



Kent Academic Repository

Sun, Li, Zhang, Ming, Jin, Wenxin, Sun, Zhongtian, Huang, Zhenhao, Peng, Hao, Su, Sen and Yu, Philip (2026) *Heterophily-agnostic hypergraph neural networks with Riemannian Local Exchanger*. In: *Proceedings of the ACM Web Conference 2026*. WWW '26: Proceedings of the ACM Web Conference 2026. . pp. 1092-1103. ACM ISBN 979-8-4007-2307-0.

Downloaded from

<https://kar.kent.ac.uk/114414/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1145/3774904.3792435>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY (Attribution)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

Heterophily-Agnostic Hypergraph Neural Networks with Riemannian Local Exchanger

Li Sun*
lsun@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Beijing, China

Zhongtian Sun
Z.Sun-256@kent.ac.uk
University of Kent
Canterbury, United Kingdom

Sen Su
susen@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Beijing, China

Ming Zhang
zhangming@ncepu.edu.cn
North China Electric Power
University
Beijing, China

Zhenhao Huang
huangzhenhao@ncepu.edu.cn
North China Electric Power
University
Beijing, China

Philip Yu
psyu@uic.edu
University of Illinois
Chicgo, United States

Wenxin Jin
jwx2265714977@outlook.com
North China Electric Power
University
Beijing, China

Hao Peng
penghao@buaa.edu.cn
Beihang University
Beijing, China

Abstract

Hypergraphs are the natural description of higher-order interactions among objects, widely applied in social network analysis, cross-modal retrieval, etc. Hypergraph Neural Networks (HGNNs) have become the dominant solution for learning on hypergraphs. Traditional HGNNs are extended from message passing graph neural networks, following the homophily assumption, and thus struggle with the prevalent heterophilic hypergraphs that call for long-range dependence modeling. Existing solutions enlarge the message flow through the hypergraph bottleneck, mitigating the oversquashing issue and capturing long-range dependence. However, they often accelerate the loss of representation distinguishability in the repeated aggregations, leading to oversmoothing. This dilemma motivates an interesting question: **Can we develop a unified mechanism that is agnostic to both homophilic and heterophilic hypergraphs?** In this paper, we achieve the best of both worlds through the lens of Riemannian geometry, which provides the potential to adjust the message passing behavior in different regions. The key insight lies in the connection between oversquashing and hypergraph bottleneck within the framework of Riemannian manifold heat flow. Building on this, we propose the novel idea of locally adapting the bottlenecks of different subhypergraphs. The core innovation of the proposed mechanism is the design of an **adaptive local (heat) exchanger**. Specifically, it captures the rich long-range dependencies via the Robin condition, and preserves the representation distinguishability via source terms,

thereby enabling heterophily-agnostic message passing with theoretical guarantees. Based on this theoretical foundation, we present a novel **Heat-Exchanger with Adaptive Locality** for Hypergraph Neural Network (**HealHGNN**), designed as a node-hyperedge bidirectional systems with linear complexity in the number of nodes and hyperedges. Extensive experiments on both homophilic and heterophilic cases show that HealHGNN achieves the state-of-the-art performance.

CCS Concepts

• **Computing methodologies** → **Neural networks**; • **Information systems** → **Web mining**;

Keywords

Hypergraph Neural Network, Heterophily, Riemannian Geometry, Long-range Dependency, Manifold Heat Flow.

ACM Reference Format:

Li Sun, Ming Zhang, Wenxin Jin, Zhongtian Sun, Zhenhao Huang, Hao Peng, Sen Su, and Philip Yu. 2026. Heterophily-Agnostic Hypergraph Neural Networks with Riemannian Local Exchanger. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3774904.3792435>

1 INTRODUCTION

Hypergraphs have emerged as a powerful tool for capturing higher-order interactions that involve more than two objects, and routinely find themselves in crucial applications, such as social network analysis and cross-modal retrieval [2, 22]. Hypergraph Neural Networks (HGNNs) have become the dominant solution for learning on hypergraphs. HGNNs [18] are typically extended from the message-passing graph neural networks, such as GCN [23] and GAT [42], using hypergraph Laplacian, and thus work under the homophily assumption that connections are formed between nodes

*Corresponding Author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *WWW '26, Dubai, United Arab Emirates*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2307-0/2026/04
<https://doi.org/10.1145/3774904.3792435>

of similar labels. Also, these methods inherit a fundamental drawback of conventional message-passing mechanism: oversmoothing (OSM) [6], where node representations become indistinguishable after repeated message aggregations, and oversquashing (OSQ) [48], where signals from distant nodes are compressed.

So far, HGNNs [13, 21, 43] are primarily designed for homophilic hypergraphs, while heterophilic hypergraphs are also ubiquitous in the real-world scenarios [25]. Learning on heterophilic hypergraphs necessitates the interaction between distant nodes that can be belonged to the same class. To capture long-range dependence, existing solutions [46, 48] make effort to enlarge the message flow through the hypergraph bottleneck, where messages are often squashed. However, they often accelerate the loss of representation distinguishability during repeated aggregations, as the drawback of oversmoothing is typically inherited. This dilemma motivates an interesting question: **Can we develop a unified mechanism that is agnostic to both homophilic and heterophilic hypergraphs?**

In this paper, we achieve the best of both worlds through the lens of Riemannian geometry. Our study begins with an interesting theoretical insight: oversquashing is inherently related to hypergraph bottleneck. Concretely, within the framework of Riemannian heat flow, hypergraph message passing is given as the heat flow governed by the Laplace–Beltrami operator. The heat kernel’s spectral decomposition, together with Cheeger–Buser’s Inequality, provides the fact that spectral gap λ is connected to hypergraph bottlenecks. Then, we rigorously prove that the increase of λ enlarges the bottleneck, promoting long-range propagation, meanwhile accelerating the decay of Dirichlet energy. That is, the increase of λ facilitates message flow but leads to severe oversmoothing.

Besides providing a systematic understanding of the dilemma, Riemannian geometry also enables the potential for heterophily-agnostic message passing. Traditional Euclidean spaces often struggle with seeking a single, tradeoff value of λ , given a global, uniform geometry everywhere in its space. On the contrary, a general Riemannian manifold allows for diversified geometries and behaviors in different regions, and thus affords more flexible, localized control: by treating a subhypergraph as a submanifold, we can modulate its spectral gap in each submanifold to regulate information flow adaptively, allowing for the locally adjustment of different regions. Herein, we visit the concept of boundary conditions and source terms, and construct a novel **Adaptive Local Exchanger**. Specifically, it performs local bottleneck control in spirit of Robin condition, and preserves the representation distinguishability through the energy sustenance of source terms, thereby generating informative representations with theoretical guarantees.

Based on the aforementioned construction, we propose a novel **Heat-Exchanger with Adaptive Locality for Hypergraph Neural Networks**, referred to as (**HeatHGNN**), as a simple yet effective neural implementation of the new message passing mechanism with Adaptive Local Exchanger. As shown in Figure 1, our approach adaptively controls bottlenecks locally and retards the Dirichlet Energy from vanishing, so that fine-grained local adjustments translate into effective global propagation. Furthermore, reflecting the two-stage message passing, we model nodes and hyperedges as coupled systems with bidirectional interactions, capturing richer high-order dependencies. Finally, we derive a clean node-wise formulation of the updating rule, allowing for a satisfactory scalability




Figure 1: (A). Traditional dynamics in HGNN. (B). Illustration of our adaptive heat exchanger; bd denotes boundary nodes. with the linear computational complexity with respect to the graph size.

In summary, key contributions of our work are listed as follows,

- We rethink the heterophily of hypergraph representation learning, and introduce a unified framework of Riemannian heat flow to an interesting problem of **heterophily-agnostic hypergraph neural network design**.
- The key insight is that we discover the inherent connection between oversquashing and hypergraph bottlenecks through the Riemannian geometry perspective, so that we develop an **adaptive local exchanger** for the new message-passing mechanism that is agnostic to both homophilic and heterophilic hypergraphs.
- Based on the proposed mechanism, we propose a novel **HeatHGNN**, a node-hyperedge bidirectional system, which achieves heterophily-agnostic message passing with linear complexity in the number of nodes plus edges and shows superior performance on a variety of real-world homophilic and heterophilic hypergraphs.

2 PRELIMINARY

This section introduces the basic concepts of hypergraphs and hypergraph learning, defines the hypergraph Laplacian, and presents two quantitative metrics: the Cheeger-constant-based bottleneck indicator for oversquashing and the Dirichlet-energy measure of global inter-node feature disparity for oversmoothing. Important notations are summarized in Appendix A.

Hypergraph. A hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{H})$ generalizes the concept of a simple graph by allowing each hyperedge to connect an arbitrary subset of vertices. Here, $\mathcal{V}$ denotes the vertex set, and $\mathcal{E}$ denotes the hyperedge set. The connectivity structure can be represented by an incidence matrix $\mathcal{H} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$, where $\mathcal{H}_{v,e} = 1$ if $v \in e$, and $\mathcal{H}_{v,e} = 0$ otherwise. Let $\mathcal{D}_{\mathcal{V}} \in \mathbb{R}_{\geq 0}^{|\mathcal{V}| \times |\mathcal{V}|}$ and $\mathcal{D}_{\mathcal{E}} \in \mathbb{R}_{\geq 0}^{|\mathcal{E}| \times |\mathcal{E}|}$ be the diagonal degree matrices of vertices and hyperedges, where $\mathcal{D}_{v,i}^{\mathcal{V}}$ counts hyperedges containing v_i , and $\mathcal{D}_{k,k}^{\mathcal{E}}$ counts vertices in e_k . The normalized hypergraph Laplacian is commonly defined as $\mathcal{L} = \mathcal{I} - \mathcal{D}_{\mathcal{V}}^{-\frac{1}{2}} \mathcal{H} \mathcal{D}_{\mathcal{E}}^{-1} \mathcal{H}^T \mathcal{D}_{\mathcal{V}}^{-\frac{1}{2}}$ [18]. We say that $\mathcal{G}$ is a r -graph (or r -uniform) if every hyperedge contains exactly r vertices.




Figure 2: Overall architecture of the proposed heat kernel based hypergraph neural network.

Cheeger Constant. For a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{H})$, the Cheeger constant $\phi(\mathcal{G})$ [47] is defined through the following construction: for any nonempty proper vertex subset S , its volume $\text{vol}(S)$ is the sum of degrees $\sum_{v \in S} d(v)$, while its edge boundary $\partial_e S$ consists of hyperedges connecting S and $\mathcal{V} \setminus S$, counted with multiplicity: $|e \cap S| \cdot |e \cap (\mathcal{V} \setminus S)|$ for each $e \in \partial_e S$. The boundary size $|\partial_e S|$ sums these multiplicities over all hyperedges. The Cheeger constant $\phi(\mathcal{G})$ is then defined as:

$$\phi(\mathcal{G}) = \min \left\{ \frac{|\partial_e S|}{\text{vol}(S)} \mid \emptyset \neq S \subset \mathcal{V}, \text{vol}(S) \leq \text{vol}(\mathcal{V} \setminus S) \right\}, \quad (1)$$

This value measures the minimal cross-boundary connectivity strength between partitions of the hypergraph, serving as a key quantitative indicator of bottlenecks in information propagation. Furthermore, we define the boundary vertex set relative to S as

$$\partial_v S := \{v \in S \mid \exists e \in \mathcal{E} : v \in e \wedge e \cap (\mathcal{V} \setminus S) \neq \emptyset\}. \quad (2)$$

This set contains all vertices in S that belong to at least one hyperedge crossing the boundary between S and its complement. Correspondingly, we define the set of internal hyperedges induced by S as $S_e := \{e \in \mathcal{E} \mid e \subseteq S\}$.

Dirichlet Energy. Here we define the Dirichlet energy on hypergraphs following [13], which we use to quantify the smoothness of node embeddings with respect to the hypergraph structure. Let $X_{\mathcal{V}}^{(k)} \in \mathbb{R}^{|\mathcal{V}| \times d}$ be the node embedding matrix at the k -th HGNN layer. The normalized Dirichlet energy is

$$E(X_{\mathcal{V}}^{(k)}) = \frac{1}{2} \sum_{e \in \mathcal{E}} \frac{1}{d_e} \sum_{u, v \in e} \|d_v^{-1/2} x_v^{(k)} - d_u^{-1/2} x_u^{(k)}\|_2^2. \quad (3)$$

Intuitively, $E(X_{\mathcal{V}}^{(k)})$ measures the variation of embeddings across the nodes incident to each hyperedge: smaller energy indicates smoother (more homogeneous) features, whereas larger energy reflects stronger inter-node disparities. Consequently, this quantity serves as a natural proxy for tracking smoothing effects across layers in HGNNs.

3 THEORETICAL INSIGHTS

In heterophilous hypergraphs, informative signals must cross mixed-label regions and narrow structural bottlenecks, while increasing

depth risks unwanted homogenization. To address both, we first link Cheeger Constant to the spectral gap, quantifying the size of bottleneck. Next, by casting hypergraph propagation as a heat equation, we introduce the Robin condition to adjust bottleneck size and add source terms to prevent the vanishing of nodes disparity. Altogether, this act as an adaptive local exchanger and yields a unified, spectral view: a larger λ improves cross-region heat exchange, yet also accelerates global smoothing—revealing a principled trade-off, all theoretically guaranteed by our analysis.

3.1 Connection between Oversquashing and Hypergraph Bottleneck

We consider the manifold heat flow (the functional space over Riemannian manifold), and discover that oversquashing is inherently related to hypergraph bottleneck.

3.1.1 Bottleneck Size to Spectral Gap. In a hypergraph, structural bottlenecks arise when a small set of vertices or hyperedges forms a cut that separates the vertex set into two large parts. Analogous to a narrow tube regulating heat transfer between two thermodynamic systems, any information exchange between the two sides during message passing must traverse this limited interface, thereby restricting the overall flow. The severity of the bottleneck can be quantified by the hypergraph Cheeger Constant which measures the sparsity of the smallest cut relative to the volumes of the two sides. A smaller $\phi(\mathcal{G})$ indicates a narrower “communication channel” between large regions of the hypergraph. Consequently, we have the following theorem, which establishes the relationship between $\phi(\mathcal{G})$ and the spectral gap λ .

THEOREM 3.1 (CHEEGER INEQUALITY [47]). *In a connected r -uniform hypergraph $\mathcal{G}$, the severity of structural bottlenecks is captured by the hypergraph Cheeger constant $\phi(\mathcal{G})$. The relationship between $\phi(\mathcal{G})$ and the spectrum of the normalized hypergraph Laplacian $\mathcal{L}$ is given by the Cheeger Inequality for r -uniform hypergraphs:*

$$\frac{\phi^2(\mathcal{G})}{2(r-1)^2} \leq \lambda_1 \leq \frac{2}{r-1} \phi(\mathcal{G}), \quad (4)$$

where λ_1 denotes the first non-zero eigenvalue of $\mathcal{L}$.

Thus, a small λ_1 indicates a severe bottleneck. In multi-hop aggregation schemes, these bottlenecks directly contribute to the oversquashing phenomenon.

3.1.2 Heat Equation View of HGNNs. Revisit feature propagation on a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{H})$ via the normalized hypergraph Laplacian $\mathcal{L} \in \mathbb{R}^{n \times n}$. A typical layer updates is given below.

$$X^{(k)} = \sigma(\mathcal{D}_V^{-1/2} \mathcal{H} \mathcal{D}_E^{-1} \mathcal{H}^\top \mathcal{D}_V^{-1/2} X^{(k-1)} \Theta^{(k-1)}). \quad (5)$$

Given that the spectral gap λ of the hypergraph Laplacian relates to information bottlenecks and oversquashing, it is natural to draw an analogy between feature propagation and heat flow on Riemannian manifold governed by the Laplace-Beltrami operator [41]. Within the k -th temporal layer interval $[t_k, t_{k+1})$, the node feature $u_k(t)$ evolves according to the heat equation

$$\partial_t u_k(t, \mathbf{x}) = -\mathcal{L} u_k(t, \mathbf{x}), \quad u_k(t_k, \mathbf{x}) = Z_k, \quad (6)$$

where Z_k is the input feature at the beginning of the k -th layer. The closed-form solution is

$$u_k(t) = e^{-(t-t_k)\mathcal{L}} Z_k, \quad (t \geq t_k), \quad (7)$$

For sufficiently small time step $\Delta t = t_{k+1} - t_k$, the heat kernel $K(t) := e^{-t\mathcal{L}}$ admits the first-order approximation $e^{-\Delta t \mathcal{L}} \approx I - \Delta t \mathcal{L}$, yielding the discrete propagation rule, where $\mathcal{G}_{\theta_k} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a learnable mapping (e.g., a neural network).

$$u_k = (I - \Delta t \mathcal{L}) Z_k, \quad Z_k = \mathcal{G}_{\theta_k}(u_{k-1}). \quad (8)$$

This form unifies hypergraph neural networks, extending classical manifold heat flow theory to discrete hypergraph settings.

3.2 Adaptive Local Exchanger

3.2.1 Bottleneck Control via Robin Condition. Building upon the previously defined boundary node set $\partial_v \mathcal{S}$, we view inter subhypergraph information exchange as heat transfer between thermodynamic systems. The boundary nodes act as the interface (a “tube”) between systems; their propagation paradigm directly determines the effective tube width, i.e., the bottleneck size. Thus, by imposing boundary conditions, we can control the bottleneck and regulate information flux at the boundary as needed. Specifically, consider the heat equation on the hypergraph induced by node subset $\mathcal{S}$:

$$\begin{cases} \partial_t u_{\mathcal{S}}(t, \mathbf{x}) = -\mathcal{L}_{\mathcal{S}} u_{\mathcal{S}}(t, \mathbf{x}), & \mathbf{x} \in \mathcal{S}, \\ \alpha u(\mathbf{x}) + \beta \frac{\partial u}{\partial n}(\mathbf{x}) = 0, & \mathbf{x} \in \partial_v \mathcal{S}, \end{cases} \quad (9)$$

where $\mathcal{L}_{\mathcal{S}}$ is the Laplacian restricted to $\mathcal{S}$. The boundary conditions capture realistic propagation dynamics—particularly at the interface between a subset $\mathcal{S}$ and its complement $\mathcal{V} \setminus \mathcal{S}$. A Robin condition [16] combines a direct constraint α on the feature value at the boundary with a constraint β on the directional derivative $\frac{\partial u}{\partial n}$ along the outward normal. Here, $\frac{\partial u}{\partial n}$ represents the rate at which information flows across the boundary. By tuning the coefficients α and β , this condition interpolates between the Dirichlet ($\beta \rightarrow 0$) and Neumann ($\alpha \rightarrow 0$) extremes, providing a flexible mechanism to control feature propagation through boundary nodes and mitigate oversquashing in bottleneck regions.

3.2.2 Energy Injection via Source Terms. In the presence of an external source term, the diffusion dynamics become $\partial_t u_k(t, \mathbf{x}) = -\mathcal{L} u_k(t, \mathbf{x}) + f(t, \mathbf{x})$. In this setting, the Dirichlet energy no longer decays exponentially as in the homogeneous case. Instead, the forcing term continually injects energy into the system, mitigating the dissipation caused by the Laplacian.

3.2.3 Unified Heat Exchanger. To simultaneously address oversmoothing and oversquashing, we aim to unify the heat equation with a source term and the Robin condition into a single formulation. Recall the problem $\partial_t u(t, \mathbf{x}) = -\mathcal{L} u(t, \mathbf{x})$ with boundary condition $h(u(t, \mathbf{x})) = \gamma(t, \mathbf{x})$. There exists a function $w(t, \mathbf{x})$, such that $h \circ w = \gamma$ on $\partial_v \mathcal{S}$. By introducing the change of variables $v = u - w$, the nonhomogeneous boundary condition $h(u(t, \mathbf{x})) = \gamma(t, \mathbf{x})$, $\mathbf{x} \in \partial_v \mathcal{S}$ can be equivalently transformed into a homogeneous form B.4:

$$\begin{cases} \partial_t v(t, \mathbf{x}) = -\mathcal{L} v(t, \mathbf{x}) + \Gamma(w), & \mathbf{x} \in \mathcal{S}, t > 0, \\ h(\mathbf{x}) = 0, & \mathbf{x} \in \partial_v \mathcal{S}, t > 0, \end{cases} \quad (10)$$

where $\Gamma(w) = -\mathcal{L} w(t, \mathbf{x}) - \partial_t w(t, \mathbf{x})$ and $v(0, \mathbf{x}) = Z(\mathbf{x}) - w(0, \mathbf{x})$. **Intuitive Explanation.** In our adaptive local heat exchanger, the Robin condition regulates cross-region heat flow like a tunable valve, while the source term serves as a booster that injects energy into the system. Together, they tighten or loosen local bottlenecks and slow the decay of Dirichlet energy, so that fine-grained local adjustments translate into effective global propagation.

3.3 Guarantees on Heterophily-agnostic Message Passing

3.3.1 Dirichlet Energy to Spectral Gap. From a physical view, the second law of thermodynamics states that in an isolated thermodynamic system, heat distribution irreversibly evolves toward equilibrium. This process mirrors the mechanism of oversmoothing in deep HGNNs, where layer-wise diffusion gradually homogenizes node features and rapidly decays Dirichlet energy, flattening feature differences across hyperedges.

THEOREM 3.2 (SPECTRAL GAP AND DIRICHLET ENERGY, B.2). *Let $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ be the eigenvalues of the hypergraph Laplacian $\mathcal{L}$, then the Dirichlet Energy satisfies*

$$E(u_k(t)) \leq e^{-2\lambda_1(t-t_k)} E(Z_k). \quad (11)$$

In particular, the Dirichlet energy decays at least exponentially with rate $2\lambda_1$.

Consistent with Theorem 3.1, oversmoothing can also be characterized by the spectral gap: a larger λ_1 accelerates the exponential decay of Dirichlet energy, thereby intensifying smoothing. At the same time, a larger λ helps mitigate oversquashing. Taken together, these effects provide a theoretical validation of the intrinsic trade-off between OSQ and OSM.

3.3.2 Guarantees of Our Constructions. These boundary conditions modify the spectral properties of the restricted Laplacian $\mathcal{L}_{\mathcal{S}}$, where locally adjusting λ controls the amount of messages that flow out of the bottlenecks. Specifically, Dirichlet fixes boundary values and maximizes exchange, Neumann blocks exchange with zero normal flux, while Robin provides a tunable finite-conductance

interface that interpolates between them. In the presence of an external source term, the diffusion dynamics become $\partial_t u_k(t, \mathbf{x}) = -\mathcal{L} u_k(t, \mathbf{x}) + f(t, \mathbf{x})$. In this setting, the Dirichlet energy no longer decays exponentially as in the homogeneous case. Instead, the forcing term continually injects energy into the system, mitigating the dissipation caused by the Laplacian.

THEOREM 3.3 (FLEXIBILITY OF THE ROBIN CONDITION, B.1). *Consider a thin tube-like region S of length L , whose cross-section at position $s \in [0, L]$ is a small ball of radius $\varepsilon(s)$. The boundary of this region is given by $\partial S = \{0, L\} \times \partial \mathcal{B}_{\varepsilon(s)}$. Let λ_D , λ_N , and λ_R denote the spectral gaps of the Laplace operator under Dirichlet, Neumann, and Robin boundary conditions, respectively. When the radius is constant, i.e., $\varepsilon(s) = \varepsilon_0$, the spectral gaps satisfy the ordering $\lambda_D \geq \lambda_R \geq \lambda_N$.*

Furthermore, the source term selectively preserves initial information and amplifies disparity; together with the Robin boundary, it yields a heterophily-agnostic propagation mechanism.

THEOREM 3.4 (MODE PRESERVATION UNDER SOURCE ENERGY INJECTION, B.3). *Consider the diffusion equation with source term $\partial_t u_k(t, \mathbf{x}) = -\mathcal{L} u_k(t, \mathbf{x}) + f(t, \mathbf{x})$. Expanding in the Laplacian eigenbasis gives $u(t) = \sum_i \hat{u}_i(t) \psi_i$, $f(t) = \sum_i \hat{f}_i(t) \psi_i$ with coefficients $\hat{u}_i(t) = \langle u(t), \psi_i \rangle$, $\hat{f}_i(t) = \langle f(t), \psi_i \rangle$, and each mode satisfies $\dot{\hat{u}}_i(t) = -\lambda_i \hat{u}_i(t) + \hat{f}_i(t)$. If the source persistently excites mode i over a time step $\Delta t = t_{k+1} - t_k$, i.e., $\hat{f}_i(s) \geq \eta_i > 0$ for all $s \in [t_k, t_{k+1}]$, then its amplitude is bounded below by*

$$|\hat{u}_i(t_{k+1})| \geq \frac{\eta_i}{\lambda_i} (1 - e^{-\lambda_i \Delta t}) - e^{-\lambda_i \Delta t} |\hat{u}_i(t_k)|. \quad (12)$$

showing that persistent energy injection prevents the mode from vanishing.

4 HEALHGNN: A UNIFIED APPROACH

We first discretize the hypergraph heat diffusion system with boundary conditions and source terms to a layer-wise update form. Let $X_V(k)$ and $X_E(k)$ denote the vertex and hyperedge representations at the k -th layer, respectively,

$$\underbrace{\begin{bmatrix} X_V(k+1) \\ X_E(k+1) \end{bmatrix}}_{\text{updated state}} = \underbrace{f\left(\begin{bmatrix} X_V(k) \\ X_E(k) \end{bmatrix}\right)}_{\text{diffusion process}}, \quad \underbrace{\begin{bmatrix} X_V(0) \\ X_E(0) \end{bmatrix}}_{\text{initial state}} = \begin{bmatrix} Z_V \\ Z_E \end{bmatrix}. \quad (13)$$

4.1 Node-Hyperedge Bidirectional System

Inspired by the continuous setting, we impose the nonhomogeneous Robin condition on hypergraph in a form directly mirroring its PDE counterpart:

$$\alpha_i \mathbf{x}_i + \beta_i \partial_n^{\text{disc}} \mathbf{x}_i = \gamma_i, \quad i \in \partial_v S, \quad (14)$$

where $\alpha_i = \alpha(\mathbf{x}_i)$, $\beta_i = \beta(\mathbf{x}_i)$, and $\gamma_i = \gamma(\mathbf{x}_i)$ are coefficient functions depending on the current node states. The discrete normal derivative $\partial_n^{\text{disc}} \mathbf{x}_i$ is defined as $\partial_n^{\text{disc}} \mathbf{x}_i := \sum_{j \sim i, j \in \mathcal{V} \setminus \partial_v S} (\mathbf{x}_i - \mathbf{x}_j)$, representing the net feature flux from boundary node i to its interior neighbors. Subject to Eq. (14) on $\partial_v S$, the heat equation in Eq. (10) governs the evolution on the interior nodes, ensuring a nontrivial steady state. The update of node embeddings on hypergraphs can be written in block-matrix form as

$$\begin{cases} \text{Interior:} & \mathcal{L}^{S \rightarrow S} X_S + \mathcal{L}^{\partial S \rightarrow S} X_{\partial S} = F_S, \\ \text{Boundary:} & \text{diag}(\alpha) X_{\partial S} + \text{diag}(\beta) \mathcal{L}^{S \rightarrow \partial S} X_S = \Gamma_{\partial S}. \end{cases} \quad (15)$$

Here, X_S and $X_{\partial S}$ denote the feature matrix of interior and boundary nodes, respectively. The matrices $\mathcal{L}^{S \rightarrow S}$ and $\mathcal{L}^{\partial S \rightarrow S}$ correspond to the discrete Laplacian restricted to interior and boundary contributions. For nodes, the Laplacian is defined as $\mathcal{L}_V = I - \mathcal{D}_V^{-1/2} \mathcal{H} \mathcal{D}_E^{-1} \mathcal{H}^\top \mathcal{D}_V^{-1/2}$. Analogously, for the two-stage propagation process involving hyperedge features, we define the hyperedge Laplacian $\mathcal{L}_E = I - \mathcal{D}_E^{-1/2} \mathcal{H}^\top \mathcal{D}_V^{-1} \mathcal{H} \mathcal{D}_E^{-1/2}$. With this definition, the same block-matrix update can be applied to hyperedges. This unified formulation Eq. (15) describes feature diffusion for both nodes and hyperedges, incorporating interior-boundary interactions and weighted contributions through $\text{diag}(\alpha)$ and $\text{diag}(\beta)$.

Building upon the separate propagation dynamics for nodes and hyperedges, we introduce bidirectional coupling between the two domains via the propagation operators $\mathcal{P}_{V \rightarrow E} = \mathcal{D}_E^{-1/2} \mathcal{H}^\top \mathcal{D}_V^{-1/2}$ and $\mathcal{P}_{E \rightarrow V} = \mathcal{D}_V^{-1/2} \mathcal{H} \mathcal{D}_E^{-1/2}$. Specifically, interior and boundary nodes in one domain receive feature contributions from the other domain, modulated by the coupling coefficients λ and μ . This design allows information to propagate not only within each domain through the Laplacian operators but also across domains, thereby capturing richer topological dependencies and enhancing interactions between vertex-level and hyperedge-level representations.

4.2 A Clean Node-wise Formulation

Formally, the coupled update rules can be expressed as a compact block-matrix form, which is convenient for unified analysis and computation.

$$A \begin{bmatrix} X_V \\ X_E \end{bmatrix} = \begin{bmatrix} \Gamma_V \\ \Gamma_E \end{bmatrix}, \quad X = \begin{bmatrix} X_S \\ X_{\partial S} \end{bmatrix}, \quad \Gamma = \begin{bmatrix} F_S \\ \Gamma_{\partial S} \end{bmatrix}. \quad (16)$$

$$A = \begin{bmatrix} \mathcal{L}_V^{S \rightarrow S} & \mathcal{L}_V^{\partial S \rightarrow S} & -\lambda \mathcal{P}_{E \rightarrow V}^{S \rightarrow S} & -\lambda \mathcal{P}_{E \rightarrow V}^{\partial S \rightarrow S} \\ \text{diag}(\beta_v) \mathcal{L}_V^{S \rightarrow \partial S} & \text{diag}(\alpha_v) & -\lambda \mathcal{P}_{E \rightarrow V}^{\partial S \rightarrow \partial S} & -\lambda \mathcal{P}_{E \rightarrow V}^{\partial S \rightarrow \partial S} \\ -\mu \mathcal{P}_{V \rightarrow E}^{S \rightarrow S} & -\mu \mathcal{P}_{V \rightarrow E}^{\partial S \rightarrow S} & \mathcal{L}_E^{S \rightarrow S} & \mathcal{L}_E^{\partial S \rightarrow S} \\ -\mu \mathcal{P}_{V \rightarrow E}^{S \rightarrow \partial S} & -\mu \mathcal{P}_{V \rightarrow E}^{\partial S \rightarrow \partial S} & \text{diag}(\beta_e) \mathcal{L}_E^{S \rightarrow \partial S} & \text{diag}(\alpha_e) \end{bmatrix} \quad (17)$$

Let $\mathbf{D}$, $\mathbf{U}$, and $\mathbf{V}$ denote the diagonal, upper, and lower blocks of the system matrix $\mathbf{A}$, respectively. Following the Jacobi iteration scheme, we parameterize the problem-dependent coefficients via learnable mappings

$$\text{diag}(\alpha) = \alpha(\mathbf{X}^{(k)}; \theta_\alpha), \quad \text{diag}(\beta) = \beta(\mathbf{X}^{(k)}; \theta_\beta), \quad \Gamma = \gamma(\mathbf{X}^{(k)}; \theta_\gamma). \quad (18)$$

The update rule is then augmented with a nonlinearity $\sigma(\cdot)$ and a learnable transformation $\varphi^{(k+1)}(\cdot)$ implemented as a multi-layer perceptron (MLP), yielding

$$\mathbf{X}^{(k+1)} = \sigma\left(\mathbf{D}^{-1}(\mathbf{U} + \mathbf{V}) \varphi^{(k+1)}(\mathbf{X}^{(k)})\right) + \mathbf{D}^{-1} \Gamma. \quad (19)$$

This formulation inherits the convergence structure of the Jacobi method while allowing adaptive, data-driven refinement through the learnable parameters θ_α , θ_β , θ_γ and $\varphi^{(k+1)}$. Let $S_{ij} := \sum_{e: i, j \in e} \frac{1}{d_e}$ denote the total inverse degree of hyperedges shared by nodes i and j and $p_i = \frac{\beta_v}{\alpha_v}$. We define the indicator function $I_i := \mathbb{I}(i \in S)$, which equals 1 if node $i \in S$ and 0 otherwise. The feature update

for node i at iteration $k + 1$ is given by

$$x_i^{(k+1)} = \underbrace{\sum_{j: j \sim i} \mathcal{P}_{\mathcal{V} \rightarrow \mathcal{V}}(i, j) x_j^{(k)}}_{\text{node-to-node aggregation}} + \underbrace{\sum_{e: e \sim i} \mathcal{P}_{\mathcal{E} \rightarrow \mathcal{V}}(i, e) x_e^{(k)}}_{\text{edge-to-node aggregation}} + \gamma_i, \quad (20)$$

where $j \sim i$ denotes that i and j share at least one hyperedge, and $e \sim i$ denotes that node i belongs to hyperedge e . Specifically, the node-to-node propagation operator is defined as

$$\mathcal{P}_{\mathcal{V} \rightarrow \mathcal{V}}(i, j) = \frac{1}{\sqrt{d_i d_j}} S_{ij} \left(I_i I_j + I_i + p_i (1 - I_i) I_j \right), \quad (21)$$

The edge-to-node propagation operator is

$$\mathcal{P}_{\mathcal{E} \rightarrow \mathcal{V}}(i, e) = \mu \frac{1}{\sqrt{d_i d_e}} \left(I_i + \frac{1}{\alpha_v} (1 - I_i) \right), \quad (22)$$

where μ is a scaling coefficient, and α_v controls the relative influence from hyperedges to boundary versus interior nodes.

For hyperedges, the update follows a similar form with appropriately defined edge-to-edge and node-to-edge propagation operators. We omit the detailed expressions here for brevity; the full derivation and definitions are provided in B.5, and the complete architecture of HealHGNN is shown in Figure 2.

5 EXPERIMENTS

Objectives. We evaluate our proposed method on a variety of hypergraph learning tasks and compare it with strong baselines. Specifically, we aim to address the following questions: **RQ1:** Can our model effectively handle both homogeneous and heterogeneous hypergraphs, as well as scale to large datasets? **RQ2:** Does our approach alleviate oversquashing in the hypergraph transfer setting and on the Long-Range Graph Benchmark(LRGB), enabling efficient long-range information propagation? **RQ3:** Can our method mitigate oversmoothing when stacking deep layers, thereby preserving expressive node and hyperedge representations? **RQ4:** How critical are the proposed components to the overall effectiveness of our method? Codes are available in <https://github.com/Mingzhang21/HealHGNN>. Further implementation notes are given in Appendix D.

Experimental Settings. We conduct experiments on ten publicly available datasets spanning various applications: Cora, Cora-CA, DBLP-CA, CiteSeer, and PubMed [49], as well as NTU2012 [5], House [8], Congress, Senate [19], and Walmart [1]. Detailed descriptions of these baselines, along with the datasets, can be found in Appendix C. All experiments are conducted on a machine equipped with an NVIDIA GeForce RTX 4090 GPU (24 GB memory), a 16-core Intel(R) Xeon(R) Gold 6430 CPU, and 120 GB RAM.

5.1 Hypergraph Node Classification (RQ1)

Setup. We evaluate the performance of our model on the node classification task using eight real-world hypergraph datasets, and adopt classification accuracy (ACC) as the evaluation metric. These datasets span diverse domains, scales, and CE homophily levels, and are widely used benchmarks in hypergraph learning. To ensure a fair comparison with baselines, we follow the experimental protocol of [43], where each dataset is randomly split into 50% training, 25% validation, and 25% test nodes. Each case is repeated 10 times with different random splits, and we report the mean accuracy along

with the standard deviation.

Results. Our model achieves competitive performance across all real-world datasets—matching strong baselines on homophilic hypergraphs and delivering clear gains on heterophilic ones. Notably, on the Senate dataset we outperform the second-best method by 5%, underscoring the benefits of our heterophily-agnostic design.

5.2 Synthetic Hypergraph under Heterophily

Setup. The oversmoothing problem can hurt the performance of hypergraph networks, especially on heterophilic graphs. To investigate this, we evaluate hypergraph networks under varying heterophily. Following [48], we use synthetic hypergraphs generated with a contextual stochastic block model. Each hypergraph contains 5000 nodes divided into two classes and 1000 hyperedges, with the heterophily level controlled by adjusting the proportion of nodes from class 0.

Results. As observed in Table 2, when the heterophily level is below three, our method performs slightly better than the other models. However, as the heterophily level increases, our approach demonstrates a significant improvement. This trend is reflected in Table 1, where our method outperforms the baselines in heterophilic hypergraphs, showcasing its ability to handle challenging heterophilic environments and mitigate the oversmoothing effect.

5.3 Hypergraph Transfer (RQ2)

Setup. Inspired by [11][48], we consider a hypergraph transfer task, where the objective is to transfer a label from a source node to a target node across hyperedges. Nodes are initialized with random-valued features, and we assign labels “1” and “0” to source and target nodes, respectively. We design four types of hypergraph distributions (see Figure 3 for a visual illustration): HyperEdgeSingle (HES(n)) uses a single hyperedge of size n ; HyperEdgePath (HEP(n, m)) places the source in an n -node hyperedge and connects it to the target via a length- m chain of degree-2 hyperedges; HyperEdgeRing (HER(n, m)) forms a ring of n -node hyperedges, each sharing one node with the next, with the source–target distance controlled by m ; and HyperEdgeDumbbell(n, m) consists of two HER(n, m) rings linked by a bridge of four degree-2 hyperedges. Importantly, the number of model layers is set equal to the required transfer distance.

Results. Figure 5 reports the performance of different models on the hypergraph transfer tasks. Overall, HGNN[18] exhibits clear limitations: due to its reliance on standard hypergraph Laplacian propagation, node features tend to become homogenized over long distances, making it difficult to effectively transfer the label information from the source node to the target node. Similarly, the baseline methods struggle as the transfer distance increases, since their varying degrees of feature compression further hinder accurate long-range information transfer. In contrast, HealHGNN demonstrates consistently strong performance. On HES tasks, it maintains high accuracy, as nodes within the same hyperedge can still preserve discriminative information. More importantly, on HEP, HER, and HED, where the structures are more complex and the required propagation distance is longer (as illustrated in Fig 3), thereby substantially increasing task difficulty, HealHGNN consistently outperforms all other models. This highlights its robustness and effectiveness in


Table 1: Node classification on homophilic and heterophilic hypergraphs. Top three models are coloured by First, Second, Third.

Type	Homophilic hypergraphs				Heterophilic hypergraphs			
Dataset	DBLP-CA	Citeseer	Pubmed	NTU2012	Congress	House	Senate	Walmart
HGNN [18]	91.03 $\pm$ 0.20	62.45 $\pm$ 1.16	86.44 $\pm$ 0.44	87.72 $\pm$ 1.35	91.26 $\pm$ 1.15	61.39 $\pm$ 2.96	48.59 $\pm$ 4.52	62.00 $\pm$ 0.24
HyperGCN [49]	89.98 $\pm$ 0.43	72.76 $\pm$ 1.12	82.84 $\pm$ 0.67	86.36 $\pm$ 1.86	87.43 $\pm$ 1.46	60.13 $\pm$ 1.76	50.13 $\pm$ 3.21	48.74 $\pm$ 2.81
HNHN [12]	86.78 $\pm$ 0.29	72.64 $\pm$ 1.57	84.21 $\pm$ 0.64	86.13 $\pm$ 0.76	90.89 $\pm$ 0.31	67.80 $\pm$ 2.59	50.93 $\pm$ 6.33	51.18 $\pm$ 0.35
HCHA [3]	90.92 $\pm$ 0.22	72.42 $\pm$ 1.42	84.56 $\pm$ 0.31	87.12 $\pm$ 1.23	90.43 $\pm$ 1.20	61.36 $\pm$ 2.53	48.62 $\pm$ 4.41	62.35 $\pm$ 0.26
UniGCNII [21]	91.69 $\pm$ 0.19	73.05 $\pm$ 2.21	87.31 $\pm$ 0.45	89.30 $\pm$ 1.33	92.43 $\pm$ 0.71	67.25 $\pm$ 2.57	49.30 $\pm$ 4.25	54.45 $\pm$ 0.37
AllDeepSets [7]	91.27 $\pm$ 0.27	70.83 $\pm$ 1.63	87.04 $\pm$ 0.52	88.69 $\pm$ 1.24	91.80 $\pm$ 1.53	67.82 $\pm$ 2.40	48.17 $\pm$ 5.67	64.55 $\pm$ 0.33
AllSetTransformer [7]	91.53 $\pm$ 0.23	73.08 $\pm$ 1.20	87.66 $\pm$ 0.43	88.43 $\pm$ 1.43	92.16 $\pm$ 1.05	69.83 $\pm$ 5.22	69.33 $\pm$2.20	65.46 $\pm$0.25
Deep-HGNN [6]	91.76 $\pm$ 0.28	74.07 $\pm$ 1.64	87.42 $\pm$ 0.22	87.12 $\pm$ 1.13	92.47 $\pm$ 1.06	75.26 $\pm$1.76	68.39 $\pm$ 4.79	58.64 $\pm$ 0.74
ED-HNN [43]	91.90 $\pm$0.19	73.70 $\pm$ 1.38	87.58 $\pm$ 0.65	88.07 $\pm$ 1.28	93.28 $\pm$0.59	72.45 $\pm$ 2.28	64.79 $\pm$ 5.14	66.79 $\pm$0.41
SheafHyperGNN [13]	91.59 $\pm$ 0.24	74.71 $\pm$ 1.23	87.68 $\pm$ 0.60	-	91.81 $\pm$ 1.60	73.84 $\pm$ 2.30	68.73 $\pm$ 4.68	-
PhenomNN [44]	91.86 $\pm$0.22	74.45 $\pm$ 0.96	78.12 $\pm$ 0.24	85.40 $\pm$ 0.42	88.24 $\pm$ 1.47	70.71 $\pm$ 2.35	67.70 $\pm$ 5.24	63.43 $\pm$ 0.43
FrameHGNN [24]	-	74.72 $\pm$2.10	88.73 $\pm$0.42	89.98 $\pm$2.02	-	72.82 $\pm$ 2.22	67.61 $\pm$ 5.27	-
KHGNN [46]	91.21 $\pm$ 0.37	74.80 $\pm$1.10	88.47 $\pm$0.47	89.60 $\pm$1.64	92.55 $\pm$1.03	74.31 $\pm$2.87	71.23 $\pm$3.94	65.17 $\pm$ 0.79
HealHGNN (Ours)	91.98 $\pm$0.31	75.06 $\pm$1.16	88.80 $\pm$0.31	90.16 $\pm$1.78	93.45 $\pm$0.79	77.18 $\pm$2.35	76.06 $\pm$4.13	68.22 $\pm$0.56

Table 2: Heterophily Experiments

Methods	Heterophily Level						
	1	2	3	4	5	6	7
HGNN	97.8	85.7	78.4	70.5	66.5	64.9	60.4
HyperGCN	85.1	77.1	75.9	72.9	69.5	65.1	60.7
HCHA	98.4	83.8	79.3	74.1	72.1	69.1	64.3
AllDeepSets	99.8	87.8	82.4	77.3	73.4	71.8	68.4
AllSetTransformer	99.9	86.5	83.3	78.5	74.2	72.4	69.2
ED-HNN	99.9	91.3	88.4	85.1	81.7	77.3	76.5
SensHNN	100	93.7	90.1	85.7	82.4	79.4	77.6
HealHGNN	100	94.1	90.8	87.1	84.3	82.2	79.8


capturing both local and global structural dependencies.

**Figure 3: Illustration of the Topological Structures**

5.4 Long-Range Graph Benchmark (RQ2)

Setup. We further examine the capacity of HealHGNN to model long-range dependencies using the Long-Range Graph Benchmark (LRGB) [15]. In this benchmark, we consider two peptide datasets, Peptides-func and Peptides-struct, which are specifically designed to evaluate models on graph tasks requiring the capture of long-range interactions within molecular graphs.

Results. As shown in Figure 4, graph transformer-based methods such as GT [14] and GRIT [26] outperform classical MPNNs including GCN, HGNN [18], and GUMP [27], owing to their ability to leverage global receptive fields. KHGNN [46] extends the receptive field by aggregating information along k -hop shortest paths, but this comes at the cost of computational overhead. In contrast,

**Figure 4: Performance comparison on LRGB tasks.**

HealHGNN achieves superior performance while maintaining linear complexity $O(|V| + |\mathcal{E}|)$. By incorporating boundary awareness, it ensures that essential signals can be transferred efficiently across distant regions of the hypergraph.

5.5 Neural Network Depth (RQ3)

Setup. We compare HealHGNN against a range of representative baselines. To examine the effect of model depth on performance, we conduct experiments on four benchmark datasets: two heterophilic hypergraphs (House and Senate) and two homophilic hypergraphs (Citeseer and Cora). The number of layers is varied across $\{1, 2, 4, 8, 16, 32\}$ to study the relative performance of shallow versus deep architectures. The results are summarized in Figure 6. **Results.** Overall, our method demonstrates stability as the number of layers increases, showing no significant performance fluctuations and even a slight upward trend. As discussed previously, the effect of oversmoothing becomes more severe in heterophilic hypergraphs, amplifying the performance gap in favor of HealHGNN.

5.6 Ablation Study (RQ4)

To evaluate the contribution of each component in HealHGNN, we create several variants by replacing the learnable coefficients γ , β , and hyperparameter μ with fixed constants (γ_0 , β_0 , μ_0). Setting $\gamma_v = 0$, $\beta_v = 0$, or $\mu = 0$ disables external inputs, boundary interactions, or hyperedge modulation, respectively, while $\gamma_v = \beta_v = \mu = 0$ reduces HealHGNN to a standard HGNN.




Figure 5: Information Transfer Performance on the Four Hypergraph Types with Hyperedge Size 4.




Figure 6: Comparison of method performance across varying neural network depths on four datasets.

Table 3: Ablation Study

Model	Citeseer	Cora-CA	House	Senate
HealHGNN	74.56±1.16	84.78±1.57	77.18±2.35	76.06±4.13
γ_0, β_0, μ_0	73.67±0.96	83.36±1.23	75.73±1.79	75.13±2.17
$\gamma_v = 0$	71.15±1.58	82.11±1.34	72.27±2.34	71.19±3.16
$\beta_v = 0$	66.46±1.17	81.41±1.18	65.79±3.78	62.17±0.89
$\mu = 0$	72.46±1.67	82.43±1.07	74.28±2.79	75.23±2.17
$\gamma_v, \beta_v, \mu = 0$	63.04±1.72	79.16±1.96	62.16±2.36	50.89±3.36
HGNN	62.45±1.16	78.64±1.65	61.39±2.96	48.59±4.52

6 RELATED WORK

Hypergraph Representation Learning. Hypergraph learning has emerged as a powerful framework for modeling high-order relations that cannot be adequately captured by pairwise graphs. Foundational models such as HGNN [18] and HyperGCN [49] established the basis for hypergraph neural networks by extending convolutional operators to the hypergraph setting. Subsequent approaches, including HNNH [12] and HCHA [3], enhanced representational capacity by incorporating nonlinear transformations and attention mechanisms. Recent works further advanced the field along theoretical, algorithmic, and efficiency dimensions. ORCHID [9] unified Ollivier–Ricci curvature on hypergraphs, providing new analytical tools. HyperGCL [45] leveraged contrastive learning to improve generalization. To address scalability, LightHGNN [17] distilled HGNNs into lightweight MLPs, while TF-HNN [40] reduced training cost via precomputed structural information.

Oversmoothing and Oversquashing. Deep hypergraph networks face two major challenges—oversmoothing and oversquashing—that constrain their expressive power. To alleviate oversmoothing, various architectures aim to preserve feature diversity during deep propagation. Deep-HGNN [6] employs high-order polynomial filters, while ED-HNN [43] and UniGCNII [21] extend diffusion and message-passing to maintain discriminative representations. Frame-HGNN [24], SheafHyperGNN [13], PhenomNN [44], and HDS [50] further introduce framelet transforms, sheaf Laplacians, or dynamic modeling to reinject node-specific signals and replenish Dirichlet

energy. Methods such as SensHNN [48] and KHGNN [46] enhance message flow via direct interactions between distant nodes. Notably, these rewiring-oriented approaches may exacerbate oversmoothing, as denser pathways tend to homogenize node features, highlighting a key trade-off in hypergraph network design.

Riemannian Geometry and Graph Neural Network. Euclidean space has long been the default of graph representation learning, but Riemannian manifolds have emerged as a compelling alternative [10, 20, 28–38]. The natural extension to *hypergraphs* on manifolds is also gaining traction recently. [51] generates hypergraphs via sparse representation and reformulates orthogonality-constrained optimization as unconstrained Riemannian optimization on the Grassmann manifold. [39] embeds hypergraphs in hyperspherical manifolds to capture higher-order geometric relations through sphere-based message passing.

7 CONCLUSION

This paper presents a heterophily-agnostic hypergraph neural network (HealHGNN) through Riemannian heat flow. We start with the theoretical discovery that long-range dependence is related to hypergraph bottlenecks, and then establish the Adaptive Local Exchanger that locally adjusts subhypergraph bottlenecks in the functional space, creating a new message passing mechanism with theoretical guarantees. Consequently, we formulate HealHGNN that bidirectionally propagates messages from nodes and hyperedges with linear complexity. Extensive experiments across multiple benchmarks demonstrate that HealHGNN consistently improves performance over competitive baselines.

Acknowledgments

This work is supported in part by NSFC under grants 62202164, U25B2029, 62072052 and 62322202, and the National Key Research and Development Program of China (2024YFF0907401). Philip S. Yu is supported in part by NSF under grants III-2106758, and POSE-2346158.

References

- [1] Ilya Amburg, Nate Veldt, and Austin Benson. 2020. Clustering in graphs and hypergraphs with categorical edge labels. In *Proceedings of the web conference 2020*. 706–717.
- [2] Alessia Antelmi, Gennaro Cordasco, Mirko Polato, Vittorio Scarano, Carmine Spagnuolo, and Dingqi Yang. 2023. A survey on hypergraph representation learning. *Comput. Surveys* 56, 1 (2023), 1–38.
- [3] Song Bai, Feihu Zhang, and Philip HS Torr. 2021. Hypergraph convolution and hypergraph attention. *Pattern Recognition* 110 (2021), 107637.
- [4] Peter Buser. 1982. A note on the isoperimetric constant. *Annales Scientifiques de l'École Normale Supérieure* 15, 2 (1982).
- [5] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. 2003. On visual similarity based 3D model retrieval. In *Computer graphics forum*, Vol. 22. Wiley Online Library, 223–232.
- [6] Guanxi Chen, Jiying Zhang, Xi Xiao, and Yang Li. 2022. Preventing oversmoothing for hypergraph neural networks. *arXiv preprint arXiv:2203.17159* (2022).
- [7] Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenkovic. 2021. You are allset: A multiset function framework for hypergraph neural networks. *arXiv preprint arXiv:2106.13264* (2021).
- [8] Philip S Chodrow, Nate Veldt, and Austin R Benson. 2021. Generative hypergraph clustering: From blockmodels to modularity. *Science Advances* 7, 28 (2021), eabh1303.
- [9] Corinna Coupette, Sebastian Dalleiger, and Bastian Rieck. 2022. Ollivier-ricci curvature for hypergraphs: A unified framework. *arXiv preprint arXiv:2210.12048* (2022).
- [10] Haitz Sáez de Ocáriz Borde, Anees Kazi, Federico Barbero, and Pietro Liò. 2023. Latent Graph Inference using Product Manifolds. In *Proceedings of the Eleventh International Conference on Learning Representations*. OpenReview.net.
- [11] Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, and Michael M Bronstein. 2023. On over-squashing in message passing neural networks: The impact of width, depth, and topology. In *International conference on machine learning*. PMLR, 7865–7885.
- [12] Yihe Dong, Will Sawin, and Yoshua Bengio. 2020. Hnbn: Hypergraph networks with hyperedge neurons. *arXiv preprint arXiv:2006.12278* (2020).
- [13] Iulia Duta, Giulia Cassarà, Fabrizio Silvestri, and Pietro Liò. 2023. Sheaf hypergraph networks. *Advances in Neural Information Processing Systems* 36 (2023), 12087–12099.
- [14] Vijay Prakash Dwivedi and Xavier Bresson. 2020. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699* (2020).
- [15] Vijay Prakash Dwivedi, Ladislav Rampásek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. 2022. Long range graph benchmark. *Advances in Neural Information Processing Systems* 35 (2022), 22326–22340.
- [16] Ofodike A Ezekoye. 2016. Conduction of heat in solids. In *SFPE handbook of fire protection engineering*. Springer, 25–52.
- [17] Yifan Feng, Yihe Luo, Shihui Ying, and Yue Gao. 2024. LightHGNN: Distilling hypergraph neural networks into MLPs for 100x faster inference. In *The Twelfth International Conference on Learning Representations*.
- [18] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. 2019. Hypergraph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 3558–3565.
- [19] James H Fowler. 2006. Legislative cosponsorship networks in the US House and Senate. *Social networks* 28, 4 (2006), 454–465.
- [20] Octavian-Eugen Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic Neural Networks. In *Advances in Neural Information Processing Systems* 31. 5350–5360.
- [21] Jing Huang and Jie Yang. 2021. Unignn: a unified framework for graph and hypergraph neural networks. *arXiv preprint arXiv:2105.00956* (2021).
- [22] Sunwoo Kim, Soo Yong Lee, Yue Gao, Alessia Antelmi, Mirko Polato, and Kijung Shin. 2024. A survey on hypergraph neural networks: an in-depth and step-by-step guide. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 6534–6544.
- [23] TN Kipf. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [24] Ming Li, Yujie Fang, Yi Wang, Han Feng, Yongchun Gu, Lu Bai, and Pietro Lio. 2025. Deep hypergraph neural networks with tight framelets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 18385–18392.
- [25] Ming Li, Yongchun Gu, Yi Wang, Yujie Fang, Lu Bai, Xiaosheng Zhuang, and Pietro Lio. 2025. When hypergraph meets heterophily: New benchmark datasets and baseline. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 18377–18384.
- [26] Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K Dokania, Mark Coates, Philip Torr, and Ser-Nam Lim. 2023. Graph inductive biases in transformers without message passing. In *International Conference on Machine Learning*. PMLR, 23321–23337.
- [27] Haiquan Qiu, Yatao Bian, and Quanming Yao. 2024. Graph unitary message passing. *arXiv preprint arXiv:2403.11199* (2024).
- [28] Li Sun, Zhenhao Huang, Hao Peng, Yujie Wang, Chunyang Liu, and Philip S Yu. 2024. Lsenet: Lorentz structural entropy neural network for deep graph clustering. *arXiv preprint arXiv:2405.11801* (2024).
- [29] Li Sun, Zhenhao Huang, Qiqi Wan, Hao Peng, and Philip S Yu. 2024. Spiking graph neural network on riemannian manifolds. *Advances in Neural Information Processing Systems* 37 (2024), 34025–34055.
- [30] Li Sun, Zhenhao Huang, Hua Wu, Junda Ye, Hao Peng, Zhengtao Yu, and Philip S Yu. 2023. Deep Ricci: Self-supervised graph structure-feature co-refinement for alleviating over-squashing. In *2023 IEEE International Conference on Data Mining (ICDM)*. IEEE, 558–567.
- [31] Li Sun, Zhenhao Huang, Ming Zhang, and Philip S Yu. 2025. Deeper with Riemannian Geometry: Overcoming Oversmoothing and Oversquashing for Graph Foundation Models. *arXiv preprint arXiv:2510.17457* (2025).
- [32] Li Sun, Zhenhao Huang, Suyang Zhou, Qiqi Wan, Hao Peng, and Philip Yu. 2025. RiemannGFM: Learning a graph foundation model from riemannian geometry. In *Proceedings of the ACM on Web Conference 2025*. 1154–1165.
- [33] Li Sun, Junda Ye, Hao Peng, and Philip S Yu. 2022. A self-supervised riemannian gnn with time varying curvature for temporal graph learning. In *Proceedings of the 31st ACM international conference on information & knowledge management*. 1827–1836.
- [34] Li Sun, Zhongbao Zhang, Pengxin Ji, Jian Wen, Sen Su, and Philip S Yu. 2019. DNA: Dynamic social network alignment. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 1224–1231.
- [35] Li Sun, Zhongbao Zhang, Junda Ye, Hao Peng, Jiawei Zhang, Sen Su, and Philip S Yu. 2022. A self-supervised mixed-curvature graph neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4146–4155.
- [36] Li Sun, Zhongbao Zhang, Jiawei Zhang, Feiyang Wang, Yang Du, Sen Su, and Philip S Yu. 2020. Perfect: A hyperbolic embedding for joint user and community alignment. In *2020 IEEE international conference on data mining (ICDM)*. IEEE, 501–510.
- [37] Li Sun, Zhongbao Zhang, Jiawei Zhang, Feiyang Wang, Hao Peng, Sen Su, and Philip S Yu. 2021. Hyperbolic variational graph neural network for modeling dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4375–4383.
- [38] Xiangguo Sun, Hongzhi Yin, Bo Liu, Hongxu Chen, Jiuxin Cao, Yingxia Shao, and Nguyen Quoc Viet Hung. 2021. Heterogeneous hypergraph embedding for graph classification. In *Proceedings of the 14th ACM international conference on web search and data mining*. 725–733.
- [39] Zhongtian Sun, Anoushka Harit, Jongmin Yu, Jingyun Wang, and Pietro Liò. 2025. Advanced hypergraph mining for web applications using sphere neural networks. In *Companion Proceedings of the ACM on Web Conference 2025*. 1316–1320.
- [40] Bohan Tang, Zexi Liu, Keyue Jiang, Siheng Chen, and Xiaowen Dong. 2024. Training-Free Message Passing for Learning on Hypergraphs. *arXiv preprint arXiv:2402.05569* (2024).
- [41] Loring W. Tu. 2011. *An Introduction to Manifolds*. Springer New York.
- [42] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [43] Peihao Wang, Shenghao Yang, Yunyu Liu, Zhangyang Wang, and Pan Li. 2022. Equivariant hypergraph diffusion neural operators. *arXiv preprint arXiv:2207.06680* (2022).
- [44] Yuxin Wang, Quan Gan, Xipeng Qiu, Xuanjing Huang, and David Wipf. 2023. From hypergraph energy functions to hypergraph neural networks. In *International Conference on Machine Learning*. PMLR, 35605–35623.
- [45] Tianxin Wei, Yuning You, Tianlong Chen, Yang Shen, Jingrui He, and Zhangyang Wang. 2022. Augmentations in hypergraph contrastive learning: Fabricated and generative. *Advances in neural information processing systems* 35 (2022), 1909–1922.
- [46] Linhuang Xie, Shihao Gao, Jie Liu, Ming Yin, and Taisong Jin. 2025. K-hop hypergraph neural network: A comprehensive aggregation approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 21679–21687.
- [47] Leyou Xu and Bo Zhou. 2024. Normalized Laplacian eigenvalues of hypergraphs. *Graphs and Combinatorics* 40, 5 (2024), 92.
- [48] Naganand Yadati. 2024. Oversquashing in Hypergraph Neural Networks. In *The Third Learning on Graphs Conference*.
- [49] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. 2019. Hypergen: A new method for training graph convolutional networks on hypergraphs. *Advances in neural information processing systems* 32 (2019).
- [50] Jielong Yan, Yifan Feng, Shihui Ying, and Yue Gao. 2024. Hypergraph dynamic system. In *The twelfth international conference on learning representations*.
- [51] Murong Yang, Shihui Ying, Xin-Jian Xu, and Yue Gao. 2025. Multi-view Spectral Clustering on the Grassmannian Manifold With Hypergraph Representation. *arXiv preprint arXiv:2503.06066* (2025).

A GLOSSARY OF NOTATIONS

We provide a summary of the notation employed in this paper, along with their respective descriptions, in Table 4.

Table 4: Notations.

Notation	Description
$\mathcal{M}$	A Riemannian manifold
$h_{\mathcal{M}}$	The Cheeger's constant of $\mathcal{M}$
ξ	A small value variable
$\varepsilon(\cdot)$	The radius function
λ_i, ψ_i	The i -th eigenvalue and eigenfunction of the Laplace-Beltrami operator.
∂_t	The partial derivative w.r.t. t
$u(t, \mathbf{x})$	The value of the heat equation solution
$\frac{\partial}{\partial n}$	The derivative in the direction of the outward normal
$K(t)$	The heat kernel
λ	The spectral gap of a closed manifold
$\mathcal{G}$	A hypergraph with node set $\mathcal{V}$, edge set $\mathcal{E}$ and incidence matrix $\mathbf{H}$
$\mathcal{L}$	The normalized hypergraph Laplacian
φ	A multi-layer perceptron (MLP)
Z	The initial condition or function for the heat equation
$\mathbb{I}$	The indicator function
$E(X)$	The Dirichlet energy of feature matrix X
α, β	The coefficients of the Robin condition
S	The interior node set of hypergraph $\mathcal{G}$
S_e	The interior edge set of hypergraph $\mathcal{G}$
$\partial_0 S$	The boundary node set of hypergraph $\mathcal{G}$
$\partial_e S$	The boundary edge set of hypergraph $\mathcal{G}$
$\Gamma_{\partial S}$	The source term of nonhomogeneous Robin condition for boundary sets
$\mathbf{F}_S$	The steady state of the heat equation for interior sets
Γ	The concatenation of $\mathbf{F}_S$ and $\Gamma_{\partial S}$

B Theorems and Proofs

This section provides detailed proofs of the theorems and the derivation of the heat-equation solution.

B.1 Proofs of Theorem 3.3

Theorem 3.3 (Flexibility of the Robin Condition). Consider a thin tube-like region S of length L , whose cross-section at position $s \in [0, L]$ is a small ball of radius $\varepsilon(s)$. The boundary of this region is given by $\partial S = \{0, L\} \times \partial \mathcal{B}_{\varepsilon(s)}$. Let λ_D, λ_N , and λ_R denote the spectral gaps of the Laplace operator under Dirichlet, Neumann, and Robin boundary conditions, respectively. When the radius is constant, i.e., $\varepsilon(s) = \varepsilon_0$, the spectral gaps satisfy the ordering $\lambda_D \geq \lambda_R \geq \lambda_N$.

PROOF. Let S be a thin tube-like region of length L , whose cross-section at position $s \in [0, L]$ is a small ball of radius $\varepsilon(s)$. Assume that $\varepsilon(s) : [0, L] \rightarrow (0, \delta)$ with $L \gg \delta$ for a small $\delta > 0$. First, we recall the Cheeger and Buser's Inequality for a closed manifold $\mathcal{M}$ [4]: if the Ricci curvature of $\mathcal{M}$ satisfies

$$\text{Ric}_{\mathcal{M}} \geq -(n-1)\xi^2, \quad \xi \geq 0, \quad (23)$$

then the spectral gap λ of the Laplace-Beltrami operator satisfies

$$\frac{h_{\mathcal{M}}^2}{4} \leq \lambda < C(\xi \cdot h_{\mathcal{M}} + h_{\mathcal{M}}^2), \quad (24)$$

where $h_{\mathcal{M}}$ is the Cheeger constant and C is a constant. This inequality shows that the spectral gap is closely related to the "bottleneck" of the manifold. Taking S as a submanifold embedded in a closed manifold, the spectral gap $\lambda \sim \inf \varepsilon^{n-1}$. For the thin tube S , when the cross-section radius is constant, i.e., $\varepsilon(s) = \varepsilon_0$, the effect of boundary conditions on the spectral gap can be analyzed through the behavior of the heat kernel. Under the Dirichlet condition, the heat kernel decays at a rate governed by $\exp(-\varepsilon_0^{-2}t)$, reflecting the strong suppression of oscillations at the boundary. In

contrast, under the Neumann condition, the decay is determined by $\exp(-L^{-2}t)$, which corresponds to much slower attenuation along the longitudinal direction of the tube. Since ε_0 is small, we have ε_0^{-2} large and ε_0^{n-1} small. Combining these scaling observations with Cheeger-type estimates, it follows that the spectral gaps satisfy the ordering

$$\lambda_D \geq \lambda_R \geq \lambda_N. \quad (25)$$

□

B.2 Proofs of Theorem 3.2

Theorem 3.2 (Spectral Gap and Dirichlet Energy). Let $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ be the eigenvalues of the hypergraph Laplacian $\mathcal{L}$, then the Dirichlet Energy $E(x) := x^\top \mathcal{L}x$ satisfies

$$E(u_k(t)) \leq e^{-2\lambda_1(t-t_k)} E(Z_k). \quad (26)$$

In particular, the Dirichlet energy decays at least exponentially with rate $2\lambda_1$, where λ_1 is the spectral gap.

PROOF. Since $\mathcal{L}$ is symmetric, there exists an orthonormal eigenbasis $\{\psi_i\}$ with $\mathcal{L}\psi_i = \lambda_i\psi_i$. We expand the initial state as

$$Z_k = \sum_i c_i \psi_i, \quad c_i := \psi_i^\top Z_k. \quad (27)$$

The closed-form solution of the diffusion equation is

$$u_k(t) = e^{-(t-t_k)\mathcal{L}} Z_k = \sum_i c_i e^{-\lambda_i(t-t_k)} \psi_i. \quad (28)$$

Now consider the Dirichlet energy:

$$\begin{aligned} E(u_k(t)) &= u_k(t)^\top \mathcal{L} u_k(t) \\ &= \left(\sum_i c_i e^{-\lambda_i \Delta t} v_i \right)^\top \left(\sum_j \lambda_j c_j e^{-\lambda_j \Delta t} v_j \right), \end{aligned} \quad (29)$$

where $\Delta t := t - t_k \geq 0$. Using orthonormality ($\psi_i^\top \psi_j = \delta_{ij}$), this reduces to

$$E(u_k(t)) = \sum_i \lambda_i |c_i|^2 e^{-2\lambda_i \Delta t}. \quad (30)$$

Since for all $i \geq 1$, we have $\lambda_i \geq \lambda_1$, it follows that

$$E(u_k(t)) \leq e^{-2\lambda_1 \Delta t} \sum_{i \geq 1} \lambda_i |c_i|^2, \quad (31)$$

The $i = 0$ term vanishes because $\lambda_0 = 0$. Hence

$$E(u_k(t)) \leq e^{-2\lambda_1(t-t_k)} E(Z_k). \quad (32)$$

This shows that the Dirichlet energy decays at least exponentially with rate $2\lambda_1$, where λ_1 is the spectral gap of the Laplacian. □

B.3 Proofs of Theorem 3.4

Theorem 3.4 (Mode Preservation under Source Energy Injection). Consider the diffusion equation with source term $\partial_t u_k(t, \mathbf{x}) = -\mathcal{L} u_k(t, \mathbf{x}) + f(t, \mathbf{x})$. Expanding in the Laplacian eigenbasis gives $u(t) = \sum_i \hat{u}_i(t) \psi_i$, $f(t) = \sum_i \hat{f}_i(t) \psi_i$ with coefficients $\hat{u}_i(t) = \langle u(t), \psi_i \rangle$, $\hat{f}_i(t) = \langle f(t), \psi_i \rangle$, and each mode satisfies $\dot{\hat{u}}_i(t) = -\lambda_i \hat{u}_i(t) + \hat{f}_i(t)$. If the source persistently excites mode i over a time step $\Delta t = t_{k+1} - t_k$, i.e., $\hat{f}_i(s) \geq \eta_i > 0$ for all $s \in [t_k, t_{k+1}]$, then its amplitude is bounded below by

$$\hat{u}_i(t_{k+1}) \geq \frac{\eta_i}{\lambda_i} (1 - e^{-\lambda_i \Delta t}) + e^{-\lambda_i \Delta t} \hat{u}_i(t_k). \quad (33)$$

showing that persistent energy injection prevents the mode from vanishing.

PROOF. Let $\mathcal{L}$ be a symmetric hypergraph Laplacian with eigenpairs (λ_i, ψ_i) . Expanding the node features $u(t)$ and source term $f(t)$ in the Laplacian eigenbasis, we have

$$u(t) = \sum_i \hat{u}_i(t) \psi_i, \quad f(t) = \sum_i \hat{f}_i(t) \psi_i, \quad (34)$$

$$\hat{u}_i(t) = \langle u(t), \psi_i \rangle, \quad \hat{f}_i(t) = \langle f(t), \psi_i \rangle. \quad (35)$$

Projecting the diffusion equation $\partial_t u(t) = -\mathcal{L}u(t) + f(t)$ onto ψ_i yields a decoupled ODE for each mode

$$\frac{d}{dt} \hat{u}_i(t) = -\lambda_i \hat{u}_i(t) + \hat{f}_i(t), \quad \hat{u}_i(0) = \langle u_0, \psi_i \rangle. \quad (36)$$

Considering the discrete time step $\Delta t = t_{k+1} - t_k$, the solution is given by Duhamel's formula:

$$\hat{u}_i(t_{k+1}) = e^{-\lambda_i \Delta t} \hat{u}_i(t_k) + \int_{t_k}^{t_{k+1}} e^{-\lambda_i(t_{k+1}-s)} \hat{f}_i(s) ds. \quad (37)$$

Assume that the source injects persistent energy into mode i during this interval, i.e., $\hat{f}_i(s) \geq \eta_i > 0$ for all $s \in [t_k, t_{k+1}]$. Then the integral term can be lower-bounded as

$$\int_{t_k}^{t_{k+1}} e^{-\lambda_i(t_{k+1}-s)} \hat{f}_i(s) ds \geq \eta_i \int_{t_k}^{t_{k+1}} e^{-\lambda_i(t_{k+1}-s)} ds \quad (38)$$

$$= \frac{\eta_i}{\lambda_i} (1 - e^{-\lambda_i \Delta t}). \quad (39)$$

Combining this with the decaying contribution from the previous step, we obtain the lower bound for the mode amplitude at t_{k+1} :

$$\hat{u}_i(t_{k+1}) \geq \frac{\eta_i}{\lambda_i} (1 - e^{-\lambda_i \Delta t}) + e^{-\lambda_i \Delta t} \hat{u}_i(t_k). \quad (40)$$

This shows that as long as the source injects persistent energy during each time step, the amplitude of the mode cannot vanish completely and thereby alleviates oversmoothing. $\square$

B.4 Proofs of Equivalence

Let $w(t, \mathbf{x})$ satisfy $h \circ w = \gamma$ on $\partial_v S$. Define the lifted unknown $v := u - w$. By linearity,

$$\partial_t v = \partial_t u - \partial_t w, \quad -\mathcal{L}v = -\mathcal{L}u + \mathcal{L}w. \quad (41)$$

Assume the original problem is

$$\begin{cases} \partial_t u(t, \mathbf{x}) = -\mathcal{L}u(t, \mathbf{x}), & \mathbf{x} \in S, \\ h(u) = \gamma, & \mathbf{x} \in \partial_v S, \\ u(0, \mathbf{x}) = Z(\mathbf{x}). \end{cases} \quad (42)$$

Substituting $u = v + w$ gives

$$\partial_t v = -\mathcal{L}v + \underbrace{(-\mathcal{L}w - \partial_t w)}_{:= \Gamma(w)}, \quad (43)$$

and the boundary becomes $h(v) = h(u) - h(w) = \gamma - \gamma = 0$ on $\partial_v S$; the initial condition is $v(0, \mathbf{x}) = Z(\mathbf{x}) - w(0, \mathbf{x})$. Conversely, if v solves

$$\begin{cases} \partial_t v(t, \mathbf{x}) = -\mathcal{L}v(t, \mathbf{x}) + \Gamma(w), \\ h(v) = 0, \\ v(0, \mathbf{x}) = Z(\mathbf{x}) - w(0, \mathbf{x}). \end{cases} \quad (44)$$

then setting $u := v + w$ yields $\partial_t u = -\mathcal{L}u$, $h(u) = \gamma$ on $\partial_v S$, and $u(0, \mathbf{x}) = Z(\mathbf{x})$. Therefore the two formulations are equivalent.

B.5 Derivation of the Heat Equation Solution

Recall the coupled update rules:

$$\begin{cases} \mathcal{L}_V^{S \rightarrow S} X_S + \mathcal{L}_V^{S \rightarrow \partial_v S} X_{\partial_v S} - \lambda (\mathcal{P}_{\mathcal{E} \rightarrow V}^{S \rightarrow S} X_{S_e} + \mathcal{P}_{\mathcal{E} \rightarrow V}^{S \rightarrow \partial_v S} X_{\partial_e S}) = F_S, \\ \text{diag}(\alpha_v) X_{\partial_v S} + \text{diag}(\beta_v) \mathcal{L}_V^{S \rightarrow \partial_v S} X_S - \lambda (\mathcal{P}_{\mathcal{E} \rightarrow V}^{S \rightarrow \partial_v S} X_{S_e} + \mathcal{P}_{\mathcal{E} \rightarrow V}^{S \rightarrow \partial_v S} X_{\partial_e S}) = \Gamma_{\partial_v S}, \\ \mathcal{L}_E^{S_e \rightarrow S_e} X_{S_e} + \mathcal{L}_E^{S_e \rightarrow \partial_e S} X_{\partial_e S} - \mu (\mathcal{P}_{V \rightarrow E}^{S \rightarrow S_e} X_S + \mathcal{P}_{V \rightarrow E}^{S \rightarrow \partial_e S} X_{\partial_e S}) = F_{S_e}, \\ \text{diag}(\alpha_e) X_{\partial_e S} + \text{diag}(\beta_e) \mathcal{L}_E^{S_e \rightarrow \partial_e S} X_{S_e} - \mu (\mathcal{P}_{V \rightarrow E}^{S \rightarrow S_e} X_S + \mathcal{P}_{V \rightarrow E}^{S \rightarrow \partial_e S} X_{\partial_e S}) = \Gamma_{\partial_e S}. \end{cases} \quad (45)$$

Then it can be expressed as a compact block-matrix form, which is convenient for unified analysis and computation.

$$A \begin{bmatrix} X_V \\ X_E \end{bmatrix} = \begin{bmatrix} \Gamma_V \\ \Gamma_E \end{bmatrix}, \quad X = \begin{bmatrix} X_S \\ X_{\partial S} \end{bmatrix}, \quad \Gamma = \begin{bmatrix} F_S \\ \Gamma_{\partial S} \end{bmatrix}. \quad (46)$$

For the heat diffusion system, let $\mathbf{D}$, $\mathbf{U}$, and $\mathbf{V}$ denote the diagonal, upper, and lower blocks of the system matrix $\mathbf{A}$, respectively. At the k -th iteration, the learnable diagonal terms are updated as

$$\text{diag}(\alpha_v) = \alpha(X_V^{(k)}; \theta_\alpha), \quad \text{diag}(\beta_v) = \beta(X_V^{(k)}; \theta_\beta), \quad \Gamma = \gamma(X_V^{(k)}; \theta_\gamma). \quad (47)$$

$$\text{diag}(\alpha_e) = \alpha(X_E^{(k)}; \theta_\alpha), \quad \text{diag}(\beta_e) = \beta(X_E^{(k)}; \theta_\beta), \quad \Gamma = \gamma(X_E^{(k)}; \theta_\gamma). \quad (48)$$

The feature update rule is then given by the block iteration scheme

$$X^{(k+1)} = \mathbf{D}^{-1}(\mathbf{V} + \mathbf{U})X^{(k)} + \mathbf{D}^{-1}\Gamma \quad (49)$$

Accordingly, the update of node i at iteration $k+1$ takes the form

$$x_i^{(k+1)} = \underbrace{\sum_{j: j \sim i} (\mathcal{P}_{V \rightarrow V})_{ij} x_j^{(k)}}_{\text{node-to-node aggregation}} + \underbrace{\sum_{e: e \sim i} (\mathcal{P}_{E \rightarrow V})_{ei} x_e^{(k)}}_{\text{edge-to-node aggregation}} + \gamma_i, \quad (50)$$

where $j \sim i$ denotes that i and j share at least one hyperedge, and $e \sim i$ means that node i belongs to hyperedge e . The entries of the propagation operators are specified as follows. For node-to-node aggregation,

$$(\mathcal{P}_{V \rightarrow V})_{ij} = \begin{cases} \frac{2}{\sqrt{d_i d_j}} \sum_{e: i, j \in e} \frac{1}{d_e} & i \sim j, i \in S, j \in S, \\ \frac{1}{\sqrt{d_i d_j}} \sum_{e: i, j \in e} \frac{1}{d_e} & i \sim j, i \in S, j \in \partial_v S, \\ \frac{\beta_v}{\alpha_v} \frac{1}{\sqrt{d_i d_j}} \sum_{e: i, j \in e} \frac{1}{d_e} & i \sim j, i \in \partial_v S, j \in S, \\ 0 & \text{otherwise.} \end{cases} \quad (51)$$

whereas for edge-to-node aggregation,

$$(\mathcal{P}_{E \rightarrow V})_{ij} = \begin{cases} \frac{\mu}{1} \frac{1}{\sqrt{d_i d_j}} & i \sim j, i \in S, j \in S_e, \\ \frac{\mu}{\alpha_v} \frac{1}{\sqrt{d_i d_j}} & i \sim j, i \in \partial_v S, j \in S_e, \\ 0 & \text{otherwise.} \end{cases} \quad (52)$$

For compactness, define $S_{ij} := \sum_{e: i, j \in e} \frac{1}{d_e}$ as the total inverse degree of hyperedges shared by nodes i and j , where d_e is the degree of hyperedge e , and $p_i = \frac{\beta_v}{\alpha_v}$. We define the indicator function $I_i := \mathbb{I}(i \in S)$, which equals 1 if node i belongs to the interior set S

and 0 otherwise. Then the node-to-node propagation operator can be written as

$$\mathcal{P}_{\mathcal{V} \rightarrow \mathcal{V}}(i, j) = \frac{1}{\sqrt{d_i d_j}} S_{ij} \left(I_i I_j + I_i + p_i (1 - I_i) I_j \right), \quad (53)$$

The edge-to-node propagation operator is

$$P_{\mathcal{E} \rightarrow \mathcal{V}}(i, e) = \mu \frac{1}{\sqrt{d_i d_e}} \left(I_i + \frac{1}{\alpha_v} (1 - I_i) \right), \quad (54)$$

where μ is a scaling coefficient, and α_v regulates the relative influence of hyperedges on boundary versus interior nodes. Analogous to the node update, the feature update for hyperedge e at iteration $k + 1$ can be written as

$$x_e^{(k+1)} = \underbrace{\sum_{f: f \sim e} (\mathcal{P}_{\mathcal{E} \rightarrow \mathcal{E}})_{ef} x_f^{(k)}}_{\text{edge-to-edge aggregation}} + \underbrace{\sum_{i: i \sim e} (\mathcal{P}_{\mathcal{V} \rightarrow \mathcal{E}})_{ie} x_i^{(k)}}_{\text{node-to-edge aggregation}} + \gamma_e, \quad (55)$$

where $f \sim e$ indicates that hyperedges e and f share at least one common node, and $i \sim e$ indicates that node i belongs to hyperedge e . The entries of the edge-to-edge propagation operator are defined by

$$(\mathcal{P}_{\mathcal{E} \rightarrow \mathcal{E}})_{ef} = \begin{cases} \frac{2}{\sqrt{d_e d_f}} \sum_{i: i \in e \cap f} \frac{1}{d_i} & e \sim f, e \in \mathcal{S}_e, f \in \mathcal{S}_e, \\ \frac{1}{\sqrt{d_e d_f}} \sum_{i: i \in e \cap f} \frac{1}{d_i} & e \sim f, e \in \mathcal{S}_e, f \in \partial_e \mathcal{S}, \\ \frac{\beta_e}{\alpha_e} \frac{1}{\sqrt{d_e d_f}} \sum_{i: i \in e \cap f} \frac{1}{d_i} & e \sim f, e \in \partial_e \mathcal{S}, f \in \mathcal{S}_e, \\ 0 & \text{otherwise.} \end{cases} \quad (56)$$

For compactness, define

$$S_{ef} := \sum_{i: i \in e \cap f} \frac{1}{d_i}, \quad q_e := \frac{\beta_e}{\alpha_e}, \quad J_e := \mathbb{I}(e \in \mathcal{S}_e). \quad (57)$$

Then the edge-to-edge propagation operator can be expressed as

$$\mathcal{P}_{\mathcal{E} \rightarrow \mathcal{E}}(e, f) = \frac{1}{\sqrt{d_e d_f}} S_{ef} (J_e J_f + J_e + q_e (1 - J_e) J_f), \quad (58)$$

Similarly, the contribution from nodes to hyperedge e is given by

$$(P_{\mathcal{V} \rightarrow \mathcal{E}})_{ie} = \begin{cases} \frac{\lambda}{1} \frac{1}{\sqrt{d_i d_e}}, & i \in e, e \in \mathcal{S}_e, \\ \frac{\lambda}{\alpha_e} \frac{1}{\sqrt{d_i d_e}}, & i \in e, e \in \partial_e \mathcal{S}, \\ 0, & \text{otherwise.} \end{cases} \quad (59)$$

Compact form:

$$\mathcal{P}_{\mathcal{V} \rightarrow \mathcal{E}}(i, e) = \lambda \frac{1}{\sqrt{d_i d_e}} (J_e + \frac{1}{\alpha_e} (1 - J_e)). \quad (60)$$

C Datasets and Baselines

C.1 Datasets

To demonstrate the generality of our approach, we evaluate it on benchmark datasets spanning different domains and structural types.

Cora/DBLP-CA are co-authorship graphs where nodes are authors and hyperedges are joint publications (Cora-CA from the Cora citation network; DBLP-CA from the DBLP database). **Cite-seer** and **Pubmed** are co-citation hypergraphs of scientific articles, linking papers that are cited together. **Congress** models US legislative collaboration with nodes as members and hyperedges as bills connecting sponsors and co-sponsors across chambers. **Senate** uses the same bill-co-sponsor construction for the Senate, with party labels on nodes. **House** links representatives who serve on the same committee, with nodes labeled by party.

C.2 Baselines

- **ED-HNN** [43] approximates continuous equivariant hypergraph diffusion operators and combines hypergraph star expansions with standard message passing.
- **SheafHyperGNN** [13] incorporate cellular sheaves into hypergraphs, introducing linear and nonlinear sheaf Laplacians.
- **PhenomNN** [44] learns node embeddings as minimizers of hypergraph-regularized energies, trained end-to-end via supervised bilevel optimization.
- **HDS** [50] formulates HGNNs as ODE-driven dynamics with control-diffusion steps.
- **SensHNN** [48] mitigates oversquashing by k-hop message passing that permits distant vertices to interact while retaining local cues.
- **FrameHGNN** [24] uses framelet-based hypergraph convolutions with low/high-pass filters plus residual/identity mappings to curb oversmoothing.
- **KHGNN** [46] aggregates along k-hop shortest paths between nodes and hyperedges via bisection nested convolutions (on HyperGINE) to model long-range structure.
- **TF-HNN** [40] introduces a training-free message passing module (TF-MP-Module) that precomputes hypergraph structural information during preprocessing.

D REPRODUCIBILITY STATEMENT

In order to guarantee the reproducibility of our findings, we present a comprehensive list of the hyperparameter settings employed during our experiments. For all datasets, we employed the GeLU activation function, kept the coupling coefficient to be 1 and utilized LayerNorm. The specific configurations used for our model in the hypergraph node classification task are provided in Table 5.

Table 5: Hyperparameter settings for node classification task.

Dataset	n_layer	hid_dim	tau	dropout	lr	w_decay
Cora-CA	3	128	0.001	0.7	5e-05	5e-07
DBLP-CA	5	128	0.01	0.7	1e-03	0.0
Cite-seer	3	512	0.1	0.8	5e-03	5e-05
Pubmed	3	64	0.5	0.6	5e-03	5e-05
Congress	5	512	0.1	0.35	5e-03	0.0
House	5	512	0.5	0.35	3e-04	0.0
Senate	5	512	0.5	0.35	3e-04	0.0
NTU2012	3	512	0.5	0.3	3e-05	0.0
Walmart	5	256	0.5	0.35	3e-04	0.0