

Transferable Hypergraph Attack via Injecting Nodes into Pivotal Hyperedges

Meixia He¹, Peican Zhu^{1*}, Le Cheng^{1,2}, Yangming Guo^{3*}, Manman Yuan⁴, Keke Tang^{5,6*}

¹School of Artificial Intelligence, Optics and Electronics (iOPEN), Northwestern Polytechnical University

²School of Computer Science, Northwestern Polytechnical University

³School of Cybersecurity, Northwestern Polytechnical University

⁴School of Computer Science, Inner Mongolia University

⁵Cyberspace Institute of Advanced Technology, Guangzhou University

⁶Huangpu Research School of Guangzhou University

meixia_he@mail.nwpu.edu.cn, ericcan@nwpu.edu.cn, chengle@mail.nwpu.edu.cn,

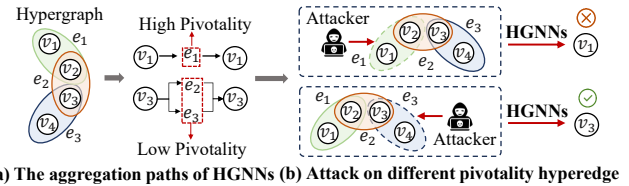
yangming_g@nwpu.edu.cn, yuanman@imu.edu.cn, tangbohutbh@gmail.com

Abstract

Recent studies have demonstrated that hypergraph neural networks (HGNNs) are susceptible to adversarial attacks. However, existing methods rely on the specific information mechanisms of target HGNNs, overlooking the common vulnerability caused by the significant differences in hyperedge pivotality along aggregation paths in most HGNNs, thereby limiting the transferability and effectiveness of attacks. In this paper, we present a novel framework, i.e., **Transferable Hypergraph Attack** via Injecting Nodes into Pivotal Hyperedges (TH-Attack), to address these limitations. Specifically, we design a hyperedge recognizer via pivotality assessment to obtain pivotal hyperedges within the aggregation paths of HGNNs. Furthermore, we introduce a feature inverter based on pivotal hyperedges, which generates malicious nodes by maximizing the semantic divergence between the generated features and the pivotal hyperedges features. Lastly, by injecting these malicious nodes into the pivotal hyperedges, TH-Attack improves the transferability and effectiveness of attacks. Extensive experiments are conducted on six authentic datasets to validate the effectiveness of TH-Attack and the corresponding superiority to state-of-the-art methods.

Introduction

With real-world networks become increasingly complex and diverse, higher-order structures such as hypergraphs have emerged as powerful tools for encapsulating intricate interaction information within graph data (Battiston et al. 2021; Antelmi et al. 2023; Jin et al. 2019). Furthermore, to effectively capture higher-order features within hypergraphs, Hypergraph Neural Networks (HGNNs) have been introduced (Feng et al. 2019; Bai, Zhang, and Torr 2021; Kim et al. 2024; Li et al. 2025), outperforming Graph Neural Networks (GNNs) (Kipf and Welling 2017; Cheng et al. 2024b; Veličković et al. 2017; Cheng et al. 2024a; He et al. 2025a; Wang et al. 2025b) in capturing comprehensive feature information and achieving remarkable results across various downstream graph tasks like node classification, source detection etc. Despite their success, HGNNs have been shown



(a) The aggregation paths of HGNNs (b) Attack on different pivotality hyperedge

Figure 1: An illustration of motivation. (a) In HGNNs, the pivotality of hyperedges along the information aggregation paths varies significantly. (b) Attacking hyperedges with high pivotality significantly impacts HGNN performance, resulting in incorrect predictions for nodes that depend on this hyperedge for feature information.

to be as susceptible to adversarial attacks as GNNs (Hu et al. 2023; Chen et al. 2023; Hu et al. 2023; Chen et al. 2023; Wei et al. 2020), raising concerns about their security.

Attacks against GNNs could be divided into modification attacks and injection attacks based on the attack methodology. Modification attacks (Chang et al. 2020; Jin et al. 2023; Zhu et al. 2024a) alter node/edge attributes or connections to degrade GNN performance, while injection attacks (Sun et al. 2019; Zou et al. 2021; Zhang, Bao, and Pan 2024; Zhu et al. 2024b) introduce malicious nodes/edges into the original graph. However, the higher-order relationships in hypergraphs prevent direct application of graph attacks to HGNNs. Consequently, hypergraph attacks against HGNNs are divided into hypergraph modification attacks and hypergraph injection attacks. For instance, HyperAttack (Hu et al. 2023) modifies hyperedge connections using gradient guidance, and MGHGA (Chen et al. 2023) uses surrogate models for untargeted feature modifications, both focusing on gradient-based modification attacks. Given that injection attacks demand lower privileges and are more practically implementable, recent strategies have emerged, including IE-Attack (He et al. 2025b) and H3NI (Shi et al. 2025). IE-Attack proposed injecting homogeneous nodes generated through Kernel Density Estimation (KDE) into elite hyperedges. H3NI (Shi et al. 2025) employed genetic algorithms to select the hyperedges that injected into the nodes.

*P. Zhu, Y. Guo, K. Tang are joint corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Although existing hypergraph attacks have achieved some advancements, they fail to find the common vulnerability caused by the significant differences in hyperedge pivotality along aggregation paths in most HGNNs, as illustrated in Figure 1(a). For instance, node v_1 aggregates higher-order features solely through hyperedge e_1 , and node v_3 has two distinct aggregation paths. Therefore, attacking the e_1 would directly disrupt its information transmission in this path, rendering the HGNNs incapable of accurately predicting v_1 , as illustrated in Figure 1(b). In contrast, node v_3 has two aggregation paths, enabling it to retain partial true feature information even if e_3 is attacked. Thus, attacking e_3 does not affect HGNN and the correct result is still predicted. This illustrates that hyperedge e_1 exhibits a higher pivotality compared to hyperedges e_2 and e_3 . Moreover, existing attack methods typically attack hyperedges with low pivotality, neglecting the common vulnerability where most HGNNs have pivotal hyperedges in their information aggregation paths. This limits their transferability and effectiveness to HGNNs with different architectures.

To tackle these challenges, we propose the **Transferable Hypergraph Attack** via Injecting Nodes into Pivotal Hyperedges (TH-Attack). First, we analyze the pivotality of hyperedges in the information aggregation paths of HGNNs and design a hyperedge recognizer via pivotality assessment to identify pivotal hyperedges. Considering the propagation characteristics of pivotal hyperedges within information aggregation paths, we develop a feature inverter that generates malicious nodes by maximizing semantic divergence from pivotal hyperedges. Finally, by injecting malicious nodes into pivotal hyperedges, the structural integrity of these hyperedges is compromised. This disruption impedes feature propagation along information aggregation paths, resulting in information aggregation failure across most HGNNs and significantly degrading their performance. We validate the effectiveness of TH-Attack on six publicly available datasets. Extensive experimental results demonstrate that TH-Attack achieves superior attack performance and transferability across multiple HGNNs, outperforming state-of-the-art hypergraph injection attack methods.

Overall, our contributions are summarized as:

- We are the first to identify and formalize the common vulnerability caused by significant differences in hyperedge pivotality along aggregation paths in most HGNNs.
- We propose the transferable hypergraph attack by injecting malicious nodes into the pivotal hyperedges located on information aggregation paths.
- We demonstrate the transferability and effectiveness of our proposed method over baseline approaches through extensive experimental validation.

Related Work

Hypergraph Neural Networks To effectively model the higher-order interactions in complex systems, the hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has been proposed and widely adopted in graph data analysis (Jin et al. 2019; Antelmi et al. 2023). Building upon this foundation, Hypergraph Neural Network (HGNN) (Feng et al. 2019) was proposed to capture more

comprehensive feature representations compared to Graph Neural Networks (GNNs). By incorporating techniques such as hypergraph convolution (Bai, Zhang, and Torr 2021), hypergraph attention (Yadati et al. 2019), and adaptive hyperedge modeling (Zhang, Zou, and Ma 2019), the feature aggregation process in HGNNs had been significantly refined, demonstrating exceptional performance in tasks like node classification. Most of these methods had a common feature of satisfying the node-hyperedge-node feature aggregation mechanism, effectively overcoming the representation limitations of traditional graph structures. The feature extraction capabilities of HGNNs have driven their adoption across diverse domains, with large-scale deployments in applications such as recommendation systems (Wang et al. 2018; Zeng et al. 2023) and biological networks (Yu et al. 2023).

As HGNNs are increasingly used in critical areas like medical diagnosis and financial risk, examining their adversarial robustness has become an urgent security concern.

Adversarial Attack against HGNNs While adversarial attacks on graphs had matured with methods ranging from modification to injection (Chang et al. 2020; Jin et al. 2023; Wang et al. 2020; Ju et al. 2023; Tao et al. 2021, 2023; Wang et al. 2025a), attacking HGNNs demanded distinct strategies due to their intrinsic higher-order interactions. This has led to two analogous but structurally divergent attack categories: hypergraph modification attacks and injection attacks. For instance, the HyperAttack (Hu et al. 2023) modified hyperedge link states of target nodes using gradient guidance, while MGHGA (Chen et al. 2023) targeted untargeted feature modifications during the training phase via surrogate models. Both relied on gradient-based modifications of hyperedge structures. Given the lower privilege requirements and feasibility of injection attacks, methods like IE-Attack (He et al. 2025b) and H3NI (Shi et al. 2025) were developed. IE-Attack generated homogeneous nodes via Kernel Density Estimation (KDE) for injection into elite hyperedges, enhancing stealth and effectiveness through hyperedge group identity. H3NI introduced a black-box node injection, using genetic algorithm and budget models to tackle hyperedge selection challenges specific to hypergraphs.

However, existing hypergraph attack methods did not explore varying pivotality of hyperedges in the information aggregation paths of HGNNs, limiting the transferability and effectiveness of attack methods.

Preliminary and Problem Statement

Hypergraph Neural Networks A hypergraph, represented as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, extends the concept of a standard graph $G = (V, E)$ by permitting hyperedges $e_j \in \mathcal{E}$ to connect two or multiple nodes, such as $e_2 = \{v_2, v_3\}$ shown in Figure 2. The number of nodes is N , and the number of hyperedges is M . The node features are represented by the matrix $\mathcal{X} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{F}|}$, where $|\mathcal{F}|$ is the feature dimension. The structure of the hypergraph is captured by its incidence matrix $H \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{E}|}$, defined as:

$$H_{ij} = \begin{cases} 1, & \text{if } v_i \in e_j, \\ 0, & \text{if } v_i \notin e_j. \end{cases} \quad (1)$$

The feature aggregation of HGNNs at l -th layer is represented as follows:

$$\mathcal{X}^{(l+1)} = \sigma \left(D_{\mathcal{V}}^{-\frac{1}{2}} H W D_{\mathcal{E}}^{-1} H^T D_{\mathcal{V}}^{-\frac{1}{2}} \mathcal{X}^{(l)} \Theta^{(l)} \right), \quad (2)$$

where $D_{\mathcal{V}}$ is the node degree matrix with elements $D_{\mathcal{V},ii} = \sum_j \mathcal{W}_{jj} H_{ij}$, $D_{\mathcal{E}}$ is the hyperedge degree matrix with elements $D_{\mathcal{E},jj} = \sum_i H_{ij}$, \mathcal{W} is the hyperedge weight matrix, $\Theta^{(l)}$ is the layer-specific trainable weight parameters for l -th layer, and $\sigma(\cdot)$ denotes a non-linear activation function.

Hypergraph Node Injection Attack Hypergraph Node Injection Attacks compromise HGNNs by injecting m malicious nodes $\mathcal{V}_{mal} = \{v_{mal}^1, \dots, v_{mal}^m\}$ into existing hyperedges of the original hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. This approach preserves legitimate nodes and hyperedges while strategically injecting nodes into existing hyperedges, augmenting them with carefully crafted features \mathcal{X}_{mal} . The attacked hypergraph becomes $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$, where $\hat{\mathcal{V}} = \mathcal{V} \cup \mathcal{V}_{mal}$ and $\hat{\mathcal{E}} = e_j \cup v_{mal}^m \mid e_j \in \mathcal{E}$ denoting nodes injected into hyperedge, such as $e_j = \{v_1, v_2\} \rightarrow \{v_1, v_2, v_{mal}^m\}$. Expanded incidence matrix is $\hat{H} \in \mathbb{R}^{(|\hat{\mathcal{V}}|+m) \times |\hat{\mathcal{E}}|}$, where $\hat{H}_{(|\hat{\mathcal{V}}|+m)j} = 1$ indicate injected nodes $v_{mal}^m \in \mathcal{V}_{mal}$ added to hyperedge $e_j \in \hat{\mathcal{E}}$. The attack is articulated as follows:

$$\min \mathcal{L}_{\text{atk}}(f_{\theta^*}(\hat{\mathcal{G}})), \quad \text{s.t. } \|\hat{\mathcal{G}} - \mathcal{G}\| \leq \Phi, \quad (3)$$

where $\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{atk}}(f_{\theta}(\hat{\mathcal{G}}))$, represents the HGNNs trained on attacked data. Φ is the budget for the number of injected nodes, determined by the perturbation rate η and the number of nodes N .

Existing Challenges of Attack in HGNNs Existing attack methods for HGNNs focus on identifying target hyperedges for alteration or injecting nodes. Hypergraph modification attacks leverage model's gradient information to obtain target hyperedges, while hypergraph injection attacks rely on hypergraph structural characteristics. However, these methods rely on the specific information mechanisms of target HGNNs, overlooking the common vulnerability arising from the significant differences in hyperedge pivotality along aggregation paths in most HGNNs. This oversight limits the transferability and effectiveness of attacks.

Therefore, this study faces two key challenges: (1) investigating varying pivotality of hyperedges in information aggregation paths about HGNNs and developing an effective pivotal hyperedge identification mechanism to maximize the disruptive impact of injected malicious nodes; (2) leveraging the common vulnerabilities in the aggregation structures of diverse HGNNs to design a feature inverter, enabling precise disruption of aggregation paths while preserving the stealthiness of injected nodes.

Method

Hyperedge Recognizer via Pivotality Assessment To tackle the first challenge of attacking HGNNs highlighted in this paper, we first analyze the feature aggregation process of HGNNs, as illustrated in Eq. (2). Specifically, the

node-hyperedge aggregation process is:

$$\mathbf{Z}^{(l)} = \sigma \left(D_{\mathcal{E}}^{-1} H^T D_{\mathcal{V}}^{-\frac{1}{2}} \mathcal{X}^{(l)} \Theta^{(l)} \right), \quad (4)$$

here, $\mathbf{Z}^{(l)} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{F}|}$ represents the hyperedge feature matrix at the l -th layer. This process is defined as $\mathbf{z}_j^{(l)} = \frac{1}{|e_j|} \sum_{v_i \in e_j} \frac{1}{\sqrt{d_{v_i}}} \mathbf{x}_i^{(l)} \Theta^{(l)}$. Furthermore, the hyperedge-node aggregation process is:

$$\mathcal{X}^{(l+1)} = \sigma \left(D_{\mathcal{V}}^{-\frac{1}{2}} H W \mathbf{Z}^{(l)} \right). \quad (5)$$

In component form, this is expressed as: $\mathbf{x}_i^{(l+1)} = \frac{1}{\sqrt{d_{v_i}}} \sum_{e_j \ni v_i} w_{e_j} \mathbf{z}_j^{(l)}$. From the aforementioned feature aggregation processes, it is evident that the update of node feature representations necessitates traversing the node-hyperedge-node aggregation pathway. However, within hypergraph structures, nodes are not limited to belonging to just one hyperedge. Some nodes may acquire higher-order feature representations through multiple hyperedges, while others may only do so through a few. As illustrated in Figure 3, nodes v_1 and v_4 obtain higher-order features solely via hyperedges e_1 and e_2 , respectively. In contrast, nodes v_2 and v_3 acquire such features through two hyperedges.

Consequently, these hyperedges within aggregation pathways involving a limited number of hyperedges have higher pivotality. Disrupting these hyperedge could impair the information propagation process, leading to error results. On this basis, we propose a hyperedge recognizer via pivotality assessment. For a node $v_i \in \mathcal{V}$, we define its isolation degree as its hyperdegree:

$$d_h(v_i) = |\{e_j \in \mathcal{E} \mid v_i \in e_j\}|. \quad (6)$$

Here, node v_j is in pivotal hyperedges if its isolation degree is in pivotality level τ : $d_h(v_i) \leq \tau$. We furthermore analyze the impact of different τ values on attack performance in the experimental section.

The theoretical foundation of the aforementioned pivotality assessment lies in the perturbation propagation characteristics of aggregation paths. And we reveal the following fundamental principles.

Theorem 1 (*Perturbation Amplification in High Pivotality Hyperedges*) *When node v_i aggregates information through the highly pivotal hyperedges (i.e., $d_h(v_i) \leq \tau$), the lower bound of its feature perturbation is given by:*

$$\|\Delta \mathbf{x}_i^{(l+1)}\|_2 \geq \frac{1}{\sqrt{d_{v_i}}} \min_{e_j \ni v_i} w_{e_j} \cdot \|\Delta \mathbf{z}_j^{(l)}\|_2. \quad (7)$$

Proof: Considering the most vulnerable scenario where information propagates through higher pivotality hyperedges, the perturbation amplification effect follows from:

$$\left\| \sum_{e_j \ni v_i} \frac{w_{e_j}}{\sqrt{d_{v_i}}} \Delta \mathbf{z}_j^{(l)} \right\|_2 \geq \frac{1}{\sqrt{d_{v_i}}} \min_{e_j \ni v_i} w_{e_j} \cdot \|\Delta \mathbf{z}_j^{(l)}\|_2,$$

where $\|\Delta \mathbf{x}_i^{(l+1)}\|_2 = \left\| \sum_{e_j \ni v_i} \frac{w_{e_j}}{\sqrt{d_{v_i}}} \Delta \mathbf{z}_j^{(l)} \right\|_2$.

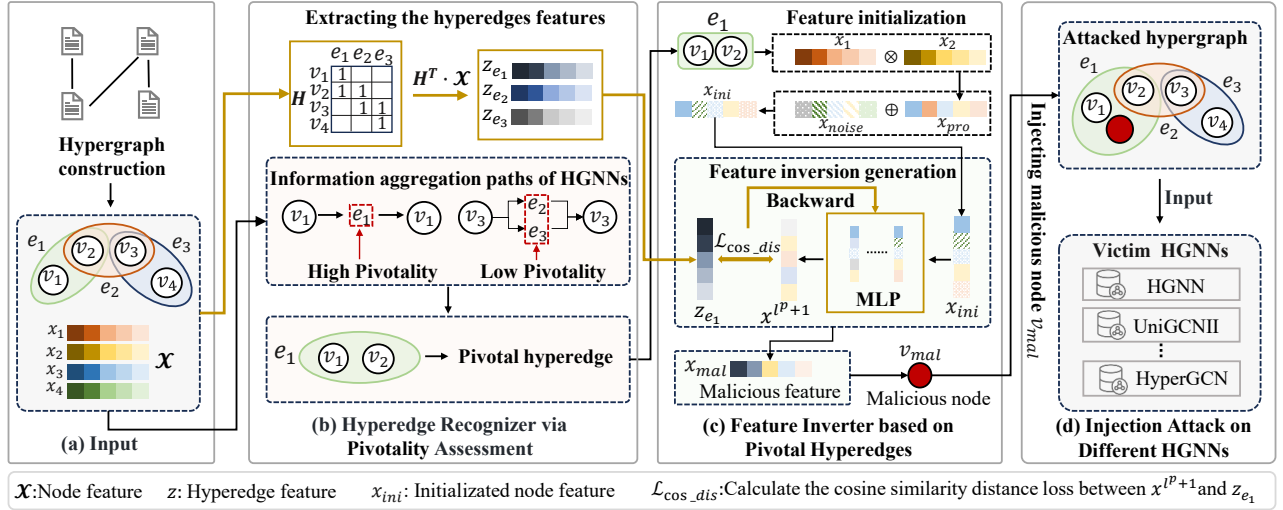


Figure 2: Framework of TH-Attack. (a) A hypergraph constructed from a regular graph as input for the model. (b) We propose a hyperedge recognizer via pivotality assessment to identify pivotal hyperedges. (c) Generating malicious nodes using a feature inverter based on pivotal hyperedges. (d) By injecting malicious nodes into the pivotal hyperedge to obtain an attacked hypergraph, which serves as input for different HGNNs and disrupts the performance of HGNNs.

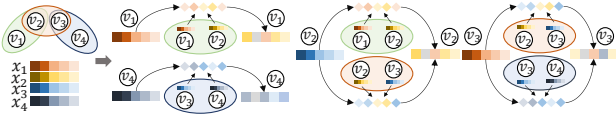


Figure 3: The information aggregation paths of HGNNs.

Theorem 2 (*Perturbation Amplification in Low Pivotality Hyperedges*) When node v_k aggregates information through a lower pivotality hyperedges (i.e., $d_h(v_k) > \tau$), the upper bound of its feature perturbation is given by:

$$\|\Delta \mathbf{x}_k^{(l+1)}\|_2 \leq \frac{1}{\sqrt{d_{v_k}}} \sum_{e_j \ni v_k} w_{e_j} \|\Delta \mathbf{z}_j^{(l)}\|_2. \quad (8)$$

Proof: Applying the triangle inequality to the aggregation formula:

$$\left\| \sum_{e_j \ni v_k} \frac{w_{z_j}}{\sqrt{d_{v_k}}} \Delta \mathbf{z}_j^{(l)} \right\|_2 \leq \frac{1}{\sqrt{d_{v_k}}} \sum_{e_j \ni v_k} w_{e_j} \|\Delta \mathbf{z}_j^{(l)}\|_2. \quad (9)$$

Theorems 1 and 2 collectively establish the theoretical basis for the pivotality assessment mechanism. Nodes with higher pivotality hyperedges in aggregation paths exhibit perturbation amplification, where attacking associated hyperedges amplifies feature perturbation $\|\Delta \mathbf{x}_i\|_2$ by at least $\frac{1}{\sqrt{d_{v_i}}} \min w_{e_j}$. Conversely, nodes in hyperedges with lower pivotality demonstrate perturbation attenuation, where disturbance energy is dissipated through redundant pathways. This validates the scientific rationale for selecting pivotal nodes via the pivotality level τ and constructing the attack pivotal hyperedges $\mathcal{E}_{all-piv}$ from their associated hyperedges.

Therefore, we obtain the pivotal hyperedges in information aggregation paths of HGNNs:

$$\mathcal{E}_{all-piv} = \{e_j \in \mathcal{E} \mid \exists v_i \in e_j \text{ such that } d_h(v_i) \leq \tau\}. \quad (10)$$

Then, we randomly select a corresponding number of hyperedges as final pivotal hyperedges \mathcal{E}_{A-P} from $\mathcal{E}_{all-piv}$ based on this budget, typically not exceeding 5% of all nodes.

Feature Inverter based on Pivotal Hyperedges In the information aggregation path of HGNNs, hyperedges act as critical intermediaries within the ‘‘node-hyperedge-node’’ feature propagation chain, playing a central role in regulating the flow of information. Traditional feature generators focus on learning the feature distribution of target nodes to generate features that conform to this distribution, while neglecting the propagation characteristics of the target nodes. To effectively disrupt the information propagation properties of pivotal hyperedges, we propose a feature inverter based on pivotal hyperedges.

For each pivotal hyperedge $e_j \in \mathcal{E}_{A-P}$, we first generate an initial feature $\mathbf{x}_{ini}^{(j)}$ by combining the internal node features of the hyperedge with random noise:

$$\mathbf{x}_{ini}^{(j)} = \mathbf{x}_{pro}^{(j)} \oplus \mathcal{N}(0, \mu^2), \quad (11)$$

where $\mathbf{x}_{pro}^{(j)} = \prod_{v_i \in e_j} \mathbf{x}_i$ represents an element-wise product along the feature dimension, exemplified by $e_1 = \{v_1, v_2\}$ and $\mathbf{x}_{pro}^{(1)} = \mathbf{x}_1 \otimes \mathbf{x}_2$ in Figure 2(c). $\mathcal{N}(0, \mu^2)$ represents Gaussian noise with zero mean and variance μ^2 . This design enhances feature diversity while preserving statistical correlations with original hyperedge features. The $\mathbf{x}_{ini}^{(j)}$ matches the dimensionality of node features, providing an initial condition for subsequent feature generation.

Furthermore, we utilize the Multi-Layer Perceptron (MLP) (Riedmiller and Lernen 2014) to enhance the initial confusion features, maximizing the feature confusion effect while maintaining the concealment of the topological struc-

ture. The formula is defined as follows:

$$\mathbf{x}^{(l^p+1)} = \sigma(\mathbf{W}^{(l^p)}\mathbf{x}^{(l^p)} + \mathbf{b}). \quad (12)$$

Herein, $\mathbf{W}^{(l^p)}$ is the weight of the l^p -th layer. When $l^p = 0$, $\mathbf{x}^{(l^p)}$ is the $\mathbf{x}_{ini}^{(j)}$. \mathbf{b} is the offset and σ is the LeakyReLU activation function. Subsequently, the generated malicious feature $\mathbf{x}_{mal}^{(j)} = \text{Softmax}(\mathbf{x}^{(l^p+1)})$ is obtained.

To maximize the semantic divergence between the generated features and the pivotal hyperedge, while constraining the feature deviation using a differentiable threshold regularization term, we design a loss function:

$$\mathcal{L}_{cos.dis} = \cos(\mathbf{x}_{mal}^{(j)}, \mathbf{z}_{e_j}) + \lambda \cdot \mathcal{L}_{reg}, \quad (13)$$

where $\mathcal{L}_{reg} = \max(\cos(\mathbf{x}_{mal}^{(j)}, \mathbf{z}_{e_j}) - t, 0)$, and t is the similarity threshold. $\mathbf{z}_{e_j} \in \mathbf{Z}$ is the hyperedge feature obtained by $\mathbf{Z} = H^T \cdot \mathcal{X}$, which is a simplified version derived from Eq. (4) relies on the hypergraph structure, without knowledge of model parameters. The similarity constraint ensures that the generated features effectively disrupt the pivotal hyperedge, while the $\lambda \cdot \mathcal{L}_{reg}$ promotes feature smoothness to maintain stealthiness. Model parameters of inverter are updated via backpropagation, allowing the generated features to evolve along the direction of the loss gradient. The hyperparameter λ controls the trade-off between attack strength and stealthiness. Consequently, we obtain the malicious nodes $v_{mal}^{(j)}$ with malicious feature $\mathbf{x}_{mal}^{(j)}$.

Injection Attack on Different HGNNs For each pivotal hyperedge $e_j \in \mathcal{E}_{A-P}$, we generate a malicious node $v_{mal}^{(j)}$, which is then injected into the pivotal hyperedge e_j , as illustrated in Figure 2(d). Consequently, for the hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, all selected pivotal hyperedges are denoted as \mathcal{E}_{A-P} according to Eq. (10). Thus, the attacked hypergraph $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$ is obtained by injecting malicious nodes \mathcal{V}_{mal} generated through feature inverter into the selected pivotal hyperedges \mathcal{E}_{A-P} , where $\hat{\mathcal{V}} = \mathcal{V} \cup \mathcal{V}_{mal}$, and $\hat{\mathcal{E}} = \{e_j \cup v_{mal}^{(j)} \mid e_j \in \mathcal{E}\}$, denoting nodes injected into hyperedge $e_j = \{v_1, v_2\} \rightarrow \{v_1, v_2, v_{mal}^{(j)}\}$. The incidence matrix $\hat{H} \in \mathbb{R}^{(|\hat{\mathcal{V}}|+m) \times |\hat{\mathcal{E}}|}$ is expanded as:

$$\hat{H}_{ij} = \begin{cases} 1, & \text{if } v_i \in e_j \text{ and } v_i \in \mathcal{V}, \\ 1, & \text{if } v_i \in \mathcal{V}_{mal} \text{ and } e_j \in \mathcal{E}_{A-P}, \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

It is evident that the node dimension of \hat{H} increases by m , while the hyperedge dimension remains unchanged. Additionally, the feature matrix \mathcal{X} has been augmented with the injected malicious nodes features \mathcal{X}_{mal} for malicious nodes \mathcal{V}_{mal} . Therefore, we obtain the attacked data for the attack on HGNNs, which includes the attacked incidence matrix \hat{H} and the disturbed node features $\hat{\mathcal{X}} = \{\mathcal{X}, \mathcal{X}_{mal}\}$.

It is crucial to emphasize that our injection attack operates under a black-box scenario, achieving transferable attacks by feeding the attacked hypergraph into diverse HGNNs. The efficacy of this black-box attack stems from precise exploitation of inherent vulnerabilities in HGNNs. Attackers

require no knowledge of the target model’s specific parameters θ or architectural details, leveraging the attacked hypergraph $\hat{\mathcal{G}}$ as the input of HGNNs. When $\hat{\mathcal{G}}$ is processed by any HGNN model f_θ , its aggregation mechanism:

$$\hat{\mathcal{X}}^{(l+1)} = \sigma\left(\hat{D}_V^{-\frac{1}{2}} \hat{H} \mathcal{W} \hat{D}_E^{-1} \hat{H}^T \hat{D}_V^{-\frac{1}{2}} \hat{\mathcal{X}}^{(l)} \Theta^{(l)}\right). \quad (15)$$

The spectral radius of the hyperedge degree matrix satisfies $\rho(\hat{D}_E^{-1}) \leq \rho(D_E^{-1}) + \mathcal{O}\left(\frac{M}{\min_j [D_E]_{jj}^2}\right)$, leading to miscalibration in information weight allocation. The gradient propagation of malicious features is constrained by the spectral norm of the operator $\hat{D}_V^{-1/2} \hat{H} \mathcal{W} \hat{D}_E^{-1}$: $\|\nabla_{\mathbf{x}_{mal}} \hat{\mathcal{X}}^{(l)}\|_F \leq C \cdot \|\hat{D}_V^{-1/2} \hat{H} \mathcal{W} \hat{D}_E^{-1}\|_2 \cdot \|\hat{\mathcal{X}}^{(l-1)} \Theta^{(l-1)}\|_F$. Such attacks consistently induce performance degradation across various models (HGNN, HyperGCN, . . . , UniGCNII), indicating that their generalization stems from perturbations to the hypergraph structure rather than specific model parameters.

Experiments

Experimental Setting

Datasets Six benchmark datasets are adopted in this paper, including citation networks (Cora, Citeseer, and Pubmed) (Yadati et al. 2019), Co-authorship networks (Cora-CA, DBLP) (Chien et al. 2021). Additionally, we evaluate our approach on the ModelNet40 (Wu et al. 2015), which is widely used in computer vision and graphics. The hypergraph construction process follows the guidelines provided in HGNNs (Feng et al. 2019; Yang et al. 2022). In addition, the partitioning strategy of the dataset follows the Allset (Chien et al. 2021), which are divided into training/validation/test sets.

Parameter Setting This study examines the impact of four key parameters on the experimental outcomes: the attack budget η , the pivotality level τ , and the hyperparameters λ and t used to constrain the magnitude of feature deviation. A detailed sensitivity analysis of these parameters is provided in the experimental section. Additionally, various parameters used in testing attacks on different HGNNs refer to the Allset (Chien et al. 2021).

Evaluating Metrics This paper utilizes the *Accuracy* and *Macro_F1* as evaluation metrics for the performance of TH-Attack. The *Accuracy* and *Macro_F1* indicates the node classification performance of HGNNs, where a lower *Accuracy* and *Macro_F1* signifies a more effective attack.

Baselines We set up five baselines, including random methods (Random, DICE (Waniek et al. 2018)), gradient methods (FGA (Chen et al. 2018), IGA (Wu et al. 2019)), and hypergraph injection methods (IE-Attack (He et al. 2025b)). Additionally, both the TH-Attack and baselines are conducted in a black-box attack scenario against HGNNs (CEGCN and CEGAT (Chien et al. 2021), HGNN (Feng et al. 2019), HyperGCN (Yadati et al. 2019), UniGCNII (Huang and Yang 2021)). All experiments are conducted on a workstation equipped with four NVIDIA RTX 3090 GPUs, which are conducted under the same parameters settings.

Datasets	Models	Clean		Random		DICE		FGA		IGA		IE-Attack		TH-Attack	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
Cora	CEGCN	74.84	72.65	72.06	69.43	71.95	68.99	72.00	69.44	72.26	69.88	<u>70.73</u>	<u>67.92</u>	67.27	63.71
	CEGAT	74.80	72.44	71.99	68.85	72.73	70.21	72.99	70.67	71.68	69.65	<u>71.44</u>	<u>68.45</u>	66.98	60.83
	HGNN	76.41	73.97	74.71	72.19	74.45	71.66	74.41	72.19	73.84	71.33	<u>73.20</u>	<u>69.33</u>	36.08	22.94
	HyperGCN	75.95	70.98	73.14	69.35	73.93	69.81	72.62	68.55	71.09	66.16	<u>68.37</u>	<u>59.86</u>	31.55	16.09
	UniGCNII	80.08	77.99	<u>75.82</u>	<u>72.84</u>	77.23	74.61	76.65	74.67	76.01	73.60	76.57	74.15	39.42	21.78
Cora-CA	CEGCN	75.92	73.64	72.62	70.76	<u>72.38</u>	<u>69.69</u>	74.28	72.43	74.71	72.59	74.80	73.25	50.39	36.71
	CEGAT	76.16	74.69	<u>73.79</u>	<u>71.10</u>	73.77	71.57	73.98	72.33	74.57	72.51	76.01	73.50	60.05	51.31
	HGNN	82.45	80.34	78.90	76.91	79.08	76.98	78.55	76.74	<u>77.13</u>	<u>74.94</u>	80.79	78.72	32.03	18.59
	HyperGCN	76.32	71.99	75.61	72.36	75.75	71.98	72.45	70.23	<u>68.93</u>	<u>65.08</u>	73.82	65.95	17.72	13.02
	UniGCNII	84.68	83.18	79.53	78.21	80.15	78.14	80.57	79.17	<u>79.07</u>	<u>77.49</u>	83.95	82.43	32.72	13.42
Citeseer	CEGCN	68.79	64.25	<u>66.12</u>	<u>61.66</u>	66.20	62.11	67.62	62.69	66.19	61.87	68.68	63.37	52.21	44.03
	CEGAT	70.48	65.15	<u>66.13</u>	<u>61.57</u>	67.60	62.68	67.05	62.68	66.26	61.05	68.77	63.59	62.85	56.63
	HGNN	71.01	66.45	67.92	63.42	67.80	62.91	67.28	<u>62.61</u>	<u>67.01</u>	62.77	68.03	63.77	24.17	17.94
	HyperGCN	70.78	66.51	68.25	63.85	67.96	64.32	67.83	<u>63.01</u>	<u>67.75</u>	63.59	68.75	64.78	20.59	12.38
	UniGCNII	72.18	67.62	69.85	65.36	69.89	65.54	71.44	66.82	<u>69.23</u>	<u>64.80</u>	70.95	65.88	27.00	15.85
Pubmed	CEGCN	86.13	85.72	<u>83.27</u>	<u>82.73</u>	83.36	82.82	83.29	82.75	83.39	83.08	85.72	85.34	48.87	36.81
	CEGAT	85.81	85.45	83.24	82.81	83.29	82.78	83.09	<u>82.46</u>	<u>82.25</u>	81.86	85.89	85.60	52.34	39.58
	HGNN	84.28	84.16	<u>80.99</u>	<u>80.73</u>	81.29	80.92	81.99	81.68	81.60	81.42	84.53	84.48	40.96	30.51
	HyperGCN	76.33	73.29	73.28	69.02	<u>70.29</u>	<u>65.83</u>	83.23	82.99	74.17	70.82	75.69	72.39	35.14	28.00
	UniGCNII	87.86	87.64	84.22	83.81	<u>83.95</u>	<u>83.46</u>	85.09	84.70	84.15	83.75	87.62	87.40	44.13	31.96
DBLP	CEGCN	87.59	86.98	84.14	83.57	84.15	83.79	84.05	<u>83.55</u>	<u>84.03</u>	83.72	87.60	87.05	74.22	70.10
	CEGAT	88.37	87.85	84.64	83.86	84.77	83.89	84.05	83.55	<u>83.74</u>	<u>83.25</u>	88.01	87.50	77.62	75.64
	HGNN	91.03	84.15	86.54	82.37	83.78	81.62	83.19	80.67	82.73	<u>80.13</u>	<u>81.37</u>	84.63	64.97	59.46
	HyperGCN	89.54	89.10	83.84	82.71	<u>81.72</u>	<u>78.79</u>	85.85	85.35	81.83	80.26	82.64	79.37	46.23	14.82
	UniGCNII	91.89	91.61	87.89	87.22	87.79	86.98	88.40	87.97	<u>87.52</u>	<u>86.66</u>	88.31	89.16	38.76	19.87
ModelNet40	CEGCN	89.82	86.99	83.92	78.61	83.74	78.59	83.61	78.32	<u>82.39</u>	<u>75.94</u>	85.94	81.13	78.24	67.49
	CEGAT	92.31	89.78	87.85	82.23	88.16	82.52	88.20	84.10	<u>84.07</u>	<u>78.00</u>	88.34	83.99	78.18	66.09
	HGNN	95.48	93.93	89.13	83.04	89.24	<u>82.98</u>	89.73	<u>84.67</u>	89.12	84.97	93.82	91.72	72.91	62.60
	HyperGCN	86.45	81.48	83.88	71.67	83.67	71.35	<u>77.36</u>	<u>63.05</u>	80.76	66.31	80.53	67.71	50.11	24.16
	UniGCNII	97.86	97.03	93.45	<u>88.89</u>	93.51	89.44	<u>93.36</u>	90.71	93.48	89.38	96.65	95.67	53.50	30.14

Table 1: Comparison of *Accuracy* (*Acc*) (%) and *Macro-F1* (*F1*) (%) of TH-Attack and baselines. The results are averaged over 10 runs on different random seeds. The best values are in **bold**, and the second-best values are underlined.

Model Performance and Parameter Analysis

Performance Comparison with State-of-the-art Methods

Table 1 presents the node classification results of the proposed TH-Attack compared to several baselines across six datasets targeting five HGNNs. The Random and DICE methods employ a strategy of randomly selecting hyperedges for node injection, whereas FGA and IGA utilize gradient guidance to select hyperedges for node injection. All baselines are under a black-box poisoning scenario similar to proposed TH-Attack. It highlights that the original design of IE-Attack was for gray-box evasion attacks. For the sake of comparability, we adapt it to a black-box poisoning scenario, aligning it with TH-Attack.

From Table 1, TH-Attack demonstrates significantly superior attack effectiveness across multiple datasets and models compared to baseline methods. For instance, on the Cora, TH-Attack reduces the *Accuracy* of HGNN from 76.41% to 36.08%, a decrease of 40.33 percentage points, which far exceeds the reduction achieved by random methods (Random/DICE) of approximately 2-3 percentage points and gradient-guided methods (FGA/IGA) of about 1-2 percentage points. Notably, on complex datasets such as ModelNet40, baselines induce only limited performance degradation. Furthermore, TH-Attack exhibits pronounced advantages in cross-model transferability. For instance, IE-Attack mainly focuses on HyperGCN, demonstrating slightly higher effectiveness on the Cora dataset for

HyperGCN compared to other HGNNs. Consequently, TH-Attack’s strategy of injecting nodes generated by the feature inverter into pivotal hyperedges proves to be highly effective across various HGNNs, showing exceptional attack performance and high transferability.

Analysis of the Pivotality Level τ We evaluate attack performance across different pivotality level τ , as shown in Table 2. The level τ is determined by the number of hyperedges the node traverses in the aggregation paths. For example, if the hyperedge count ranges from $\{1, 2, \dots, 12\}$, it is divided into six segments. At level $\tau = 1$, the hyperedge count is $\{1, 2\}$, indicating the highest pivotality. Higher τ correspond to lower pivotality. The results indicate that when $\tau \geq 2$, the *Accuracy* of HGNN generally increases across datasets, suggesting a decline in attack performance. Nevertheless, this experiment confirms that injecting nodes into

Datasets	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$	$\tau = 6$
Cora	39.15	37.16	39.90	43.63	39.95	42.13
Cora-CA	30.00	29.95	31.15	32.36	30.54	30.78
Citeseer	22.52	21.59	22.12	24.65	25.16	24.19
Pubmed	41.10	40.45	42.47	42.98	43.16	43.97

Table 2: The *Accuracy* (%) of TH-Attack under different pivotality level τ against HGNN.

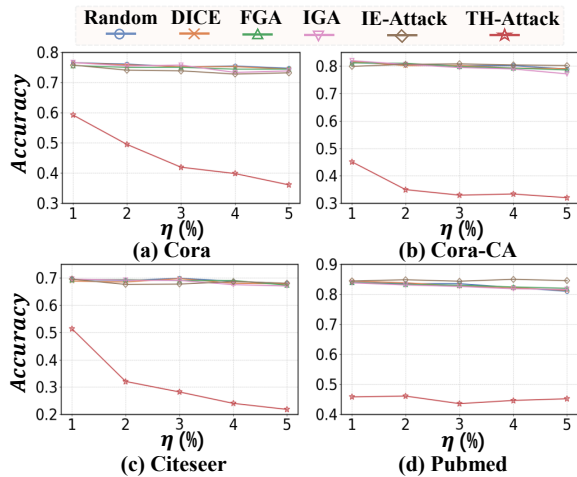


Figure 4: The *Accuracy* of TH-Attack compared to baselines under different η against HGNN.

pivotal hyperedges within information aggregation paths enhances the attack’s performance against HGNN. This is because smaller τ values indicate higher pivotality hyperedges in information aggregation path, resulting in more malicious interference with the aggregation process.

Analysis of the Perturbation Rate η We examine results for η ranging from 1% to 5%, as illustrated in Figure 4. Overall, with η increases, the performance of HGNN under TH-Attack deteriorates significantly more than with baselines. Notably, this experiment demonstrates that our TH-Attack maintains superior attack performance even with minimal node injections. For instance, when $\eta = 1\%$, only 23 nodes are injected into the Cora, resulting in the 0.1717 drop in HGNN *Accuracy* attacked by TH-Attack, compared to a maximum 5% reduction with baselines. Therefore, our strategy of injecting carefully crafted malicious nodes into pivotal hyperedges along information aggregation paths significantly enhances hypergraph attack performance.

Modules	Models	Cora	Cora-CA	Citeseer	Pubmed
w/o HR	HGNN	41.19	38.30	28.69	45.16
	HyperGCN	34.87	24.51	21.43	36.47
	UniGCNII	43.48	40.00	34.05	47.14
w/o FI	HGNN	61.80	58.40	54.26	44.25
	HyperGCN	38.65	26.57	43.93	38.97
	UniGCNII	73.85	77.72	66.94	47.95
w/o CDL	HGNN	61.05	59.31	54.45	55.85
	HyperGCN	59.42	54.66	52.64	52.76
	UniGCNII	59.80	60.34	66.06	46.40
TH-Attack	HGNN	36.08	32.03	24.17	40.96
	HyperGCN	31.55	17.02	20.59	35.14
	UniGCNII	39.42	32.72	27.00	44.13

Table 3: The *Accuracy* (%) of variants for TH-Attack on different HGNNs. w/o means remove this strategy.

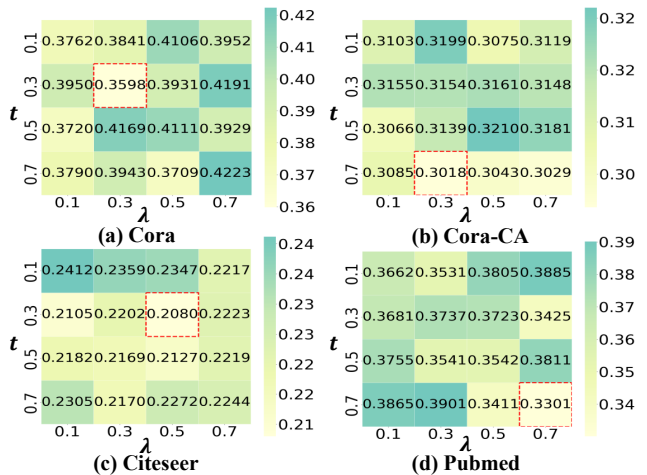


Figure 5: The *Accuracy* of TH-Attack under different λ and t against HGNN.

Analysis of the Hyperparameters λ and t We conduct attack performance concerning the regularization coefficient λ and similarity threshold t , as shown in Figure 5. TH-Attack shows optimal attack performance in $\lambda = 0.3$ and $t = 0.7$ (red box), achieving a minimum *Accuracy* of 0.3018 on Cora-CA. This optimal combination uses a low λ to enhance attack intensity by reducing regularization constraints, and a high t to preserve crucial feature perturbations by relaxing similarity thresholds. Therefore, λ and t maximize semantic differences while preventing excessive suppression, achieving a balance between attack effectiveness and stealth.

Ablation Study and Analysis

We conduct ablation experiments by removing the ‘‘Hyper-edge Recognizer (HR)’’, ‘‘Feature Inverter (FI)’’, and ‘‘Cosine Distance Loss (CDL)’’. Table 3 indicates that excluding these three strategies leads to a decline in the TH-Attack’s performance, resulting in higher *Accuracy*. Notably, the absence of ‘‘CDL’’ has a significant impact on the TH-Attack’s performance. TH-Attack enhances its attack performance and improves cross-model transferability by maximizing the influence of malicious nodes generated by the inverter on pivotal hyperedges. Therefore, the ablation experiments confirm the remarkable effectiveness and transferability of the proposed TH-Attack across different HGNNs.

Conclusion

This paper explores a common vulnerability in HGNNs due to the varying pivotality of hyperedges in the information aggregation path and introduces the TH-Attack. TH-Attack uses a hyperedge recognizer to identify pivotal hyperedges and a feature inverter to generate malicious nodes, which are then injected into these hyperedges to enhance attack transferability and effectiveness. Experiments on six real-world datasets across five HGNNs show that TH-Attack outperforms baselines. Future work will focus on attack strategies for dynamic HGNNs and cross-level knowledge techniques.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (62572400,62472117), the Guangdong Basic and Applied Basic Research Foundation (2025A1515010157), the Science and Technology Projects in Guangzhou (2025A03J0137), the National Natural Science Foundation of China, Regional Science Foundation Project (62466042).

References

- Antelmi, A.; Cordasco, G.; Polato, M.; Scarano, V.; Spagnuolo, C.; and Yang, D. 2023. A survey on hypergraph representation learning. *ACM Computing Surveys*, 56(1): 1–38.
- Bai, S.; Zhang, F.; and Torr, P. H. 2021. Hypergraph convolution and hypergraph attention. *Pattern Recognition*, 110: 107637.
- Battiston, F.; Amico, E.; Barrat, A.; Bianconi, G.; Ferraz de Arruda, G.; Franceschiello, B.; Iacopini, I.; Kéfi, S.; Latora, V.; Moreno, Y.; et al. 2021. The physics of higher-order interactions in complex systems. *Nature Physics*, 17(10): 1093–1098.
- Chang, H.; Rong, Y.; Xu, T.; Huang, W.; Zhang, H.; Cui, P.; Zhu, W.; and Huang, J. 2020. A restricted black-box adversarial framework towards attacking graph embedding models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3389–3396.
- Chen, J.; Wu, Y.; Xu, X.; Chen, Y.; Zheng, H.; and Xuan, Q. 2018. Fast gradient attack on network embedding. *arXiv preprint arXiv:1809.02797*.
- Chen, Y.; Picek, S.; Ye, Z.; Wang, Z.; and Zhao, H. 2023. Momentum gradient-based untargeted attack on hypergraph neural networks. *arXiv preprint arXiv:2310.15656*.
- Cheng, L.; Zhu, P.; Gao, C.; Wang, Z.; and Li, X. 2024a. A heuristic framework for sources detection in social networks via graph convolutional networks. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(11): 7002–7014.
- Cheng, L.; Zhu, P.; Tang, K.; Gao, C.; and Wang, Z. 2024b. GIN-SD: source detection in graphs with incomplete nodes via positional encoding and attentive fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 55–63.
- Chien, E.; Pan, C.; Peng, J.; and Milenkovic, O. 2021. You are allset: A multiset function framework for hypergraph neural networks. *arXiv preprint arXiv:2106.13264*.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3558–3565.
- He, M.; Zhu, P.; Liu, Y.; and Tang, K. 2025a. PLGNN: graph neural networks via adaptive feature perturbation and high-way links. *Complex & Intelligent Systems*, 11(7): 288.
- He, M.; Zhu, P.; Tang, K.; and Guo, Y. 2025b. Hypergraph attacks via injecting homogeneous nodes into elite hyperedges. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 282–290.
- Hu, C.; Yu, R.; Zeng, B.; Zhan, Y.; Fu, Y.; Zhang, Q.; Liu, R.; and Shi, H. 2023. Hyperattack: Multi-gradient-guided white-box adversarial structure attack of hypergraph neural networks. *arXiv preprint arXiv:2302.12407*.
- Huang, J.; and Yang, J. 2021. Unignn: a unified framework for graph and hypergraph neural networks. *arXiv preprint arXiv:2105.00956*.
- Jin, D.; Feng, B.; Guo, S.; Wang, X.; Wei, J.; and Wang, Z. 2023. Local-global defense against unsupervised adversarial attacks on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8105–8113.
- Jin, T.; Yu, Z.; Gao, Y.; Gao, S.; Sun, X.; and Li, C. 2019. Robust L2- hypergraph and its applications. *Information Sciences*, 501: 708–723.
- Ju, M.; Fan, Y.; Zhang, C.; and Ye, Y. 2023. Let graph be the go board: gradient-free node injection attack for graph neural networks via reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 4383–4390.
- Kim, S.; Lee, S. Y.; Gao, Y.; Antelmi, A.; Polato, M.; and Shin, K. 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.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations*, 1–14.
- Li, M.; Fang, Y.; Wang, Y.; Feng, H.; Gu, Y.; Bai, L.; and Lio, P. 2025. Deep hypergraph neural networks with tight framelets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18385–18392.
- Riedmiller, M.; and Lermen, A. 2014. Multi layer perceptron. *Machine learning lab special lecture, University of Freiburg*, 24.
- Shi, H.; Zeng, B.; Yu, R.; Yang, Y.; Zouxia, Z.; Hu, C.; and Shi, R. 2025. H3NI: Non-target-specific node injection attacks on hypergraph neural networks via genetic algorithm. *Neurocomputing*, 613: 128746.
- Sun, Y.; Wang, S.; Tang, X.; Hsieh, T.-Y.; and Honavar, V. 2019. Node injection attacks on graphs via reinforcement learning. *arXiv preprint arXiv:1909.06543*.
- Tao, S.; Cao, Q.; Shen, H.; Huang, J.; Wu, Y.; and Cheng, X. 2021. Single node injection attack against graph neural networks. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 1794–1803.
- Tao, S.; Cao, Q.; Shen, H.; Wu, Y.; Hou, L.; Sun, F.; and Cheng, X. 2023. Adversarial camouflage for node injection attack on graphs. *Information Sciences*, 649: 119611.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. In *International Conference on Learning Representations*, 1–12.
- Wang, J.; Luo, M.; Suya, F.; Li, J.; Yang, Z.; and Zheng, Q. 2020. Scalable attack on graph data by injecting vicious nodes. *Data Mining and Knowledge Discovery*, 34: 1363–1389.

- Wang, X.; Liu, J.; Cheng, Y.; Liu, A.; and Chen, E. 2018. Dual hypergraph regularized PCA for biclustering of tumor gene expression data. *IEEE Transactions on Knowledge and Data Engineering*, 31(12): 2292–2303.
- Wang, X.; Sun, R.; Zhang, Y.; Feng, B.; He, D.; Wang, L.; and Jin, D. 2025a. Stealthy Yet Effective: Distribution-Preserving Backdoor Attacks on Graph Classification. *arXiv preprint arXiv:2509.26032*.
- Wang, X.; Wang, Y.; He, D.; Yu, Z.; Li, Y.; Wang, L.; Dang, J.; and Jin, D. 2025b. Elevating Knowledge-Enhanced Entity and Relationship Understanding for Sarcasm Detection. *IEEE Transactions on Knowledge and Data Engineering*, 37(6): 3356–3371.
- Waniek, M.; Michalak, T. P.; Wooldridge, M. J.; and Rahwan, T. 2018. Hiding individuals and communities in a social network. *Nature Human Behaviour*, 2: 139–147.
- Wei, J.; Yaxin, L.; Han, X.; Yiqi, W.; Shuiwang, J.; and Charu, A. 2020. Adversarial attacks and defenses on graphs: A review, a tool, and empirical studies. *arXiv preprint arXiv:2003.00653*, 00653.
- Wu, H.; Wang, C.; Tyshetskiy, Y.; Docherty, A.; Lu, K.; and Zhu, L. 2019. Adversarial examples on graph data: Deep insights into attack and defense. *arXiv preprint arXiv:1903.01610*.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Yadati, N.; Nimishakavi, M.; Yadav, P.; Nitin, V.; Louis, A.; and Talukdar, P. 2019. Hypergen: A new method for training graph convolutional networks on hypergraphs. *Advances in Neural Information Processing Systems*, 32.
- Yang, C.; Wang, R.; Yao, S.; and Abdelzaher, T. 2022. Semi-supervised hypergraph node classification on hypergraph line expansion. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2352–2361.
- Yu, Y.; Yang, E.; Guo, G.; Jiang, L.; and Wang, X. 2023. Basket representation learning by hypergraph convolution on repeated items for next-basket recommendation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2415–2422.
- Zeng, Y.; Jin, Q.; Bao, T.; and Li, W. 2023. Multi-modal knowledge hypergraph for diverse image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3376–3383.
- Zhang, R.; Zou, Y.; and Ma, J. 2019. Hyper-SAGNN: a self-attention based graph neural network for hypergraphs. *arXiv preprint arXiv:1911.02613*.
- Zhang, X.; Bao, P.; and Pan, S. 2024. Maximizing malicious influence in node injection attack. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 958–966.
- Zhu, G.; Chen, M.; Yuan, C.; and Huang, Y. 2024a. Simple and efficient partial graph adversarial attack: A new perspective. *IEEE Transactions on Knowledge and Data Engineering*, 36(8): 4245–4259.
- Zhu, P.; Pan, Z.; Tang, K.; Cui, X.; Wang, J.; and Xuan, Q. 2024b. Node injection attack based on label propagation against graph neural network. *IEEE Transactions on Computational Social Systems*, 11(5): 5858–5870.
- Zou, X.; Zheng, Q.; Dong, Y.; Guan, X.; Kharlamov, E.; Lu, J.; and Tang, J. 2021. Tdgia: Effective injection attacks on graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2461–2471.