langle g_j, g_k \rangle_\omega = 0$, where $j \neq k$), which is equivalent to that $g(\cdot)$ is an orthonormal basis
 737 in the polynomial space. Therefore, we choose a general form of orthogonal polynomials with
 738 flexible enough weight functions to adapt to different graph signal density functions, i.e., Jacobi
 739 Polynomials. \square
 740

741 **Proof of Theorem 3.** We first prove the injectivity of the combination of the Jacobi Polynomial-based
 742 graph filter to show that it can preserve the full information from the original graph.
 743

744 Suppose we have a combination of Jacobi Polynomial-based graph filter, that is, $\mathcal{T}(\mathbf{x}) = \sum_{v=1}^V P_n^{(\alpha_v, \beta_v)}(\mathbf{L}_v) \mathbf{x}$, where \mathbf{x} is a graph signal, $P_n^{(\alpha_v, \beta_v)}$ is the Jacobi Polynomial-based graph
 745 filter for the v -th view, and \mathbf{L}_v is the Laplacian matrix for v -th view. Define the kernel characterization
 746 for \mathcal{T} :

$$747 \quad \ker(\mathcal{T}) = \{ \mathbf{x} \in \mathbb{R}^N \mid \sum_{v=1}^V P_n^{(\alpha_v, \beta_v)}(\mathbf{L}_v) \mathbf{x} = 0 \}$$

750 To prove the injectivity is equal to prove $\ker(\mathcal{T}) = \{0\}$ under V heterogeneous views.
 751

752 Assume the different views of a heterogeneous graph satisfy the following condition:
 753

- 754 • Spectral Disjointness: $\forall i \neq j$, $\operatorname{Eigen}(\mathbf{L}_i) \cap \operatorname{Eigen}(\mathbf{L}_j) = \emptyset$, i.e., there are no shared eigenvalues
 755 between Laplacian matrices of different views.
- Full Spectral Coverage: $\cup_{v=1}^V \operatorname{Eigen}(\mathbf{L}_v) = \mathbb{R}_{\geq 0}$, i.e., eigenvalues of Laplacian matrix cover the
 756 entire spectrum.

756 • No Common Eigenvectors: $\exists \mathbf{x} \neq 0, \mathbf{L}_v \mathbf{x} = \lambda_v \mathbf{x}, \forall v$, i.e., different Laplacian matrices have
 757 distinct eigenspaces.
 758

759 By tuning α_v, β_v of each $P_n^{(\alpha_v, \beta_v)}$, we can easily obtain root avoidance of the Jacobi Polynomial-
 760 based graph filter:

$$761 P_n^{(\alpha_v, \beta_v)}(\lambda) \neq 0, \forall \lambda \in \text{Eigen}(\mathbf{L}_v),$$

762 as $P_n^{(\alpha_v, \beta_v)}(\mathbf{x})$ have n real roots in $[-1, 1]$, and the spectra \mathbf{L}_v can be rescale to $[-1, 1]$.
 763

764 By root avoidance and spectral disjointness, we have:

$$765 P_n^{(\alpha_v, \beta_v)}(\lambda) \neq 0 \Rightarrow |P_n^{(\alpha_v, \beta_v)}(\lambda)| > \delta_v > 0,$$

767 where δ_v is a small value larger than 0, which shows the strict positivity of $|P_n^{(\alpha_v, \beta_v)}(\lambda)|$.
 768

769 Then for $\mathbf{x} \neq 0$, we expand the eigenbasis of each \mathbf{L}_v :

$$770 \mathbf{x} = \sum_{i=1}^N c_{v,i} \mathbf{u}_{v,i},$$

773 where $c_{v,i} = \mathbf{u}_{v,i}^T \mathbf{x}$, $\mathbf{u}_{v,i}$ is the i -th eigenvector of \mathbf{L}_v , and for $a \neq b$, $\mathbf{u}_{a,c}^T \mathbf{u}_{b,d} = 0, \forall b, d$.
 774

775 Therefore, we can define the v -th filtered component as:

$$777 \mathbf{z}_v = P_n^{(\alpha_v, \beta_v)}(\mathbf{L}_v) \mathbf{x} = \sum_{i=1}^N P_n^{(\alpha_v, \beta_v)}(\lambda_{v,i}) c_{v,i} \mathbf{u}_{v,i}$$

780 Thus, we can compute:

$$782 \|\mathcal{T}(\mathbf{x})\|_2^2 = \left\langle \sum_{i=0}^V \mathbf{z}_i, \sum_{j=0}^V \mathbf{z}_j \right\rangle = \sum_{v=1}^V \|\mathbf{z}_v\|_2^2 + \sum_{i \neq j} \langle \mathbf{z}_i, \mathbf{z}_j \rangle,$$

785 where $\langle \mathbf{z}_i, \mathbf{z}_j \rangle = \mathbf{z}_i^T \mathbf{z}_j = 0, \forall i \neq j$, due to orthogonality.

786 By spectral coverage, $\exists v$ and i such that:

$$788 |c_{v,i}| > 0, \lambda_{v,i} \in \text{Eigen}(\mathbf{L}_v),$$

789 for which view v and i we have:

$$791 \|\mathbf{z}_v\|_2^2 \geq |P_n^{(\alpha_v, \beta_v)}(\lambda_{v,i})|^2 |c_{v,i}|^2 > \delta_v^2 |c_{v,i}|^2 > 0$$

793 Thus, we have:

$$794 \|\mathcal{T}(\mathbf{x})\|_2^2 > 0 \Rightarrow \mathcal{T}(\mathbf{x}) \neq 0,$$

795 which means strict positivity. Then we have $\ker(\mathcal{T}) = \{0\}$, i.e., injectivity, as desired. \square

796 **Proof of Theorem 4.** We define the fused feature extractor on the target component as:

$$798 \mathcal{F}(\mathbf{x}) = \sum_{v=1}^V \Pi_S P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v) \mathbf{x},$$

801 where Π_S is the projection onto the target feature subspace, and $P_n^{(\alpha_v^*, \beta_v^*)}$ is the optimized Jacobi
 802 Polynomial filter for view v . Then the enhancement factor is:

$$804 \gamma(\mathcal{F}) = \min_{\mathbf{z} \in \mathcal{S}, \|\mathbf{z}\|_2^2=1} \|\mathcal{F}(\mathbf{z})\|_2^2,$$

806 which represents the worst-case amplification of target features by the fused extractor.

807 Assume that:

809 • The target subspace \mathcal{S} is spanned by common eigenvectors of all graph Laplacians \mathbf{L}_v , which
 means for each $\mathbf{z} \in \mathcal{S}$, $\exists i$ such that $\mathbf{L}_v \mathbf{z} = \lambda_{v,i} \mathbf{z}, \forall v$, where $\lambda_{v,i}$ is the i -th eigenvalue of \mathbf{L}_v .

810 • For each view v , the Jacobi Polynomial filter is designed such that for all eigenvalues $\lambda_{v,i}$ associated
 811 with \mathcal{S} , we have $\forall \mathbf{z} \in \mathcal{S}, P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v)\mathbf{z} = c_v \mathbf{z}$, where $c_v = P_n^{(\alpha_v^*, \beta_v^*)}(\lambda_{v,i}) > 0$.
 812 • There exists a constant $c_{\min} > 0$ such that $c_v \geq c_{\min}, \forall v$.

814 Thus, for any $\mathbf{z} \in \mathcal{S}$ with $\|\mathbf{z}\|_2^2 = 1$, we have:

$$\Pi_{\mathcal{S}} P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v)\mathbf{z} = \Pi_{\mathcal{S}}(c_v \mathbf{z}) = c_v \mathbf{z},$$

817 since $\Pi_{\mathcal{S}} \mathbf{z} = \mathbf{z} (\mathbf{z} \in \mathcal{S})$.

818 Then we can have:

$$\mathcal{F}(\mathbf{z}) = \sum_{v=1}^V c_v \mathbf{z} = \left(\sum_{v=1}^V c_v \right) \mathbf{z},$$

822 whose norm is:

$$\|\mathcal{F}(\mathbf{z})\|_2^2 = \left\| \left(\sum_{v=1}^V c_v \right) \mathbf{z} \right\|_2^2 = \left| \left(\sum_{v=1}^V c_v \right) \right| \|\mathbf{z}\|_2^2 = \sum_{v=1}^V c_v$$

826 Give that $c_v \geq c_{\min} > 0, \forall v$, we have:

$$\gamma(\mathcal{F}) = \min_{\mathbf{z} \in \mathcal{S}, \|\mathbf{z}\|_2^2=1} \|\mathcal{F}(\mathbf{z})\|_2^2 \geq V c_{\min},$$

829 which means $\gamma(\mathcal{F}) = \Theta(V)$. □

830 Next, we show its robustness to noise. Consider a signal \mathbf{x} with target component $\mathbf{x}_{target} \in \mathcal{S}$ and
 831 noise $\mathbf{e} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ uncorrelated across views. Then we have:

$$\gamma(\mathcal{F}) = \mathcal{F}(\mathbf{x}_{target}) + \mathcal{F}(\mathbf{e})$$

834 From the above proof, we have the following for the target term:

$$\|\mathcal{F}(\mathbf{x}_{target})\|_2^2 \geq \gamma(\mathcal{F}) \|\mathbf{x}_{target}\|_2^2 \geq V c_{\min} \|\mathbf{x}_{target}\|_2^2$$

837 For the noise term, we have:

$$\begin{aligned} \mathbb{E}[\|\mathcal{F}(\mathbf{e})\|_2^2] &= \mathbb{E}[\left\| \sum_{v=1}^V \Pi_{\mathcal{S}} P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v) \mathbf{e} \right\|_2^2] \\ &= \sum_{v=1}^V \mathbb{E}[\|\Pi_{\mathcal{S}} P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v) \mathbf{e}\|_2^2] \\ &\leq \sum_{v=1}^V \sigma^2 \|\Pi_{\mathcal{S}} P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v)\|_F^2, \end{aligned}$$

847 since the noise is uncorrelated across views.

848 Assume $\|\Pi_{\mathcal{S}} P_n^{(\alpha_v^*, \beta_v^*)}(\mathbf{L}_v)\|_F^2 \leq M$ for some constant M , we have:

$$\mathbb{E}[\|\mathcal{F}(\mathbf{e})\|_2^2] \leq V \sigma^2 M$$

852 By Jensen's inequality, we have:

$$\mathbb{E}[\|\mathcal{F}(\mathbf{e})\|_2^2] \leq \sqrt{\mathbb{E}[\|\mathcal{F}(\mathbf{e})\|_2^2]} \leq \sigma \sqrt{V M}$$

855 Thus, the signal-to-noise ratio is:

$$\frac{\|\mathcal{F}(\mathbf{x}_{target})\|_2^2}{\mathbb{E}[\|\mathcal{F}(\mathbf{e})\|_2^2]} \geq \frac{V c_{\min} \|\mathbf{x}_{target}\|_2^2}{\sigma \sqrt{V M}} = \frac{\sqrt{V} c_{\min} \|\mathbf{x}_{target}\|_2^2}{\sigma \sqrt{M}} = \Theta(\sqrt{V})$$

859 □

860 **Proof of Theorem 5.** The κ -stereographic model provides a unified framework for Euclidean,
 861 Hyperbolic, and Spherical geometries through a common metric, parameterized by the curvature κ :

$$ds^2 = \frac{dr^2 + r^2 d\Omega^2}{(1 + \kappa r^2)^2},$$

864 where s is the square of an infinitesimally small distance between two points in a space, r is the radial
 865 coordinate in stereographic projection, and $d\Omega^2$ is the metric on the unit sphere.
 866

867 Jacobi Polynomials $P_n^{(\alpha, \beta)}(x)$ are orthogonal polynomials on $[-1, 1]$ with respect to the weight
 868 $(1-x)^\alpha(1+x)^\beta$. They arise as eigenfunctions of the Laplace-Beltrami operator in the κ -stereographic
 869 model for specific values of κ and after appropriate coordinate transformations. Below, we derive the
 870 connections for each geometry.

871 For $\kappa > 0$, the space is spherical with radius $R = \frac{1}{\sqrt{\kappa}}$. The Laplace-Beltrami operator Δ in
 872 stereographic coordinates has eigenfunctions that can be separated into radial and angular parts. The
 873 Laplace-Beltrami operator for a radial function $f(r)$ in n dimensions is:

$$874 \quad \Delta f = \kappa \frac{(1 + \kappa r^2)^n}{r^{n-1}} \frac{d}{dr} [r^{n-1} (1 + \kappa r^2)^{2-n} \frac{df}{dr}],$$

877 whose eigenvalue equation is $\Delta f + \lambda f = 0$.

878 To solve the equation, we introduce $u = \kappa r^2 \geq 0$:

$$879 \quad \frac{df}{dr} = 2\sqrt{\kappa u} \frac{df}{du}, \frac{d}{dr} = 2\sqrt{\kappa u} \frac{d}{du}$$

882 Then the equation can be represented as:

$$884 \quad \frac{(1+u)^n}{u^{\frac{n-1}{2}}} \frac{d}{du} [u^{\frac{n}{2}} (1+u)^{2-n} \frac{df}{du}] + \frac{\lambda}{4\kappa^2} f = 0$$

886 Change variable to $x = \frac{u-1}{u+1} = \frac{\kappa r^2 - 1}{\kappa r^2 + 1}$, so $x \in [-1, 1]$. Then $u = \frac{1+x}{1-x}$ and the derivatives become:

$$888 \quad \frac{d}{du} = \frac{dx}{du} \frac{d}{dx} = -\frac{2}{(1-x)^2} \frac{d}{dx}, \frac{d^2}{du^2} = \frac{4}{(1-x)^4} \frac{d^2}{dx^2} + \frac{8}{(1-x)^3} \frac{d}{dx}$$

891 Then the equation can be simplified to the Jacobi differential equation:

$$892 \quad (1-x^2) \frac{d^2 f}{dx^2} + [\beta - \alpha - (\alpha + \beta + 2)x] \frac{df}{dx} + k(k + \alpha + \beta + 1)f = 0,$$

895 where $\alpha = \beta = \frac{n-2}{2}$, and k is the quantum number related to the eigenvalue $\lambda = k(k + n - 1)\kappa$.

896 The solutions are Jacobi Polynomials:

$$898 \quad f(r) \propto P_k^{(\frac{n-2}{2}, \frac{n-2}{2})} \left(\frac{\kappa r^2 - 1}{\kappa r^2 + 1} \right)$$

900 \square

901 For $\kappa < 0$, the space is hyperbolic with curvature radius $R = \frac{1}{\sqrt{|\kappa|}}$. The spectrum of the Laplace-
 902 Beltrami operator is continuous, and the radial eigenfunctions are not polynomials but can be
 903 expressed as Jacobi functions (analytical continuations of Jacobi Polynomials).

905 Set $\kappa = -|\kappa|$, so $1 + \kappa r^2 = 1 - |\kappa|r^2$. The eigenvalue equation is similar to the spherical case:

$$907 \quad \Delta f + \lambda f = 0, \Delta f = \kappa \frac{(1 + \kappa r^2)^n}{r^{n-1}} \frac{d}{dr} [r^{n-1} (1 + \kappa r^2)^{2-n} \frac{df}{dr}]$$

909 Then we use $u = |\kappa|r^2$ to get:

$$911 \quad \frac{(1+u)^n}{u^{\frac{n-1}{2}}} \frac{d}{du} [u^{\frac{n}{2}} (1+u)^{2-n} \frac{df}{du}] + \frac{\lambda}{4\kappa^2} f = 0$$

914 Change variable to $x = \frac{u-1}{u+1} = \frac{|\kappa|r^2 - 1}{|\kappa|r^2 + 1}$, so $x \in [-\infty, 0]$, then the equation becomes a confluent
 915 hypergeometric equation, which is solved by Jacobi functions with parameters $\alpha = \beta = \frac{n-2}{2}$:

$$917 \quad f(r) \propto P_{-\frac{1}{2}+i\sigma}^{(\frac{n-2}{2}, \frac{n-2}{2})} \left(\frac{|\kappa|r^2 - 1}{|\kappa|r^2 + 1} \right),$$

918 where $\sigma = \sqrt{\frac{\lambda}{|\kappa|} - \frac{(n-1)^2}{4}}$, $\forall \lambda > \frac{(n-1)^2|\kappa|}{4}$, and i is the imaginary unit. □

920 For $\kappa = 0$, the space is Euclidean space. The Laplace-Beltrami operator reduces to the standard
921 Laplacian, and the radial eigenfunctions are Bessel functions, which arise as a limit of Jacobi
922 Polynomials as $\kappa \rightarrow 0$.

923 When $\kappa \rightarrow 0$, the metric is $ds^2 = dr^2 + r^2 d\Omega^2$, and the radial eigenvalue equation is:
924

$$925 \frac{1}{r^{n-1}} \frac{d}{dr} (r^{n-1} \frac{df}{dr}) + \lambda f = 0$$

927 This is the spherical Bessel equation, whose solutions are Bessel functions:
928

$$929 f(r) \propto r^{-\frac{n-2}{2}} J,$$

930 where J stands for $J_n(x) = \sum_{m=0}^{\infty} \frac{(-1)^m}{m! \Gamma(m+n+1)} \left(\frac{x}{2}\right)^{2m+n}$, and $\Gamma(\cdot)$ is the gamma function.
931

932 For fixed r and k , as $\kappa \rightarrow 0$, the Jacobi Polynomial limit is:
933

$$934 \lim_{\kappa \rightarrow 0} P_k^{\left(\frac{n-2}{2}, \frac{n-2}{2}\right)} \left(\frac{\kappa r^2 - 1}{\kappa r^2 + 1}\right) \propto r^{-\frac{n-2}{2}} J_{\frac{n-2}{2}}(\sqrt{k(k+n-1)r}),$$

□

936 **Proof of Theorem 6.** First, we introduce the following theorem:
937

938 **Theorem 10** (Weierstrass approximation theorem). *Suppose $f(x)$ is a continuous real-valued function
939 defined on the real interval $[a, b]$. For every $\epsilon > 0$, there exist a polynomial p such that for all x in
940 $[a, b]$, we have $\|f(x) - p(x)\|_{\infty} < \epsilon$.*

941 Then, given a $f(x) \in C([-1, 1])$ and $\epsilon > 0$, by theorem 10, we can pick a genuine polynomial
942 $p(x) = \sum_{i=1}^N w_i x^i$ so that $\|f(x) - p(x)\|_{\infty} < \epsilon$.
943

944 And by definition, each Jacobi Polynomial $P_n^{(\alpha, \beta)}(x)$ is a genuine polynomial of degree n , so:
945

$$946 \text{Span}(P_0^{(\alpha, \beta)}, P_1^{(\alpha, \beta)}, \dots) = \{\text{all real polynomials in } x\},$$

947 which means every polynomial $p(x)$ of degree $\leq N$ can be written uniquely in the Jacobi basis:
948

$$949 p(x) = \sum_{i=1}^N c_i P_i^{(\alpha, \beta)}(x)$$

950 Hence the finite Jacobi Polynomial sum $S_N(x) = \sum_{i=1}^N c_i P_i^{(\alpha, \beta)}(x)$ satisfies:
951

$$952 \|f(x) - S_N(x)\|_{\infty} = \|f(x) - p(x)\|_{\infty} < \epsilon$$

953 Thus, any continuous f on $[-1, 1]$ can be uniformly approximated by finite linear combinations of
954 Jacobi polynomials. □

955 Then we consider the weighted space $L_w^2[-1, 1]$, where the weight function is the same as Jacobi
956 Polynomials:
957

$$958 w(x) = (1-x)^{\alpha} (1+x)^{\beta},$$

959 with $\alpha, \beta > -1$.
960

961 Let μ be the measure defined by $d\mu = w(x)dx$. Since $\alpha, \beta > -1$, the integral satisfies:
962

$$963 \int_{-1}^1 w(x)dx < \infty,$$

964 which means μ is a finite Borel measure on $[-1, 1]$.
965

966 Denote any function $q \in L_w^2[-1, 1]$, we then prove that Jacobi Polynomials can approximate it in the
967 L_w^2 norm.
968

969 Since μ is a finite Borel measure on the compact interval $[-1, 1]$, the continuous functions on $[-1, 1]$
970 area dense in $L_w^2[-1, 1]$, which is proved in (\cdot) . Thus, for any $\delta > 0$, there exists a continuous function
971 t on $[-1, 1]$ such that:
972

$$973 \|q(x) - t(x)\|_{L_w^2} < \frac{\delta}{2}$$

972 Then, following the previous proof, we have:
 973

$$974 \quad ||t(x) - p(x)|| = \sup_{x \in [-1, 1]} |t(x) - p(x)| < \epsilon, \\ 975$$

976 where $p(x)$ is finite linear combinations of Jacobi polynomials.
 977

978 For any $\epsilon > 0$, we can choose ϵ , such that:
 979

$$980 \quad \epsilon^2 \int_{-1}^1 w(x) dx < \left(\frac{\delta}{2}\right)^2 \\ 981$$

982 Thus, we have:
 983

$$984 \quad ||t(x) - p(x)||_{L_w^2}^2 = \int_{-1}^1 |t(x) - p(x)|^2 w(x) dx \leq \epsilon^2 \int_{-1}^1 w(x) dx < \left(\frac{\delta}{2}\right)^2, \\ 985$$

986 which means $||t(x) - p(x)||_{L_w^2} < \frac{\delta}{2}$.
 987

988 By triangle inequality, we can derive:
 989

$$990 \quad ||q(x) - p(x)||_{L_w^2} \leq ||q(x) - t(x)||_{L_w^2} + ||t(x) - p(x)||_{L_w^2} < \frac{\delta}{2} + \frac{\delta}{2} = \delta, \\ 991$$

992 which concludes that Jacobi Polynomials can approximate any function $q \in L_w^2[-1, 1]$. \square
 993

994 **Proof of Theorem 7.** Let $E_N(f(x)) = \min_{\deg(g(x)) \leq N} ||f(x) - g(x)||_\infty$, where $f(x) \in$
 995 $C^{r-1}([-1, 1])$, $f^{(r)}(x) \in L(w^{(\alpha, \beta)}(x))$, r is the derivative order, and $w^{(\alpha, \beta)}(x) = (1-x)^\alpha(1+x)^\beta$.

996 We then construct a positive kernel:
 997

$$998 \quad K_N^{(r)}(x, t) = \sum_{k=0}^N a_{N,k} P_k^{(\alpha, \beta)}(x) P_k^{(\alpha, \beta)}(t), \\ 999$$

1000 where $a_{N,k} \geq 0$ chosen so that for each fixed $x \in [-1, 1]$:

- 1001 • Normalization: $\int_{-1}^1 K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt = 1$,
- 1002 • Moment vanishing up to order $r-1$: $\int_{-1}^1 (t-x)^m K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt = 0, m = 1, 2, \dots, r-1$,
- 1003 • High-order moment bound: $\int_{-1}^1 |t-x|^r K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt \leq \frac{C}{N^r}$, where C depends on r, α, β

1004 The existence of such a kernel is standard, i.e., the classical Jackson kernel in orthogonal-polynomial
 1005 theory (Rudin, 1987).
 1006

1007 Then define the Jackson operator J_N by:
 1008

$$1009 \quad (J_N f)(x) = \int_{-1}^1 f(t) K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt, \\ 1010$$

1011 where $J_N f(x)$ is a polynomial of degree $\leq N$, because $K_N^{(r)}(x, \cdot)$ is a sum of Jacobi Polynomials up
 1012 to degree N . Consequently:
 1013

$$1014 \quad E_N(f(x)) \leq ||f(x) - (J_N f)(x)||_\infty$$

1015 Next, we use repeated integration of the Taylor formula. For each $t \in [-1, 1]$, there holds:
 1016

$$1017 \quad f(t) = f(x) + f'(x)(t-x) + \dots + \frac{f^{(r-1)}(x)}{(r-1)!} (t-x)^{r-1} + R_r(x, t), \\ 1018$$

1019 where the remainder can be written in integral form:
 1020

$$1021 \quad R_r(x, t) = \frac{1}{(r-1)!} \int_0^1 (1-u)^{r-1} f^{(r)}(x+u(t-x))(t-x)^r du \\ 1022$$

1026 By the moment-vanishing property of $K_N^{(r)}(x, t)$, when we subtract out the Taylor part, all terms up
 1027 to $(t - x)^{r-1}$ integrate to 0. Thus:
 1028

$$\begin{aligned} 1029 \quad f(x) - (J_N f)(x) &= \int_{-1}^1 [f(t) - \text{Taylor at } x] K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt \\ 1030 \\ 1031 \\ 1032 \\ 1033 \end{aligned}$$

$$= \int_{-1}^1 R_r(x, t) K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt$$

1034 Insert the integral form of $R_r(x, t)$:

$$1035 \quad f(x) - (J_N f)(x) = \frac{1}{(r-1)!} \int_{-1}^1 \int_0^1 (1-u)^{r-1} f^{(r)}(x+u(t-x))(t-x)^r K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) du dt$$

$$1036$$

$$1037$$

1038 Taking absolute values and using Fubini's theorem:

$$1039 \quad |f(x) - (J_N f)(x)| \leq \frac{1}{(r-1)!} \int_{-1}^1 (1-u)^{r-1} \int_0^1 |f^{(r)}(x+u(t-x))| |t-x|^r K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt du$$

$$1040$$

$$1041$$

$$1042$$

1043 Then, change variables in the inner integral, i.e., for each fixed u , the map $t \rightarrow s = x + u(t - x)$ is
 1044 linear of Jacobian $dt = \frac{ds}{u}$. Moreover $w^{(\alpha, \beta)}(t) dt \leq C w^{(\alpha, \beta)}(s) ds$, since w is smooth and $u \in [0, 1]$.
 1045 One shows:

$$1046 \quad |t-x|^r K_N^{(r)}(x, t) w^{(\alpha, \beta)}(t) dt \leq \frac{C}{N^r} w^{(\alpha, \beta)}(s) ds,$$

$$1047$$

$$1048$$

by the high-order moment bound. Hence, we have:

$$1049 \quad |f(x) - (J_N f)(x)| \leq \frac{C}{N^r} \int_0^1 (1-u)^{r-1} \frac{du}{u} \int_{-1}^1 |f^{(r)}(s)| w^{(\alpha, \beta)}(s) ds,$$

$$1050$$

$$1051$$

1052 where $\int_0^1 (1-u)^{r-1} \frac{du}{u}$ converges to a constant depending only on r . We conclude:

$$1053 \quad \|f(x) - (J_N f)(x)\|_\infty \leq |f(x) - (J_N f)(x)| \leq \frac{C}{N^r} \int_{-1}^1 |f^{(r)}(s)| w^{(\alpha, \beta)}(s) ds,$$

$$1054$$

$$1055$$

1056 which completes the proof. \square

$$1057 \quad E_N(f(x)) \leq \|f(x) - (J_N f)(x)\|_\infty \leq \frac{C}{N^r} \int_{-1}^1 |f^{(r)}(t)| w^{(\alpha, \beta)}(t) dt = \frac{C}{N} \|f^{(r)}(x)\|_{L_w^1},$$

$$1058$$

$$1059$$

1060 where $J_N f$ is a polynomial of degree $\leq N$ and C depends only on r, α, β . \square

1061 **Proof of Theorem 8.** For simplicity, we assume the total class number C is 2.

1062 We first calculate the gradient flow of \mathcal{L}_{RFACE} using the chain rule:

$$1063 \quad \nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{RFACE} = (\tilde{\mathbf{p}}_{i,c} - \mathbf{y}_{i,c}) [1 - \gamma \kappa_c \mathbf{p}_{i,c} (1 - \mathbf{p}_{i,c})]$$

$$1064$$

1065 Since $\kappa_c < 0$:

$$1066 \quad \gamma |\kappa_c| \mathbf{p}_{i,c} (1 - \mathbf{p}_{i,c}) > 0 \Rightarrow [1 + \gamma |\kappa_c| \mathbf{p}_{i,c} (1 - \mathbf{p}_{i,c})] > 1$$

$$1067$$

$$1068$$

1069 We then calculate the gradient flow of \mathcal{L}_{CE} using the chain rule:

$$1070 \quad \nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{CE} = \mathbf{p}_{i,c} - 1$$

$$1071$$

1072 We focus on the case when $\mathbf{y}_{i,c} = 1, \mathbf{p}_{i,c} \approx 0$ (probability of anomalous class is low):

$$1073 \quad |\tilde{\mathbf{p}}_{i,c} - \mathbf{y}_{i,c}| \approx |\mathbf{p}_{i,c} - 1|$$

$$1074$$

1075 Then we can have:

$$1076 \quad |\nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{RFACE}| = [1 + \gamma |\kappa_c| \mathbf{p}_{i,c} (1 - \mathbf{p}_{i,c})] |\mathbf{p}_{i,c} - 1| > |\nabla_{\mathbf{z}_{i,c}} \mathcal{L}_{CE}| = |\mathbf{p}_{i,c} - 1|,$$

$$1077$$

$$1078$$

1079 with an amplifying factor $[1 + \gamma |\kappa_c| \mathbf{p}_{i,c} (1 - \mathbf{p}_{i,c})]$ proportional to $\gamma |\kappa_c|$. \square

1079 Next, we prove the adjustment $\Delta \mathbf{p}_{i,c}$ is $(1 + \gamma |\kappa_c|)$ -Lipschitz continuous, so the topology of the latent graph embedding space is preserved.

1080 Define $g(z_{i,c}) = \Delta p_{i,c} = -\gamma \kappa_c (\mathbf{p}_{i,c} - \mathbf{y}_{i,c})$, then for two different graphs embedding $\mathbf{z}_1, \mathbf{z}_2$, we
 1081 have:

$$|g(\mathbf{z}_1) - g(\mathbf{z}_2)| = |-\gamma \kappa_c (\mathbf{p}_1 - \mathbf{p}_2)| = \gamma |\kappa_c| |\mathbf{p}_1 - \mathbf{p}_2|$$

1082 Since Sigmoid is 1-Lipschitz, we have:

$$|g(\mathbf{z}_1) - g(\mathbf{z}_2)| \leq \gamma |\kappa_c| |\mathbf{z}_1 - \mathbf{z}_2|,$$

1083 which means $g(\cdot)$ is $\gamma |\kappa_c|$ -Lipschitz.

1084 Then the adjusted logit $\tilde{\mathbf{z}}_1 = \mathbf{z}_1 + g(\mathbf{z}_1), \tilde{\mathbf{z}}_2 = \mathbf{z}_2 + g(\mathbf{z}_2)$ satisfies:

$$|\tilde{\mathbf{z}}_1 - \tilde{\mathbf{z}}_2| \leq |\mathbf{z}_1 - \mathbf{z}_2| + |g(\mathbf{z}_1) - g(\mathbf{z}_2)| \leq (1 + \gamma |\kappa_c|) |\mathbf{z}_1 - \mathbf{z}_2|,$$

1085 which means the amplification follows $(1 + \gamma |\kappa_c|)$ -Lipschitz. \square

1086 **Proof of Theorem 9.** When $\tilde{\mathbf{p}}_{i,c} = \mathbf{y}_{i,c}$, there are two cases, i.e., $\mathbf{y}_{i,c} = 0$ and $\mathbf{y}_{i,c} = 1$.

1087 For $\mathbf{y}_{i,c} = 0$, we have:

$$\tilde{\mathbf{p}}_{i,c} = 0 \Rightarrow \text{Sigmoid}(\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c}) = 0 \Rightarrow \mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} \rightarrow -\infty$$

1088 In this situation, since $\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} = \mathbf{z}_{i,c} - \gamma \kappa_c \mathbf{p}_{i,c}$, and $-\gamma \kappa_c \mathbf{p}_{i,c}$ is bounded, $\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} \rightarrow -\infty$
 1089 requires $\mathbf{z}_{i,c} \rightarrow -\infty$, which implies $\mathbf{p}_{i,c} = \text{Sigmoid}(\mathbf{z}_{i,c}) \rightarrow 0 = \mathbf{y}_{i,c}$.

1090 For $\mathbf{y}_{i,c} = 1$, we have:

$$\tilde{\mathbf{p}}_{i,c} = 1 \Rightarrow \text{Sigmoid}(\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c}) = 1 \Rightarrow \mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} \rightarrow \infty$$

1091 In this situation, since $\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} = \mathbf{z}_{i,c} - \gamma \kappa_c (\mathbf{p}_{i,c} - 1)$, and $-\gamma \kappa_c (\mathbf{p}_{i,c} - 1)$ is bounded,
 1092 $\mathbf{z}_{i,c} + \Delta \mathbf{p}_{i,c} \rightarrow \infty$ requires $\mathbf{z}_{i,c} \rightarrow \infty$, which implies $\mathbf{p}_{i,c} = \text{Sigmoid}(\mathbf{z}_{i,c}) \rightarrow 1 = \mathbf{y}_{i,c}$. \square

B DATASETS AND BASELINES

1108 **Datasets.** MCF-7, MOLT-4, PC-3, SW-620, NCI-H23, OVCAR-8, P388, SF-295, SN12C, and
 1109 UACC257 are 10 small-molecule biological activity datasets from TUDataset (Morris et al., 2020),
 1110 each corresponding to a different cancer cell line screen. Compounds are represented as heterogeneous
 1111 graphs where nodes are atom types and edges are the bonds between them. **Remarkably, we utilize**
 1112 **the original data in the TUDataset datasets, where the number of node and edge types are large, as**
 1113 **the real chemical compounds are extremely complex. The large number of node and edge types in**
 1114 **the public datasets posts additional challenges for the heterogeneous graph-level anomaly detection.**
 1115 Each compound is labeled as active or inactive against its respective cancer type; we treat inactive
 1116 compounds as normal and active ones as anomalies. Node features are one-hot encodings of the atom
 1117 labels.

1118 The above public datasets are originally graph classification datasets, whereas the datasets below are
 1119 node classification datasets. Therefore, we need to transform them into graph classification datasets.
 1120 The transformation is the same:

- 1121 • We follow the original anomalous ratio to sample n_n normal nodes and n_a anomalous nodes from
 1122 a heterogeneous graph, where $\frac{n_a}{n_n + n_a}$ is the anomalous ratio of the original dataset, to simulate the
 1123 imbalanced nature of graph-level anomaly detection tasks.
- 1124 • Use original Breadth-First Search algorithm to obtain the subgraph around n_n normal nodes and
 1125 n_a anomalous nodes. For small-scale graphs, such as DBLP and IMDB, we set small n_n and
 1126 n_a and 3 as the sampling layer number to limit the overlap between graphs, while keep enough
 1127 information in each graph. For large-scale graphs, such as PDNS and RCDD, we set large n_n
 1128 and n_a and 2 as the sampling layer number to provide diverse enough samples while reduce the
 1129 running cost for several baselines for fair comparison, because some of them might cost high
 1130 computational resources, as reported in Appendix C.
- 1131 • Use the center node label as the subgraph label.

1132 DBLP and IMDB are public datasets processed by pytorch_geometric (Fey & Lenssen, 2019). The
 1133 DBLP dataset is a subset of the computer science bibliography, comprising four node types: authors,

1134
 1135 Table 5: Datasets used in the experiments, where n_n, n_a are the normal and anomalous number of
 1136 graphs respectively, $r = \frac{n_a}{n_n+n_a}$ is the anomalous rate of the dataset, n, m are the average number
 1137 of nodes and edges in graphs respectively, T_V, R_E are the types of nodes and edges in graphs
 respectively, and d is the dimension of nodes after projection.

| Source | Type | Dataset | n_n | n_a | r | n | m | T_V | R_E | d |
|----------|---------|-------------|-------|-------|--------|--------|-------|-------|-------|-----|
| Public | Bioinfo | MCF-7 | 25476 | 2294 | 0.0826 | 26.40 | 28.53 | 46 | 129 | 46 |
| | | MOLT-4 | 36625 | 3140 | 0.0790 | 26.10 | 28.14 | 64 | 176 | 64 |
| | | PC-3 | 25941 | 1568 | 0.0570 | 26.36 | 28.49 | 45 | 133 | 45 |
| | | SW-620 | 38122 | 2410 | 0.0595 | 26.06 | 28.09 | 65 | 184 | 65 |
| | | NCI-H23 | 38296 | 2057 | 0.0510 | 26.07 | 28.10 | 65 | 182 | 65 |
| | | OVCAR-8 | 38437 | 2079 | 0.0513 | 26.08 | 28.11 | 65 | 184 | 65 |
| | | P388 | 39174 | 2298 | 0.0554 | 22.11 | 23.56 | 72 | 271 | 72 |
| | | SF-295 | 38246 | 2025 | 0.0503 | 26.06 | 28.09 | 65 | 184 | 65 |
| | | SN12C | 38049 | 1955 | 0.0489 | 26.08 | 28.11 | 65 | 184 | 65 |
| | | UACC257 | 38345 | 1643 | 0.0411 | 26.09 | 28.13 | 64 | 176 | 64 |
| Citation | | DBLP | 1197 | 745 | 0.3836 | 162.08 | 96.29 | 4 | 6 | 50 |
| | | IMDB | 1584 | 1135 | 0.4174 | 85.35 | 59.52 | 3 | 4 | 64 |
| | | PDNS | 41337 | 8663 | 0.1733 | 48.80 | 72.13 | 2 | 3 | 32 |
| | | RCDD | 50000 | 8364 | 0.1433 | 17.43 | 9.20 | 7 | 8 | 256 |
| Private | Finance | Transaction | 20000 | 437 | 0.0214 | 14.87 | 20.12 | 6 | 11 | 325 |

1154
 1155 papers, terms, and conferences. Authors are categorized by research area (database, data mining,
 1156 artificial intelligence, information retrieval) and are represented by a bag-of-words feature vector
 1157 derived from their paper keywords. For our task, authors from the database area are designated as
 1158 normal nodes, while those from data mining are treated as anomalous. The IMDB dataset is a subset
 1159 of the Internet Movie Database, containing movies, actors, and directors as node types. Movies
 1160 are classified by genre (action, comedy, drama) and are represented by bag-of-words features from
 1161 their plot keywords. In this context, action movies are the normal class, and comedy movies are the
 1162 anomalous class.

1163 PDNS and RCDD are public datasets collected from Kaggle ¹. The PDNS dataset is a cybersecurity
 1164 graph constructed from a seed set of malicious domains. Its infrastructure data is extracted from a
 1165 global passive DNS repository. The graph contains two entity types (domains and IPs) connected
 1166 by four relations (e.g., "domain resolves to IP"). Each domain node has a 10-dimensional feature
 1167 vector derived from its domain name and a binary label identifying it as malicious. We directly use
 1168 these original labels to define normal and anomalous nodes. The RCDD is a large-scale e-commerce
 1169 network from Alibaba, built for real-world risk detection. It contains 7 node types (e.g., buyer, seller)
 1170 and 7 edge types (e.g., buy, sell), though specific names are anonymized for confidentiality. In this
 1171 network, risk nodes often disguise themselves by forging relationships. Each node is described by
 1172 a 256-dimensional feature vector, and item nodes are labeled as either risk commodities or normal.
 1173 These original labels are used to designate the normal and anomalous classes.

1174 The final dataset is a proprietary financial heterogeneous graph provided by a prominent company. Its
 1175 objective is to identify sub-networks, or communities, associated with suspicious or non-compliant
 1176 activity. The graph schema is complex, comprising 6 node types (e.g., representing real users and
 1177 entities) and 11 edge types that define the intricate relationships between them. The task is naturally
 1178 a GAD problem: each entire graph is labeled as either containing a risky community or being normal.
 1179 We directly adopt these original labels to train our model to distinguish between anomalous and
 1180 normal graphs.

1181 **Baselines.** The first group is homogeneous graph classification models:

1182 • **GCN** (Kipf & Welling, 2017): A foundational graph convolutional network that performs neighbor-
 1183 hood aggregation through a spectral graph convolution-inspired operation.

1184 • **SAGE** (Hamilton et al., 2017): A scalable inductive framework that generates node embeddings
 1185 by sampling and aggregating features from a node’s local neighborhood.

1¹<https://www.kaggle.com/>

- **GAT** (Velickovic et al., 2018): Employs an attention mechanism to compute hidden representations by assigning different weights to each neighbor node.
- **GIN** (Xu et al., 2019): A theoretically powerful model designed to be as expressive as the Weisfeiler-Lehman graph isomorphism test.
- **LRGNN** (Wei et al., 2023): Addresses the limitation of shallow receptive fields by stacking multiple GNNs to capture long-range dependencies between distant nodes.
- **GRDL** (Wang & Fan, 2024): Treats node embeddings as discrete distributions within a latent space, enabling graph-level classification without a global readout function.
- **UQGNN** (Wu et al., 2025): A model that integrates uncertainty quantification into the graph representation learning process, producing confidence estimates alongside predictions.
- **UIL** (Sui et al., 2025): Provides a unified framework for invariant graph learning by enforcing both structural and semantic invariance, leading to the identification of more robust and stable node representations.

The second group is heterogeneous graph classification models:

- **HMGNN** (Yu & Gao, 2022): Models complex heterogeneous structures by constructing heterogeneous motif graphs to capture rich semantic information from multiple node and edge types.
- **muxGNN** (Melton & Krishnan, 2023): Represents graphs as multiplex networks, using separate graphs for each relation type and a coupling graph to connect node representations across these relations.
- **HeGCL** (Shi et al., 2024): A contrastive learning framework that learns node and graph embeddings by contrasting a meta-path view with a global network topology view.
- **RFAGNN** (Wu et al., 2024): Handles both heterophily and heterogeneity within a unified model using a relation-based frequency adaptive graph filter.
- **SHGLNN** (Hayat et al., 2024): Leverages hypergraphs constructed from heterogeneous graphs to model complex higher-order (intra- and inter-graph) contextual relationships.

The third group is graph-level anomaly detection models:

- **iGAD** (Zhang et al., 2022): Anomaly detection is performed by comparing input graphs against a set of prototypical neural substructure patterns.
- **GmapAD** (Ma et al., 2023): Maps entire graphs into a well-structured latent space where normal and anomalous graphs are more easily separable.
- **RumorMixer** (Xu et al., 2024): A specialized model for rumor detection that captures the echo chamber effect and platform heterogeneity inherent in social networks.
- **RQGNN** (Dong et al., 2024): Leverages the Rayleigh Quotient to combine spectral and spatial information for anomaly detection.
- **UniGAD** (Lin et al., 2024): A unified framework that integrates node-level, subgraph-level, and graph-level information for comprehensive graph anomaly detection.

C ALGORITHM AND COMPLEXITY

We first analyze the Preprocess function. As shown in Algorithm 1, in lines 1-6, we have in total of $O(N)$, where N is the number of graphs in the dataset, as we need to find the d_{\max} of all the graphs. Then, in lines 7-11, we need to do the projection for each graph in the dataset. Each will cost $O(Vdd_{\max})$, where V is the number of nodes in graph G . Hence, the total cost will be $O(Nndd_{\max})$, where n is the average number of nodes in each graph of \mathcal{G} . Therefore, the total time complexity of Preprocess is $O(Nndd_{\max})$.

Then, we analyze the time complexity of JPGNN for each graph G . As presented in Algorithm 2, in lines 1-8, the dominant cost is the summation of weighted adjacency matrices. In practice, we don't need the summation, as we can multiply the coefficients by the edge weights. Thus, the total cost is $O(E)$, where E is the number of edges in the graph. Then, for lines 9-16, the dominant cost should be line 12, which has a cost of $O(KTVEd_{hid}^3)$, where d_{hid} is the hidden dimension of the layer of the GNN. Therefore, the total time complexity of JPGNN is $O(KTVEd_{hid}^3)$. Next, we analyze the time complexity of RFACE in Algorithm 3. In lines 1-6, we only need to use basic operations with $O(1)$ time complexity. Therefore, the total time complexity of RFACE is $O(1)$.

1242

Algorithm 1: Preprocess

```

1243 Input:  $\mathcal{G}, d$ 
1244 Output:  $\mathcal{G}'$ 
1245 1  $\mathcal{G}' \leftarrow \mathcal{G};$ 
1246 2  $d_{\max} \leftarrow 0;$ 
1247 3 for  $G$  in  $\mathcal{G}$  do
1248 4   for  $\mathbf{X}_t$  in  $G.\mathcal{X}$  do
1249 5      $d_{\max} \leftarrow \max(d_{\max}, d_t);$ 
1250
1251 6  $\mathbf{P} \leftarrow \mathbf{P}_{i,j} \sim \mathcal{N}(0, \frac{1}{d}), \mathbf{P} \in \mathbb{R}^{d \times d_{\max}};$ 
1252 7 for  $G$  in  $\mathcal{G}'$  do
1253 8    $G.\mathbf{X}^{proj} \leftarrow \text{Null};$ 
1254 9   for  $\mathbf{X}_t$  in  $G.\mathcal{X}$  do
1255 10      $\mathbf{X}_t^{proj} \leftarrow \mathbf{X}_t \oplus \mathbf{0}, \mathbf{0} \in \mathbb{R}^{|V_t| \times (d_{\max} - d_t)};$ 
1256 11      $G.\mathbf{X}^{proj} \leftarrow G.\mathbf{X}^{proj} \oplus \mathbf{P}\mathbf{X}^{proj};$ 
1257
1258 12 Return  $\mathcal{G}'$ ;
1259
1260

```

1260

Algorithm 2: JPGNN

```

1261 Input:  $\mathcal{A}, \mathbf{X}^{proj}, K, T$ 
1262 Output:  $z$ 
1263 1  $\mathbf{A} \leftarrow 0;$ 
1264 2 for  $\mathbf{A}_r$  in  $\mathcal{A}$  do
1265 3    $\mathbf{A} \leftarrow \mathbf{A} + \omega_r \mathbf{A}_r;$ 
1266
1267 4  $\mathbf{L} \leftarrow \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}};$ 
1268 5  $\lambda_{\max} \leftarrow \max(\text{Eigen}(\mathbf{L}));$ 
1269 6  $\hat{\mathbf{L}} \leftarrow \frac{2}{\lambda_{\max}} \mathbf{L} - \mathbf{I};$ 
1270 7  $\mathbf{H}^{(0)} \leftarrow \sigma(\mathbf{X}^{proj} \mathbf{W}^{(0)});$ 
1271 8  $\mathbf{H}^{stack} \leftarrow \mathbf{H}^{(0)};$ 
1272 9 for  $k$  in  $\{1, \dots, K\}$  do
1273 10    $\mathbf{H}^{(k)} \leftarrow 0;$ 
1274 11   for  $t$  in  $\{0, \dots, T\}$  do
1275 12      $\mathbf{H}^{(k)} \leftarrow \mathbf{H}^{(k)} + \theta_t^{(k)} P_t^{(\alpha^{(k)}, \beta^{(k)})}(\hat{\mathbf{L}}) \mathbf{H}^{(k-1)} \mathbf{W}^{(k)};$ 
1276 13    $\mathbf{H}^{(k)} \leftarrow \sigma(\mathbf{H}^{(k)});$ 
1277 14    $\mathbf{H}^{stack} \leftarrow \mathbf{H}^{stack} \oplus \mathbf{H}^{(k)};$ 
1278
1279 15  $\mathbf{H} \leftarrow \sigma(\mathbf{H}^{stack} \mathbf{W});$ 
1280 16  $z \leftarrow \text{Pooling}(\mathbf{H});$ 
1281 17 Return  $z$ ;
1282
1283
1284

```

1285

Finally, in Algorithm 4, to clearly show the time complexity of each epoch of the training procedure, we combine the above time complexities. As shown in lines 7-10, we need to call JPGNN and RFACE $|\mathcal{G}|$ times, so the time complexity of each epoch of the training procedure of JacobiGAD is $O(NKTnmd_{hid}^3)$, where m is the average number of edges in each graph of \mathcal{G} .

1286

Compared to homogeneous graph-level classification, such as GRDL models. Its time complexity for each sample in each training epoch is $C_1 + O(K(n^2 + mn + m^2))$, as reported in their paper, where C_1 is the time complexity of the used GNN, K is the number of classes, and n, m are the number of nodes and edges in each graph. Thus, we can easily conclude $O(\text{JacobiGAD}) \leq O(\text{GRDL})$.

1287

Compared to heterogeneous graph-level classification models, such as HeGCL. Its time complexity for each sample in each training epoch is $Q|\mathcal{N}|^2 + |\mathcal{E}| + |\mathcal{E}^\Phi| + |\Phi||\mathcal{N}|$, as reported in their paper, where Q is the number of heads of attention layer, $|\mathcal{N}|, |\mathcal{E}|$ are the number of nodes and edges in

1296 **Algorithm 3:** RFACE

1297 **Input:** κ, y, z, γ, C
1298 **Output:** \mathcal{L}_{RFACE}
1299 1 $p \leftarrow \text{Sigmoid}(z);$
1300 2 $\Delta p_c \leftarrow -\gamma \kappa_c (p_c - y_c);$
1301 3 $\tilde{p} \leftarrow \text{Sigmoid}(z + \Delta p);$
1302 4 $\mathcal{L}_{RFACE} \leftarrow 0;$
1303 5 **for** c in $\{1, \dots, C\}$ **do**
1304 6 $\mathcal{L}_{RFACE} \leftarrow \mathcal{L}_{RFACE} + y_c \log(\tilde{p}_c);$
1305 7 **Return** $\mathcal{L}_{RFACE};$

1307

1308 **Algorithm 4:** JacobiGAD

1309 **Input:** $\mathcal{G}, f, K, T, \gamma, C, \epsilon, d, E$
1310 1 $\mathcal{G}' \leftarrow \text{Preprocess}(\mathcal{G}, d);$
1311 2 $\mathcal{L}_{RFACE} \leftarrow 0;$
1312 3 $\kappa \leftarrow \mathbf{0};$
1313 4 **for** c in $\{1, \dots, C\}$ **do**
1314 5 $\kappa_c \leftarrow \frac{f_c}{\max(f) + \epsilon};$
1315 6 **for** $epoch$ in $1, \dots, E$ **do**
1316 7 **for** G in \mathcal{G}' **do**
1317 8 $z \leftarrow \text{JPGNN}(G, \mathcal{A}, G, \mathbf{X}^{proj}, K, T);$
1318 9 $\mathcal{L}_{RFACE} \leftarrow \mathcal{L}_{RFACE} + \text{RFACE}(\kappa, G, y, z, \gamma, C);$
1319 10 $\mathcal{L}_{RFACE} \leftarrow -\frac{1}{|\mathcal{G}|} \mathcal{L}_{RFACE};$

1321

1322

1323 each graph, $|\Phi|$ is the number of meta-path, and $|\mathcal{E}^\Phi|$ is the number of meta-path-based edges. Thus,
1324 we can easily conclude $O(\text{JacobiGAD}) \leq O(\text{HeGCL})$.
1325

1326 Compared to GAD models, such as RQGNN. Its time complexity for each sample in each training
1327 epoch is $O(Kqnmd_{hid}^3)$, where K, q are the width and depth of the GNN, n, m are the number
1328 of nodes and edges in each graph, and d_{hid} is the hidden dimension of the layer of the GNN. Thus, we
1329 can easily conclude $O(\text{JacobiGAD}) \leq O(\text{RQGNN})$.

1330 To sum up, we compare the theoretical time complexity of JacobiGAD with representative previous
1331 works in different categories, and conclude that our JacobiGAD has practical cost for real deployment,
1332 as its time complexity is less than or equal to the previous works.

1333 Additionally, we further report the runtime and memory cost of JacobiGAD and compare them with
1334 all baselines across 3 datasets. The results in Table 6 show that JacobiGAD achieves competitive
1335 computational efficiency while maintaining state-of-the-art detection performance.

1336 Empirically, JacobiGAD’s training time is faster than most included baselines, and its total GPU
1337 memory usage stays within a comparable range. This indicates that the model scales well with both
1338 graph size and dataset difficulty. Notably, JacobiGAD maintains SOTA performance while requiring
1339 no additional memory-heavy modules. As a result, JacobiGAD provides a favorable trade-off between
1340 efficiency and accuracy: it preserves strong anomaly detection capability without incurring substantial
1341 computational cost.

1342 These observations confirm that the proposed method is not only effective but also practical for
1343 real-world heterogeneous graph-level anomaly detection scenarios where time and memory resources
1344 are often constrained.

1345

1346

D EXPERIMENTAL SETTINGS

1347

1348

1349 The hyperparameters used for training JacobiGAD are provided in Table 7. The model
1350 was tuned through an extensive grid search over the following values: learning rate $\eta \in$

1350 Table 6: Average wall-clock time (s) and total memory cost (MB).
1351

| Cost | Datasets | MCF-7 | | IMDB | | RCDD | |
|------------------------------------|------------|----------|---------|---------|---------|----------|----------|
| | | Memory | Time | Memory | Time | Memory | Time |
| Homogeneous Graph Classification | GCN | 800.88 | 82.09 | 643.80 | 12.06 | 1165.37 | 214.11 |
| | SAGE | 707.21 | 83.63 | 539.50 | 12.16 | 1065.95 | 219.77 |
| | GAT | 842.66 | 101.17 | 645.05 | 18.88 | 1167.25 | 242.17 |
| | GIN | 693.98 | 484.71 | 495.53 | 65.97 | 1027.38 | 1383.77 |
| | LRGNN | 845.38 | 1013.79 | 823.08 | 357.49 | 1344.03 | 1505.59 |
| | GRDL | 812.52 | 734.36 | 614.29 | 234.39 | 1301.55 | 1937.96 |
| | UQGNN | 1199.74 | 229.02 | 944.09 | 319.57 | 1335.56 | 267.89 |
| | UIL | 1177.79 | 332.46 | 918.78 | 60.83 | 1253.44 | 741.10 |
| Heterogeneous Graph Classification | HMGNN | 2098.57 | 342.67 | 913.00 | 121.24 | 2629.50 | 492.56 |
| | muxGNN | 5450.38 | 91.28 | 1325.74 | 27.08 | 12548.35 | 405.09 |
| | HeGCL | 17812.20 | 1933.91 | 1031.01 | 52.65 | 5156.08 | 762.66 |
| | RFAGNN | 1240.74 | 255.83 | 898.17 | 43.41 | 1743.74 | 523.85 |
| | SHGLNN | 653.80 | 571.93 | 799.83 | 105.06 | 1007.95 | 2123.87 |
| Graph-level Anomaly Detection | iGAD | 799.52 | 289.79 | 821.83 | 144.45 | 1247.28 | 801.91 |
| | GmapAD | 1335.84 | 1885.73 | 1077.30 | 85.76 | 2073.54 | 33755.93 |
| | RumorMixer | 739.81 | 6253.59 | 667.43 | 1401.91 | 1192.67 | 12094.97 |
| | RQGNN | 1177.37 | 2151.12 | 1062.88 | 248.91 | 1682.53 | 8817.05 |
| | UniGAD | 964.59 | 162.33 | 1077.91 | 110.22 | 1420.30 | 2013.64 |
| Ours | JacobiGAD | 1265.71 | 136.22 | 803.02 | 26.67 | 1250.07 | 170.07 |

1369 Table 7: Hyperparameters for different datasets, where η is learning rate, h_{dim} is hidden dimension
1370 of JPGNN layers, and K, T are the width and depth of the JPGNN.
1371

| Dataset | η | h_{dim} | K | T |
|-------------|--------|------------------|-----|-----|
| MCF-7 | 0.005 | 128 | 3 | 5 |
| MOLT-4 | 0.0001 | 256 | 4 | 5 |
| PC-3 | 0.001 | 128 | 2 | 5 |
| SW-620 | 0.005 | 64 | 4 | 4 |
| NCI-H23 | 0.001 | 256 | 3 | 5 |
| OVCAR-8 | 0.0005 | 256 | 4 | 5 |
| P388 | 0.0001 | 128 | 4 | 4 |
| SF-295 | 0.001 | 256 | 3 | 4 |
| SN12C | 0.001 | 256 | 3 | 4 |
| UACC257 | 0.0005 | 64 | 4 | 4 |
| DBLP | 0.0001 | 256 | 2 | 5 |
| IMDB | 0.005 | 128 | 2 | 5 |
| PDNS | 0.001 | 128 | 3 | 3 |
| RCDD | 0.005 | 128 | 1 | 3 |
| Transaction | 0.001 | 64 | 1 | 4 |

1388 {0.005, 0.001, 0.0005, 0.0001}, hidden dimension size $h_{\text{dim}} \in \{64, 128, 256\}$, $K \in \{1, 2, 3, 4\}$,
1389 and $T \in \{2, 3, 4, 5\}$. The optimal hyperparameter set was chosen based on the best composite
1390 performance, considering AUROC, AUPRC, Recall@k, and F1-score, on the validation set, and we
1391 report the test results for this configuration. Note that for hyperparameters in RFACE, i.e., ϵ and γ , we
1392 set them as default values $1e-8$ and 0.3 respectively, as they reach a relatively better performance.
1393 All trials were executed on an NVIDIA Quadro RTX 8000 to maintain a consistent experimental
1394 environment.

1396 E ADDITIONAL EXPERIMENTAL RESULTS

1397 To further demonstrate the robustness and generality of our method, we conduct additional experiments
1398 on 7 public graph benchmarks: MCF-7, MOLT-4, PC-3, SW-620, NCI-H23, OVCAR-8, and
1399 P388. Evaluating on this expanded set enables a more rigorous assessment of our model’s ability to
1400 generalize across different graph distributions.

1401 Across all datasets, our method consistently outperforms representative homogeneous graph clas-
1402 sification baselines, shown in Table 8, heterogeneous graph classification approaches, shown in

1404 Table 8: Average performance with multiple runs (homogeneous graph classification models).
1405

| Datasets | Metrics | GCN | SAGE | GAT | GIN | LRGNN | GRDL | UQGNN | UIL | JacobiGAD |
|----------|----------|--------|--------|--------|--------|--------|--------|--------|--------|---------------|
| MCF-7 | AUROC | 0.6720 | 0.7264 | 0.6971 | 0.7110 | 0.7070 | 0.5349 | 0.5238 | 0.7162 | 0.7679 |
| | AUPRC | 0.1460 | 0.2318 | 0.1983 | 0.1730 | 0.2434 | 0.1115 | 0.0878 | 0.2367 | 0.3403 |
| | Recall@k | 0.1808 | 0.2818 | 0.2585 | 0.2150 | 0.2825 | 0.1590 | 0.0792 | 0.2716 | 0.3769 |
| | F1-score | 0.4783 | 0.5949 | 0.5832 | 0.4872 | 0.5458 | 0.4785 | 0.4785 | 0.5503 | 0.6597 |
| MOLT-4 | AUROC | 0.6628 | 0.7155 | 0.7076 | 0.6880 | 0.7334 | 0.5257 | 0.5856 | 0.6867 | 0.7381 |
| | AUPRC | 0.1367 | 0.2051 | 0.2288 | 0.1611 | 0.2353 | 0.1208 | 0.1016 | 0.1735 | 0.3097 |
| | Recall@k | 0.1624 | 0.2649 | 0.2797 | 0.2033 | 0.2962 | 0.1948 | 0.1200 | 0.2171 | 0.3519 |
| | F1-score | 0.4803 | 0.5903 | 0.6100 | 0.4954 | 0.4906 | 0.4794 | 0.4794 | 0.4794 | 0.6507 |
| PC-3 | AUROC | 0.6717 | 0.7157 | 0.7391 | 0.7109 | 0.7389 | 0.5106 | 0.5395 | 0.7352 | 0.7677 |
| | AUPRC | 0.1017 | 0.1869 | 0.1973 | 0.1321 | 0.2289 | 0.0797 | 0.0693 | 0.2008 | 0.3064 |
| | Recall@k | 0.1318 | 0.2359 | 0.2657 | 0.1679 | 0.3050 | 0.1360 | 0.0903 | 0.2370 | 0.3603 |
| | F1-score | 0.4941 | 0.5882 | 0.6064 | 0.4853 | 0.5120 | 0.4853 | 0.4853 | 0.5229 | 0.6394 |
| SW-620 | AUROC | 0.7000 | 0.7619 | 0.7187 | 0.7270 | 0.7660 | 0.5601 | 0.5525 | 0.7202 | 0.7728 |
| | AUPRC | 0.1342 | 0.2137 | 0.1737 | 0.1371 | 0.2281 | 0.0996 | 0.0654 | 0.1484 | 0.2697 |
| | Recall@k | 0.1715 | 0.2697 | 0.2172 | 0.1853 | 0.2918 | 0.1362 | 0.0539 | 0.1777 | 0.3347 |
| | F1-score | 0.4990 | 0.5880 | 0.5695 | 0.5509 | 0.4896 | 0.4847 | 0.4847 | 0.4888 | 0.6461 |
| NCI-H23 | AUROC | 0.6950 | 0.7416 | 0.7703 | 0.7284 | 0.7812 | 0.5150 | 0.5254 | 0.7656 | 0.7900 |
| | AUPRC | 0.1064 | 0.1976 | 0.1904 | 0.1299 | 0.2056 | 0.0945 | 0.0546 | 0.1916 | 0.2927 |
| | Recall@k | 0.1296 | 0.2623 | 0.2502 | 0.1595 | 0.2632 | 0.1441 | 0.0583 | 0.2421 | 0.3417 |
| | F1-score | 0.4939 | 0.5667 | 0.5979 | 0.5452 | 0.5566 | 0.4869 | 0.4869 | 0.5056 | 0.6546 |
| OVCAR-8 | AUROC | 0.6791 | 0.7464 | 0.7296 | 0.6917 | 0.7467 | 0.5213 | 0.5148 | 0.7152 | 0.7762 |
| | AUPRC | 0.0947 | 0.1797 | 0.1840 | 0.1100 | 0.2066 | 0.0834 | 0.0530 | 0.1429 | 0.2888 |
| | Recall@k | 0.1162 | 0.2452 | 0.2652 | 0.1346 | 0.2388 | 0.1474 | 0.0585 | 0.1867 | 0.3438 |
| | F1-score | 0.4882 | 0.5728 | 0.5886 | 0.5598 | 0.4893 | 0.4868 | 0.4868 | 0.4876 | 0.6461 |
| P388 | AUROC | 0.6444 | 0.7424 | 0.7269 | 0.7391 | 0.7148 | 0.6196 | 0.5169 | 0.7375 | 0.7896 |
| | AUPRC | 0.0911 | 0.2151 | 0.2413 | 0.2131 | 0.1508 | 0.2493 | 0.0615 | 0.2175 | 0.3929 |
| | Recall@k | 0.1255 | 0.3067 | 0.3336 | 0.3009 | 0.2045 | 0.2843 | 0.0718 | 0.2980 | 0.4431 |
| | F1-score | 0.4912 | 0.4942 | 0.4990 | 0.6053 | 0.4858 | 0.4858 | 0.4858 | 0.5608 | 0.7061 |

1425 Table 9: Average performance with multiple runs (heterogeneous graph classification models).
1426

| Datasets | Metrics | HMGNN | muxGNN | HeGCL | RFAGNN | SHGLNN | JacobiGAD |
|----------|----------|--------|--------|--------|--------|--------|---------------|
| MCF-7 | AUROC | 0.3652 | 0.5570 | 0.6733 | 0.6990 | 0.5079 | 0.7679 |
| | AUPRC | 0.0646 | 0.1402 | 0.1735 | 0.1829 | 0.0774 | 0.3403 |
| | Recall@k | 0.0792 | 0.2070 | 0.2367 | 0.2186 | 0.0378 | 0.3769 |
| | F1-score | 0.4965 | 0.4848 | 0.4817 | 0.5022 | 0.4785 | 0.6597 |
| MOLT-4 | AUROC | 0.5068 | 0.5009 | 0.6675 | 0.6540 | 0.4980 | 0.7381 |
| | AUPRC | 0.1038 | 0.0809 | 0.1462 | 0.1472 | 0.0737 | 0.3097 |
| | Recall@k | 0.1242 | 0.0801 | 0.1874 | 0.1773 | 0.0483 | 0.3519 |
| | F1-score | 0.5084 | 0.5003 | 0.4794 | 0.4849 | 0.4794 | 0.6507 |
| PC-3 | AUROC | 0.5359 | 0.4511 | 0.6913 | 0.6923 | 0.5201 | 0.7677 |
| | AUPRC | 0.0903 | 0.0516 | 0.1372 | 0.1157 | 0.0547 | 0.3064 |
| | Recall@k | 0.1010 | 0.0308 | 0.1690 | 0.1456 | 0.0202 | 0.3603 |
| | F1-score | 0.4869 | 0.4909 | 0.4852 | 0.4877 | 0.4853 | 0.6394 |
| SW-620 | AUROC | 0.5392 | 0.4823 | 0.6610 | 0.6633 | 0.5012 | 0.7728 |
| | AUPRC | 0.0678 | 0.0629 | 0.1419 | 0.1114 | 0.0555 | 0.2697 |
| | Recall@k | 0.0816 | 0.0781 | 0.1978 | 0.1660 | 0.0207 | 0.3347 |
| | F1-score | 0.4854 | 0.4860 | 0.4853 | 0.4846 | 0.4847 | 0.6461 |
| NCI-H23 | AUROC | 0.3299 | 0.5302 | 0.7090 | 0.6834 | 0.5212 | 0.7900 |
| | AUPRC | 0.0357 | 0.0615 | 0.1621 | 0.1024 | 0.0490 | 0.2927 |
| | Recall@k | 0.0300 | 0.0858 | 0.2340 | 0.1482 | 0.0178 | 0.3417 |
| | F1-score | 0.4867 | 0.5160 | 0.4891 | 0.4869 | 0.4869 | 0.6546 |
| OVCAR-8 | AUROC | 0.4711 | 0.4449 | 0.6673 | 0.6691 | 0.5154 | 0.7762 |
| | AUPRC | 0.0634 | 0.0457 | 0.1188 | 0.0950 | 0.0489 | 0.2888 |
| | Recall@k | 0.0817 | 0.0481 | 0.1707 | 0.1282 | 0.0224 | 0.3438 |
| | F1-score | 0.4867 | 0.4872 | 0.4883 | 0.4960 | 0.4868 | 0.6461 |
| P388 | AUROC | 0.5734 | 0.4296 | 0.5765 | 0.7270 | 0.6333 | 0.7896 |
| | AUPRC | 0.0925 | 0.0632 | 0.0749 | 0.1745 | 0.0774 | 0.3929 |
| | Recall@k | 0.1313 | 0.1066 | 0.0964 | 0.2379 | 0.0718 | 0.4431 |
| | F1-score | 0.4858 | 0.5284 | 0.4878 | 0.5663 | 0.4858 | 0.7061 |

1449 Table 9, and graph-level anomaly detection methods, shown in Table 10. These results reinforce the
1450 effectiveness and broad applicability of our approach and confirm that the improvements are not
1451 confined to a narrow set of benchmarks but hold across a diverse collection of graph domains.
14521453

F ABLATION STUDY

14541455 In this section, we will further analyze the influence of different components in JacobiGAD. To be
1456 specific, we will investigate different components in three dimensions, that is, component deactivation
1457 (additional experiments), input replacement, and polynomial degradation.
1458

1458 Table 10: Average performance with multiple runs (GAD models).
1459

| Datasets | Metrics | iGAD | GmapAD | RumorMixer | RQGNN | UniGAD | JacobiGAD |
|----------|----------|--------|--------|------------|--------|--------|---------------|
| MCF-7 | AUROC | 0.7140 | 0.5889 | 0.3951 | 0.7332 | 0.5987 | 0.7679 |
| | AUPRC | 0.1913 | 0.1001 | 0.0663 | 0.2585 | 0.1133 | 0.3403 |
| | Recall@k | 0.2629 | 0.1147 | 0.0621 | 0.3065 | 0.1467 | 0.3769 |
| | F1-score | 0.5637 | 0.4046 | 0.4785 | 0.5768 | 0.5091 | 0.6597 |
| MOLT-4 | AUROC | 0.7111 | 0.6108 | 0.4985 | 0.7082 | 0.5880 | 0.7381 |
| | AUPRC | 0.2025 | 0.1018 | 0.0789 | 0.2248 | 0.1067 | 0.3097 |
| | Recall@k | 0.2749 | 0.1056 | 0.0786 | 0.2845 | 0.1433 | 0.3519 |
| | F1-score | 0.5766 | 0.4351 | 0.4794 | 0.6072 | 0.4915 | 0.6507 |
| PC-3 | AUROC | 0.7040 | 0.5707 | 0.3878 | 0.7260 | 0.6308 | 0.7677 |
| | AUPRC | 0.1254 | 0.0665 | 0.0448 | 0.2143 | 0.0958 | 0.3064 |
| | Recall@k | 0.1807 | 0.0755 | 0.0287 | 0.2880 | 0.1403 | 0.3603 |
| | F1-score | 0.5117 | 0.3826 | 0.4853 | 0.6207 | 0.5091 | 0.6394 |
| SW-620 | AUROC | 0.7280 | 0.6058 | 0.4230 | 0.7687 | 0.6195 | 0.7728 |
| | AUPRC | 0.1699 | 0.0760 | 0.0498 | 0.2105 | 0.0972 | 0.2697 |
| | Recall@k | 0.2254 | 0.0761 | 0.0346 | 0.2621 | 0.1480 | 0.3347 |
| | F1-score | 0.5606 | 0.4092 | 0.4847 | 0.5883 | 0.5000 | 0.6461 |
| NCI-H23 | AUROC | 0.7531 | 0.5556 | 0.4007 | 0.7817 | 0.6276 | 0.7900 |
| | AUPRC | 0.1616 | 0.0572 | 0.0408 | 0.2618 | 0.0877 | 0.2927 |
| | Recall@k | 0.2316 | 0.0615 | 0.0324 | 0.3142 | 0.1377 | 0.3417 |
| | F1-score | 0.5577 | 0.2746 | 0.4869 | 0.6258 | 0.5087 | 0.6546 |
| OVCAR-8 | AUROC | 0.7205 | 0.6147 | 0.4042 | 0.7381 | 0.5975 | 0.7762 |
| | AUPRC | 0.1449 | 0.0673 | 0.0414 | 0.1973 | 0.0734 | 0.2888 |
| | Recall@k | 0.2228 | 0.0793 | 0.0248 | 0.2596 | 0.0954 | 0.3438 |
| | F1-score | 0.5538 | 0.3971 | 0.4868 | 0.5903 | 0.4542 | 0.6461 |
| P388 | AUROC | 0.6776 | 0.5620 | 0.4369 | 0.7625 | 0.6065 | 0.7896 |
| | AUPRC | 0.1989 | 0.0643 | 0.0504 | 0.2572 | 0.0763 | 0.3929 |
| | Recall@k | 0.2828 | 0.0790 | 0.0464 | 0.3256 | 0.0892 | 0.4431 |
| | F1-score | 0.5622 | 0.4655 | 0.4858 | 0.5559 | 0.4862 | 0.7061 |

1479 Table 11: Ablation study for component deactivation.
1480

| Datasets | Metrics | JacobiGAD | w/o \mathcal{L}_{RFACE} | w/o learnable (α, β) | w/o learnable ω_r |
|----------|----------|-----------|---------------------------|---------------------------------|--------------------------|
| MCF-7 | AUROC | 0.7679 | 0.7400 | 0.7358 | 0.7181 |
| | AUPRC | 0.3403 | 0.2794 | 0.2763 | 0.2517 |
| | Recall@k | 0.3769 | 0.3232 | 0.3275 | 0.2992 |
| | F1-score | 0.6597 | 0.5945 | 0.5982 | 0.6124 |
| MOLT-4 | AUROC | 0.7381 | 0.7216 | 0.7280 | 0.7186 |
| | AUPRC | 0.3097 | 0.2927 | 0.2593 | 0.2807 |
| | Recall@k | 0.3519 | 0.3455 | 0.3174 | 0.3349 |
| | F1-score | 0.6507 | 0.6375 | 0.6302 | 0.6400 |
| PC-3 | AUROC | 0.7677 | 0.7526 | 0.7308 | 0.7584 |
| | AUPRC | 0.3064 | 0.2589 | 0.2578 | 0.2780 |
| | Recall@k | 0.3603 | 0.2986 | 0.3209 | 0.3390 |
| | F1-score | 0.6394 | 0.6087 | 0.6172 | 0.6378 |
| SW-620 | AUROC | 0.7728 | 0.7457 | 0.7480 | 0.7553 |
| | AUPRC | 0.2697 | 0.2434 | 0.2467 | 0.2568 |
| | Recall@k | 0.3347 | 0.3071 | 0.3119 | 0.3098 |
| | F1-score | 0.6461 | 0.6031 | 0.6210 | 0.6236 |
| NCI-H23 | AUROC | 0.7900 | 0.7684 | 0.7895 | 0.7735 |
| | AUPRC | 0.2927 | 0.2444 | 0.2647 | 0.2899 |
| | Recall@k | 0.3417 | 0.3158 | 0.3142 | 0.3409 |
| | F1-score | 0.6546 | 0.5967 | 0.6395 | 0.6265 |
| OVCAR-8 | AUROC | 0.7762 | 0.7521 | 0.7636 | 0.7578 |
| | AUPRC | 0.2888 | 0.2189 | 0.2532 | 0.2371 |
| | Recall@k | 0.3438 | 0.2716 | 0.3117 | 0.2925 |
| | F1-score | 0.6461 | 0.6003 | 0.6130 | 0.6022 |
| P388 | AUROC | 0.7896 | 0.7332 | 0.7699 | 0.7733 |
| | AUPRC | 0.3929 | 0.3262 | 0.3693 | 0.3729 |
| | Recall@k | 0.4431 | 0.3836 | 0.4032 | 0.4119 |
| | F1-score | 0.7061 | 0.6762 | 0.6354 | 0.6698 |

1503
1504 F.1 ADDITIONAL COMPONENT DEACTIVATION1505
1506 To further validate the contribution of each component in our framework, we conduct an extended
1507 ablation study on 7 other benchmark datasets: MCF-7, MOLT-4, PC-3, SW-620, NCI-H23, OVCAR-8,
1508 and P388.1509
1510 In this expanded evaluation of Table 11, we follow the same setting shown in Section 5.3. Across
1511 all datasets, the full model consistently achieves the highest detection scores, while removing any
major component leads to clear and reproducible degradation. These results collectively demonstrate

1512 Table 12: Ablation study for input replacement.
1513

| Datasets | Metrics | JacobiGAD | SVD | Concat | MLP |
|----------|----------|-----------|--------|--------|--------|
| MCF-7 | AUROC | 0.7679 | 0.7258 | 0.7235 | 0.7393 |
| | AUPRC | 0.3403 | 0.2549 | 0.2459 | 0.3060 |
| | Recall@k | 0.3769 | 0.2905 | 0.3028 | 0.3508 |
| | F1-score | 0.6597 | 0.5935 | 0.5972 | 0.6419 |
| MOLT-4 | AUROC | 0.7381 | 0.7257 | 0.7223 | 0.7106 |
| | AUPRC | 0.3097 | 0.2607 | 0.2745 | 0.2412 |
| | Recall@k | 0.3519 | 0.3137 | 0.3471 | 0.3052 |
| | F1-score | 0.6507 | 0.6292 | 0.6377 | 0.6163 |
| PC-3 | AUROC | 0.7677 | 0.7645 | 0.7378 | 0.7397 |
| | AUPRC | 0.3064 | 0.2669 | 0.2450 | 0.2537 |
| | Recall@k | 0.3603 | 0.3220 | 0.2880 | 0.3092 |
| | F1-score | 0.6394 | 0.5849 | 0.6113 | 0.6256 |
| SW-620 | AUROC | 0.7728 | 0.7576 | 0.7322 | 0.7365 |
| | AUPRC | 0.2697 | 0.2508 | 0.2233 | 0.2319 |
| | Recall@k | 0.3347 | 0.3133 | 0.2863 | 0.2953 |
| | F1-score | 0.6461 | 0.6209 | 0.6162 | 0.6195 |
| NCI-H23 | AUROC | 0.7900 | 0.7718 | 0.7319 | 0.7474 |
| | AUPRC | 0.2927 | 0.2556 | 0.2317 | 0.2394 |
| | Recall@k | 0.3417 | 0.3320 | 0.2858 | 0.2947 |
| | F1-score | 0.6546 | 0.6304 | 0.6046 | 0.6129 |
| OVCAR-8 | AUROC | 0.7762 | 0.7567 | 0.7444 | 0.7729 |
| | AUPRC | 0.2888 | 0.2333 | 0.2319 | 0.2462 |
| | Recall@k | 0.3438 | 0.2877 | 0.2965 | 0.2901 |
| | F1-score | 0.6461 | 0.5892 | 0.5923 | 0.5889 |
| P388 | AUROC | 0.7896 | 0.7779 | 0.7862 | 0.7439 |
| | AUPRC | 0.3929 | 0.3501 | 0.3702 | 0.3376 |
| | Recall@k | 0.4431 | 0.4054 | 0.4271 | 0.4054 |
| | F1-score | 0.7061 | 0.6592 | 0.6938 | 0.6760 |
| SF-295 | AUROC | 0.7729 | 0.7554 | 0.7563 | 0.7421 |
| | AUPRC | 0.2623 | 0.2179 | 0.1935 | 0.1820 |
| | Recall@k | 0.3210 | 0.3004 | 0.2733 | 0.2502 |
| | F1-score | 0.6356 | 0.6124 | 0.6103 | 0.5859 |
| SN12C | AUROC | 0.7797 | 0.7682 | 0.7523 | 0.7568 |
| | AUPRC | 0.2666 | 0.2457 | 0.2671 | 0.2430 |
| | Recall@k | 0.3240 | 0.3078 | 0.3291 | 0.3018 |
| | F1-score | 0.6329 | 0.6329 | 0.6380 | 0.6221 |
| UACC257 | AUROC | 0.7613 | 0.7692 | 0.7440 | 0.7505 |
| | AUPRC | 0.1995 | 0.2218 | 0.1660 | 0.1988 |
| | Recall@k | 0.2819 | 0.2901 | 0.2475 | 0.2677 |
| | F1-score | 0.6246 | 0.6214 | 0.5775 | 0.5930 |
| DBLP | AUROC | 0.9830 | 0.9679 | 0.9780 | 0.9805 |
| | AUPRC | 0.9842 | 0.9665 | 0.9796 | 0.9822 |
| | Recall@k | 0.9575 | 0.9374 | 0.9441 | 0.9508 |
| | F1-score | 0.9651 | 0.9294 | 0.9576 | 0.9623 |
| IMDB | AUROC | 0.7263 | 0.6665 | 0.6961 | 0.6905 |
| | AUPRC | 0.7619 | 0.7282 | 0.7318 | 0.7215 |
| | Recall@k | 0.7192 | 0.6824 | 0.7035 | 0.6951 |
| | F1-score | 0.6585 | 0.6046 | 0.6445 | 0.6402 |
| PDNS | AUROC | 0.8728 | 0.8697 | 0.8673 | 0.8707 |
| | AUPRC | 0.6871 | 0.6650 | 0.6801 | 0.6865 |
| | Recall@k | 0.6283 | 0.6206 | 0.6204 | 0.6247 |
| | F1-score | 0.7760 | 0.7526 | 0.7708 | 0.7676 |
| RCDD | AUROC | 0.9826 | 0.9777 | 0.9829 | 0.9806 |
| | AUPRC | 0.9332 | 0.9174 | 0.9322 | 0.9288 |
| | Recall@k | 0.8747 | 0.8550 | 0.8775 | 0.8623 |
| | F1-score | 0.9280 | 0.9161 | 0.9284 | 0.9194 |

1557 that each component contributes meaningfully to the final performance and that their combination is
1558 essential for achieving the strong detection capability of our method.
1559

1560 F.2 INPUT REPLACEMENT

1561 Next, we investigate the influence of different ways of input for JacobiGAD, i.e., SVD, Concat, and
1562 MLP.
1563

1564 As shown in Table 12, this ablation study evaluates the efficacy of the proposed input function in
1565 JacobiGAD for unifying features from different views in a heterogeneous graph by comparing it

against three common alternative methods: SVD (which may lose critical information), Concat (which creates a high-dimensional feature space), and MLP (which causes higher computational cost and may easily overfit). The results consistently demonstrate that JacobiGAD’s specialized integration method, Gaussian projection, significantly outperforms all three alternatives across the vast majority of datasets and metrics. Although there are rare, minor exceptions where an alternative method performs comparably on some datasets, the overall trend is unequivocal: the custom-designed input function in JacobiGAD is uniquely capable of effectively synthesizing heterogeneous information, which is a critical factor in the model’s superior anomaly detection performance.

F.3 POLYNOMIAL DEGRADATION

Finally, we analyze the influence of different polynomials for JacobiGAD, i.e., Gegenbauer ($\alpha = \beta = \lambda - \frac{1}{2}$), Chebyshev ($\alpha = \beta = -\frac{1}{2}$), and Legendre ($\alpha = \beta = 0$).

As shown in Table 13, the ablation study demonstrates that the choice of polynomial basis for the graph filter is critical, with the proposed Jacobi polynomials consistently outperforming Gegenbauer, Chebyshev, and Legendre polynomials across all datasets and metrics. The key drawback of these alternative polynomials is their inherent rigidity. Unlike the parameter-rich Jacobi basis, which can be adaptively tuned to fit the complex spectral characteristics of heterogeneous graphs, the fixed spectral response of Chebyshev and Legendre polynomials and the limited single-parameter flexibility of Gegenbauer polynomials render them less capable of capturing the nuanced patterns necessary for effective anomaly detection. This lack of adaptability manifests clearly in the significant performance gaps, indicating that the alternative filters struggle to generate the highly discriminative representations needed to reliably separate anomalies from normal nodes in complex graph data.

G LEARNED PARAMETERS

In this section, we will present the learned parameters of one run of our experiment to show the influence of different parameters on all datasets.

The blank slot of Table 14 is due to the best K for different datasets not being the same. As shown in Table 14, the results further demonstrate the importance of learnable (α, β) as the best performance of different datasets requires distinct combinations of (α, β) , instead of fixed parameters.

In Table 15, we present the statistical information of ω_r , due to the large number of different relations in heterogeneous datasets. We use the row Range as the start and the end of the range. For example, for the first range of MCF-7, it is formed by $[-1.4456, -1.0144]$. And the corresponding frequency is reported in the row Frequency. In this case, the frequency of $[-1.4456, -1.0144]$ is 1. Other cases can be deduced by analogy. We can be informed by Table 15 that the learnable ω_r is of vital importance for heterogeneous GAD, as the best ω_r for different datasets can distribute evenly, focus on the center part, or lie mainly on the extreme spots.

H COMPARISON WITH FOCAL LOSS

We further compare our proposed RFACE with Focal loss, a classical loss for imbalanced data, to demonstrate the effectiveness of our proposed methods.

Assume we have logits $\mathbf{z} = [z_1, \dots, z_C]$, where C is the number of classes, sigmoid per class $\mathbf{p} = [p_1, \dots, p_C]$, where $p_i = \text{Sigmoid}(z_i)$, and multi-label target $\mathbf{y} = [y_1, \dots, y_C] \in \{0, 1\}^C$, then we will investigate the gradients of Cross-Entropy loss, Focal loss, and RFACE to show the key advantages of RFACE.

For Cross-Entropy loss:

$$\mathcal{L}_{CE} = - \sum_{i=1}^C [y_i \log p_i + (1 - y_i) \log(1 - p_i)],$$

the gradient vector is:

$$\nabla_{\mathbf{z}} \mathcal{L}_{CE} = \mathbf{p} - \mathbf{y}$$

Table 13: Ablation study for polynomial degradation.

| Datasets | Metrics | JacobiGAD | Gegenbauer | Chebyshev | Legendre |
|----------|----------|-----------|------------|-----------|----------|
| MCF-7 | AUROC | 0.7679 | 0.7281 | 0.7146 | 0.6986 |
| | AUPRC | 0.3403 | 0.2787 | 0.2375 | 0.2384 |
| | Recall@k | 0.3769 | 0.3217 | 0.2847 | 0.2767 |
| | F1-score | 0.6597 | 0.6048 | 0.5889 | 0.5886 |
| MOLT-4 | AUROC | 0.7381 | 0.7180 | 0.7076 | 0.7202 |
| | AUPRC | 0.3097 | 0.2587 | 0.2476 | 0.2475 |
| | Recall@k | 0.3519 | 0.3068 | 0.3100 | 0.3132 |
| | F1-score | 0.6507 | 0.6225 | 0.6248 | 0.6224 |
| PC-3 | AUROC | 0.7677 | 0.7226 | 0.7589 | 0.7384 |
| | AUPRC | 0.3064 | 0.2057 | 0.2494 | 0.2281 |
| | Recall@k | 0.3603 | 0.2508 | 0.3092 | 0.2944 |
| | F1-score | 0.6394 | 0.5689 | 0.6061 | 0.5974 |
| SW-620 | AUROC | 0.7728 | 0.7386 | 0.7441 | 0.7381 |
| | AUPRC | 0.2697 | 0.2392 | 0.2461 | 0.2233 |
| | Recall@k | 0.3347 | 0.3105 | 0.2988 | 0.2766 |
| | F1-score | 0.6461 | 0.5995 | 0.6294 | 0.5905 |
| NCI-H23 | AUROC | 0.7900 | 0.7758 | 0.7891 | 0.7727 |
| | AUPRC | 0.2927 | 0.2413 | 0.2556 | 0.2274 |
| | Recall@k | 0.3417 | 0.3045 | 0.3296 | 0.3020 |
| | F1-score | 0.6546 | 0.5775 | 0.5679 | 0.6101 |
| OVCAR-8 | AUROC | 0.7762 | 0.7703 | 0.7733 | 0.7691 |
| | AUPRC | 0.2888 | 0.2551 | 0.2388 | 0.2367 |
| | Recall@k | 0.3438 | 0.3101 | 0.3117 | 0.3117 |
| | F1-score | 0.6461 | 0.6286 | 0.6048 | 0.6003 |
| P388 | AUROC | 0.7896 | 0.7554 | 0.7564 | 0.7656 |
| | AUPRC | 0.3929 | 0.3465 | 0.3424 | 0.3587 |
| | Recall@k | 0.4431 | 0.3952 | 0.4054 | 0.3988 |
| | F1-score | 0.7061 | 0.6855 | 0.6886 | 0.6790 |
| SF-295 | AUROC | 0.7729 | 0.7670 | 0.7578 | 0.7461 |
| | AUPRC | 0.2623 | 0.2085 | 0.1939 | 0.1919 |
| | Recall@k | 0.3210 | 0.2724 | 0.2634 | 0.2634 |
| | F1-score | 0.6356 | 0.5965 | 0.5965 | 0.6020 |
| SN12C | AUROC | 0.7797 | 0.7459 | 0.7384 | 0.7404 |
| | AUPRC | 0.2666 | 0.2284 | 0.2308 | 0.2224 |
| | Recall@k | 0.3240 | 0.3018 | 0.2864 | 0.2805 |
| | F1-score | 0.6329 | 0.6052 | 0.6080 | 0.6138 |
| UACC257 | AUROC | 0.7613 | 0.7484 | 0.7390 | 0.6997 |
| | AUPRC | 0.1995 | 0.1995 | 0.1715 | 0.1587 |
| | Recall@k | 0.2819 | 0.2708 | 0.2525 | 0.2231 |
| | F1-score | 0.6246 | 0.5656 | 0.5659 | 0.5827 |
| DBLP | AUROC | 0.9830 | 0.9778 | 0.9756 | 0.9746 |
| | AUPRC | 0.9842 | 0.9802 | 0.9750 | 0.9732 |
| | Recall@k | 0.9575 | 0.9508 | 0.9463 | 0.9418 |
| | F1-score | 0.9651 | 0.9632 | 0.9557 | 0.9539 |
| IMDB | AUROC | 0.7263 | 0.7096 | 0.7130 | 0.7060 |
| | AUPRC | 0.7619 | 0.7485 | 0.7552 | 0.7444 |
| | Recall@k | 0.7192 | 0.7119 | 0.7108 | 0.7098 |
| | F1-score | 0.6585 | 0.6527 | 0.6461 | 0.6565 |
| PDNS | AUROC | 0.8728 | 0.8724 | 0.8721 | 0.8716 |
| | AUPRC | 0.6871 | 0.6732 | 0.6844 | 0.6860 |
| | Recall@k | 0.6283 | 0.6145 | 0.6270 | 0.6241 |
| | F1-score | 0.7760 | 0.7700 | 0.7704 | 0.7702 |
| RCDD | AUROC | 0.9826 | 0.9809 | 0.9805 | 0.9815 |
| | AUPRC | 0.9332 | 0.9283 | 0.9290 | 0.9299 |
| | Recall@k | 0.8747 | 0.8741 | 0.8719 | 0.8667 |
| | F1-score | 0.9280 | 0.9279 | 0.9229 | 0.9220 |

For Focal loss:

$$\mathcal{L}_{Focal} = - \sum_{i=1}^C [y_i(1-p_i)^\gamma \log p_i + (1-y_i)p_i^\gamma \log(1-p_i)],$$

the gradient vector is:

$$\nabla_{\mathbf{z}} \mathcal{L}_{Focal} = \mathbf{s}_{Focal}(\mathbf{p} - \mathbf{y}),$$

where \mathbf{s}_{Focal} is a scalar vector for each class i , depending on the ground truth label y_i , the predicted probability with no modification p_i , and the power for measuring the difficulty of samples γ .

Table 14: Learned (α, β) for Jacobi Polynomials.

| Datasets | α | | | | β | | | |
|----------|----------|--------|--------|--------|---------|--------|--------|--------|
| MCF-7 | 0.1123 | 1.2406 | 1.2030 | | 1.7027 | 1.5336 | 1.8209 | |
| MOLT-4 | 1.9611 | 0.8162 | 0.2852 | 1.8448 | 0.3078 | 1.0801 | 0.7354 | 1.5658 |
| PC-3 | 0.8800 | 1.6410 | | | 0.0865 | 0.8657 | | |
| SW-620 | 0.4661 | 1.608 | 1.3252 | 0.0181 | 1.9602 | 0.3216 | 0.4713 | 0.5893 |
| NCI-H23 | 0.3489 | 1.5247 | 1.9206 | | 1.8042 | 1.0248 | 1.3645 | |
| OVCAR-8 | 1.3458 | 1.6198 | 0.6928 | 1.0384 | 0.1690 | 0.0049 | 0.2966 | 1.1328 |
| P388 | 1.0701 | 0.5101 | 0.1645 | 0.2895 | 1.9218 | 1.5771 | 0.1775 | 0.1930 |
| SF-295 | 0.7552 | 0.326 | 1.4094 | | 0.3687 | 0.4355 | 1.8797 | |
| SN12C | 0.7820 | 0.8886 | 0.9901 | | 0.9222 | 0.1497 | 0.3569 | |
| UACC257 | 1.7794 | 1.7520 | 1.7412 | 1.9050 | 1.0449 | 0.2968 | 1.0761 | 0.3197 |
| DBLP | 1.9142 | 0.4518 | | | 1.7499 | 0.3554 | | |
| IMDB | 0.6390 | 0.5686 | | | 0.7820 | 1.3118 | | |
| PDNS | 0.6258 | 0.5431 | 1.9307 | | 0.1602 | 1.1881 | 0.7147 | |
| RCDD | 0.4016 | | | | 0.8339 | | | |

Table 15: Learned ω_r for different relations.

| Datasets | Metrics | Range | -1.4456 | -1.0144 | -0.5832 | -0.1519 | 0.2793 | 0.7105 | 1.1417 | 1.5729 | 2.0042 | 2.4354 | 2.8666 | |
|----------|---------|-----------|---------|---------|---------|---------|--------|--------|--------|--------|--------|--------|--------|--|
| MCF-7 | Range | Frequency | 1 | 0 | 1 | 19 | 30 | 27 | 16 | 25 | 7 | 3 | | |
| MOLT-4 | Range | Frequency | 0.0200 | 0.2136 | 0.4071 | 0.6006 | 0.7942 | 0.9877 | 1.1813 | 1.3748 | 1.5684 | 1.7619 | 1.9555 | |
| PC-3 | Range | Frequency | 0.0103 | 0.2200 | 0.4297 | 0.6394 | 0.8490 | 1.0587 | 1.2684 | 1.4781 | 1.6878 | 1.8974 | 2.1071 | |
| SW-620 | Range | Frequency | -1.1965 | -0.6950 | -0.1936 | 0.3078 | 0.8093 | 1.3107 | 1.8121 | 2.3136 | 2.8150 | 3.3164 | 3.8179 | |
| NCI-H23 | Range | Frequency | -0.0007 | 0.2072 | 0.4151 | 0.6230 | 0.8309 | 1.0389 | 1.2468 | 1.4547 | 1.6626 | 1.8705 | 2.0784 | |
| OVCAR-8 | Range | Frequency | -0.0148 | 0.1907 | 0.3963 | 0.6019 | 0.8074 | 1.0130 | 1.2185 | 1.4241 | 1.6296 | 1.8352 | 2.0407 | |
| P388 | Range | Frequency | 0.0009 | 0.2002 | 0.3995 | 0.5988 | 0.7981 | 0.9974 | 1.1967 | 1.3960 | 1.5953 | 1.7946 | 1.9939 | |
| SF-295 | Range | Frequency | -0.1725 | 0.0563 | 0.2852 | 0.5140 | 0.7429 | 0.9717 | 1.2006 | 1.4294 | 1.6583 | 1.8871 | 2.1160 | |
| SN12C | Range | Frequency | -0.0906 | 0.1206 | 0.3317 | 0.5428 | 0.7540 | 0.9651 | 1.1762 | 1.3873 | 1.5985 | 1.8096 | 2.0207 | |
| UACC257 | Range | Frequency | -0.0503 | 0.1549 | 0.3601 | 0.5653 | 0.7705 | 0.9757 | 1.1809 | 1.3861 | 1.5913 | 1.7965 | 2.0017 | |
| DBLP | Range | Frequency | -0.0094 | 0.1563 | 0.3220 | 0.4878 | 0.6535 | 0.8193 | 0.9850 | 1.1507 | 1.3165 | 1.4822 | 1.6480 | |
| IMDB | Range | Frequency | -0.0099 | 0.1620 | 0.3339 | 0.5058 | 0.6777 | 0.8496 | 1.0214 | 1.1933 | 1.3652 | 1.5371 | 1.7090 | |
| PDNS | Range | Frequency | 0.3526 | 0.4796 | 0.6065 | 0.7335 | 0.8604 | 0.9874 | 1.1143 | 1.2413 | 1.3682 | 1.4952 | 1.6221 | |
| RCDD | Range | Frequency | 0.7637 | 0.8787 | 0.9938 | 1.1089 | 1.2240 | 1.3390 | 1.4541 | 1.5692 | 1.6842 | 1.7993 | 1.9144 | |

For RFACE:

$$\mathcal{L}_{RFACE} = - \sum_{i=1}^C [y_i \log p'_i + (1 - y_i) \log(1 - p'_i)],$$

$$p'_i = \text{Sigmoid}(z_i + \gamma \kappa_i \nabla_{z_i} \mathcal{L}_{CE}),$$

the gradient vector is:

$$\nabla_{\mathbf{z}} \mathcal{L}_{RFACE} = \mathbf{s}_{RFACE}(\mathbf{p}' - \mathbf{y}),$$

where \mathbf{s}_{RFACE} is a scalar vector for each class i , depending on the ground truth label y_i , the predicted probability with logit modification (indicating class frequency by κ_i and sample difficulty by $\nabla_{z_i} \mathcal{L}_{CE}$) p_i , the coefficient of modification term γ , and the class-frequency-based curvature κ_i .

As shown above, the advantage of the proposed RFACE over Focal loss stems from its distinct mechanism for addressing class imbalance in heterogeneous graph-level anomaly detection. RFACE

Table 16: Ablation study for Focal Loss.

| Datasets | Metrics | JacobiGAD | w/ Focal Loss |
|----------|----------|-----------|---------------|
| MCF-7 | AUROC | 0.7679 | 0.7544 |
| | AUPRC | 0.3403 | 0.2846 |
| | Recall@k | 0.3769 | 0.3261 |
| | Macro-F1 | 0.6597 | 0.6106 |
| MOLT-4 | AUROC | 0.7381 | 0.7273 |
| | AUPRC | 0.3097 | 0.2612 |
| | Recall@k | 0.3519 | 0.3429 |
| | Macro-F1 | 0.6507 | 0.5796 |
| PC-3 | AUROC | 0.7677 | 0.7490 |
| | AUPRC | 0.3064 | 0.2275 |
| | Recall@k | 0.3603 | 0.3050 |
| | Macro-F1 | 0.6394 | 0.6203 |
| SW-620 | AUROC | 0.7728 | 0.7497 |
| | AUPRC | 0.2697 | 0.2537 |
| | Recall@k | 0.3347 | 0.3264 |
| | Macro-F1 | 0.6461 | 0.6281 |
| NCI-H23 | AUROC | 0.7900 | 0.7417 |
| | AUPRC | 0.2927 | 0.2281 |
| | Recall@k | 0.3417 | 0.3028 |
| | Macro-F1 | 0.6546 | 0.6255 |
| OVCAR-8 | AUROC | 0.7762 | 0.7310 |
| | AUPRC | 0.2888 | 0.2259 |
| | Recall@k | 0.3438 | 0.2925 |
| | Macro-F1 | 0.6461 | 0.5880 |
| P388 | AUROC | 0.7896 | 0.7503 |
| | AUPRC | 0.3929 | 0.3136 |
| | Recall@k | 0.4431 | 0.3749 |
| | Macro-F1 | 0.7061 | 0.6431 |
| SF-295 | AUROC | 0.7729 | 0.7369 |
| | AUPRC | 0.2623 | 0.1916 |
| | Recall@k | 0.3210 | 0.2691 |
| | Macro-F1 | 0.6356 | 0.5812 |
| SN12C | AUROC | 0.7797 | 0.7682 |
| | AUPRC | 0.2666 | 0.2182 |
| | Recall@k | 0.3240 | 0.2975 |
| | Macro-F1 | 0.6329 | 0.5929 |
| UACC257 | AUROC | 0.7613 | 0.7558 |
| | AUPRC | 0.1995 | 0.1749 |
| | Recall@k | 0.2819 | 0.2414 |
| | Macro-F1 | 0.6246 | 0.5580 |
| DBLP | AUROC | 0.9830 | 0.9754 |
| | AUPRC | 0.9842 | 0.9647 |
| | Recall@k | 0.9575 | 0.9463 |
| | Macro-F1 | 0.9651 | 0.9550 |
| IMDB | AUROC | 0.7263 | 0.6874 |
| | AUPRC | 0.7619 | 0.7371 |
| | Recall@k | 0.7192 | 0.6940 |
| | Macro-F1 | 0.6585 | 0.6291 |
| PDNS | AUROC | 0.8728 | 0.8612 |
| | AUPRC | 0.6871 | 0.6611 |
| | Recall@k | 0.6283 | 0.5981 |
| | Macro-F1 | 0.7760 | 0.7297 |
| RCDD | AUROC | 0.9826 | 0.9797 |
| | AUPRC | 0.9332 | 0.9306 |
| | Recall@k | 0.8747 | 0.8741 |
| | Macro-F1 | 0.9280 | 0.9276 |

applies a class-dependent and difficulty-aware logit transformation. This transformation modifies the optimization gradient based on both class frequency and sample difficulty. As a result, RFACE reshapes the decision boundary by explicitly expanding the margins of minority anomaly classes and gently contracting the margins of dominant normal classes, based on both class and difficulty. This curvature-inspired adjustment acts like a discrete Ricci flow step, improving the geometric regularity of the representation space by amplifying deviations. In heterogeneous graph data, where anomalies arise from subtle irregularities and specialized node–edge interactions, such margin rebalancing is

1782 crucial: it ensures that minority classes receive sustained and directionally beneficial updates even
1783 when the classifier becomes confident about them, preventing premature gradient vanishing.
1784

1785 In contrast, Focal loss only rescales the Cross-Entropy gradient through a difficulty-based factor
1786 s , which offers no mechanism to correct class-frequency-induced imbalance. Because Focal loss
1787 downweights "easy" samples regardless of their class, it may inadvertently suppress minority-class
1788 gradients once the model becomes moderately confident, leading to possible overfitting to "hard"
1789 samples. Moreover, Focal loss treats all classes identically and cannot incorporate global distributional
1790 information; the optimization trajectory therefore lacks the class-dependent curvature adjustment that
1791 RFACE introduces. This makes Focal loss sensitive to the randomness of minibatch composition,
1792 more prone to instability on small anomalous sets, and often ineffective when many anomalies are
1793 not "hard" samples, where class number information rather than prediction confidence determines
1794 anomaly separability.

1794 Beyond gradient modification, RFACE also provides additional advantages, as proved in Theorems 8
1795 (RFACE will amplify information of the minority class) and 9 (RFACE will converge), while Focal
1796 loss is a heuristic-oriented loss without guarantee. These theoretical advantages are also strongly
1797 supported by our empirical results. As shown in Table 16, RFACE consistently outperforms the Focal
1798 loss variant on all datasets.

1799 Together, these theoretical insights and empirical observations demonstrate that RFACE is signifi-
1800 cantly better suited than Focal loss for heterogeneous graph-level anomaly detection, offering stronger
1801 geometric corrections, more stable optimization, better calibration, and improved exploitation of
1802 graph information.

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

1815

1816

1817

1818

1819

1820

1821

1822

1823

1824

1825

1826

1827

1828

1829

1830

1831

1832

1833

1834

1835