

LGAN: An Efficient High-Order Graph Neural Network via the Line Graph Aggregation

Lin Du¹, Lu Bai^{1*}, Jincheng Li¹, Lixin Cui², Hangyuan Du³, Lichi Zhang⁴, Yuting Chen⁵, Zhao Li⁶

¹School of Artificial Intelligence, Beijing Normal University, Beijing, China

²School of Information, Central University of Finance and Economics, Beijing, China

³School of Computer and Information Technology, Shanxi University, Taiyuan, China

⁴School of Biomedical Engineering, Shanghai Jiaotong University, Shanghai, China

⁵Centre for Learning Sciences and Technologies, The Chinese University of Hong Kong, Hong Kong, China

⁶Zhejiang Lab, Zhejiang, China

dulin@mail.bnu.edu.cn, bailu@bnu.edu.cn

Abstract

Graph Neural Networks (GNNs) have emerged as a dominant paradigm for graph classification. Specifically, most existing GNNs mainly rely on the message passing strategy between neighbor nodes, where the expressivity is limited by the 1-dimensional Weisfeiler-Lehman (1-WL) test. Although a number of k -WL-based GNNs have been proposed to overcome this limitation, their computational cost increases rapidly with k , significantly restricting the practical applicability. Moreover, since the k -WL models mainly operate on node tuples, these k -WL-based GNNs cannot retain fine-grained node- or edge-level semantics required by attribution methods (e.g., Integrated Gradients), leading to the less interpretable problem. To overcome the above shortcomings, in this paper, we propose a novel Line Graph Aggregation Network (LGAN), that constructs a line graph from the induced subgraph centered at each node to perform the higher-order aggregation. We theoretically prove that the LGAN not only possesses the greater expressive power than the 2-WL under injective aggregation assumptions, but also has lower time complexity. Empirical evaluations on benchmarks demonstrate that the LGAN outperforms state-of-the-art k -WL-based GNNs, while offering better interpretability.

1 Introduction

Recently, Graph Neural Networks (GNNs) have proven to be powerful tools for graph-structured data analysis across various domains (Bai et al. 2023; Cui et al. 2024; Qin et al. 2025). Most GNNs follow the message-passing framework, where each node updates its representation by aggregating information from its neighbors, thereby incorporating both node features and graph topology.

A fundamental limitation of such Message Passing Neural Networks (MPNNs) lies in their inability to distinguish certain non-isomorphic graphs, which has been proven to be at most as powerful as the 1-dimensional Weisfeiler-Lehman (1-WL) test (Xu et al. 2019). This limitation is usually caused by ignoring interactions among neighbors,

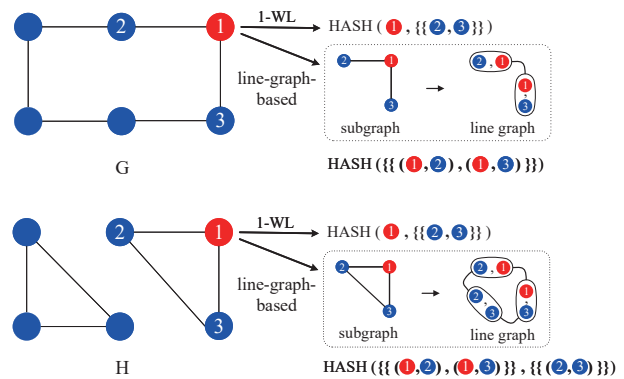


Figure 1: Motivating example: why 1-WL fails and line-graph-based aggregation succeeds.

so that the MPNNs cannot distinguish the subtle structural differences like triangles. For instance, Figure 1 exhibits a pair of graphs G and H , that are non-isomorphic but indistinguishable under the 1-WL test. The red target node 1 in both graphs aggregates messages from an identical multiset of neighbors (i.e., nodes 2 and 3), leading to identical hash updates despite distinct substructures.

To overcome the expressivity limitations of MPNNs, researchers have developed a number of alternative GNNs based on the higher-order k -WL test (Cai, Fürer, and Immerman 1992; Shervashidze et al. 2011) or stronger variants such as the k -FWL (Folklore Weisfeiler-Lehman) test. Since the k -WL test provides a more powerful framework for capturing the structural similarity of graphs, these k -WL-based GNNs simulating the higher-order isomorphism tests can achieve better performance on challenging graph learning tasks (Morris et al. 2021; Maron et al. 2020; Morris, Rattan, and Mutzel 2020; Bodnar et al. 2022). However, they rely on operations over node tuples or sets. While increasing k enables capturing higher-order node interactions (Sato 2020; Huang and Villar 2021), this design also results in greater computational overhead and reduced interpretability.

On the other hand, the line graphs from the classical graph

*Corresponding author: Lu Bai

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

theory have recently offered another promising direction for improving the expressive power of GNNs. For a line graph, each node corresponds to an edge of the original graph structure, and each edge is formed if its connected nodes corresponding to the original edges sharing the same original node. This construction inherently captures the higher-order structural information that is not directly accessible in the original graph. As shown in Figure 1, the line graph of the induced subgraph centered at node 1 encodes not only the node-neighbor relations (1, 2) and (1, 3), but also the neighbor-neighbor interactions (2, 3), enabling more expressive and structure-aware aggregation. Thus, this enriched representation can naturally distinguish graphs G and H with just one iteration. Motivated by this observation, the line-graph-based neural networks have attracted growing attention, particularly in edge-centric tasks such as link prediction (Zhang et al. 2023; Liang et al. 2025) and community detection (Chen, Li, and Bruna 2020). However, their theoretical connection to the k -WL test is still not clear, influencing the employments for other graph-level tasks.

To address the above shortcomings, we propose a novel GNN framework, namely the **Line Graph Aggregation Network (LGAN)**, for graph classification. Our key idea is to perform the information aggregation associated with the line graphs, that are constructed from the induced subgraph centered at each node. The proposed method can employ either the structural relationship between an original graph and its line graph or the favorable properties of the line graph. Overall, **the main contributions** are threefold.

1. **We introduce the LGAN, a novel line-graph-based GNN for graph-level tasks.** Unlike the previous line-graph-based methods focusing on edge-level tasks, the LGAN constructs line graphs over the node-centered subgraphs and further aggregates their messages to represent the whole graph. Moreover, the LGAN employs a dual aggregation design over the line graph to reinforce its isomorphic correspondence with the original graph.
2. **We prove the expressive power of the LGAN and its relationship to the k -WL test.** We prove that the LGAN surpasses the 2-WL test, under the assumption that aggregation-related functions are injective. Moreover, we show that the LGAN simulates the behavior of 2-FWL in a localized manner, while reducing the time complexity from cubic to nearly linear on sparse graphs.
3. **We empirically validate the performance of the LGAN on graph classification tasks.** The LGAN consistently outperforms or matches state-of-the-art baselines including the k -WL-based models. Moreover, the LGAN enables the visualization of important substructures in both synthetic and real-world graphs through edge-level attributions.

2 Preliminary Concepts

2.1 The Basic Notations

Let $G(V, E, X_V)$ be an undirected graph with node set V , edge set E , and node feature matrix $X_V = \{\mathbf{x}_v \in \mathbb{R}^d \mid v \in V\}$. The number of nodes and edges in G are denoted

by $n = |V|$ and $m = |E|$, respectively. The degree of a node $v \in V$ is denoted by d_v . The neighborhood of a node v is defined as $\mathcal{N}(v) = \{u \in V \mid (v, u) \in E\}$. A pair of isomorphic graphs G and H are denoted as $G \cong H$. There exists a bijection $\varphi : V(G) \rightarrow V(H)$ such that $(u, v) \in E(G)$ if and only if $(\varphi(u), \varphi(v)) \in E(H)$. $\{\{\cdot\}\}$ denotes a multiset, a set that allows repeated elements.

2.2 The Weisfeiler-Lehman Test

We commence by introducing the **1-WL** test (Weisfeiler and Lemman 1968), which is a well-known algorithm for graph isomorphism test. Given a labeled graph (G, ℓ) , the algorithm proceeds by iteratively updating the node representations based on the local neighborhood information. Initially, each node is assigned a color by hashing its input features. At iteration $l > 0$, each node $v \in V$ updates its color by aggregating various colors of its multiset neighbors from the previous iteration and applying a hash function to this multiset with its previous color. The update rule is given by

$$c_v^{(l)} = \text{HASH} \left(c_v^{(l-1)}, \{\{c_u^{(l-1)} \mid u \in \mathcal{N}(v)\}\} \right). \quad (1)$$

where $c_v^{(l)}$ denotes the color (i.e., the label) of node v at iteration l , and $\text{HASH}(\cdot)$ is an injective function to preserve distinctions between different inputs. This process is repeated until the color assignments converge, i.e., no further changes occur across iterations. It can be observed that the 1-WL algorithm cannot distinguish all non-isomorphic graphs, as illustrated in Figure 1.

We next introduce the **k -WL** test, a natural generalization of the 1-WL algorithm. Unlike the 1-WL that colors individual nodes, the k -WL test operates on the k -tuples of nodes and refines their labels by incorporating the structure of the induced k -dimensional neighborhoods. Given a labeled graph (G, ℓ) , the algorithm first assigns initial colors to all k -tuples of nodes based on their isomorphism types and labels. At each iteration $l > 0$, the color of a k -tuple $\mathbf{v} = (v_1, \dots, v_k) \in V^k$ is updated by aggregating the colors of its neighboring tuples. Specifically, for each position $i \in \{1, \dots, k\}$, we compute the multiset of colors of all tuples obtained by replacing the i -th element of \mathbf{v} with every possible node $w \in V$, and then update the color of \mathbf{v} . This process is defined as

$$c_{\mathbf{v},i}^{(l)} \leftarrow \{\{c_{\mathbf{v}(i \rightarrow w)}^{(l-1)} \mid w \in V\}\}, \quad (2a)$$

$$c_{\mathbf{v}}^{(l)} = \text{HASH} \left(c_{\mathbf{v}}^{(l-1)}, c_{\mathbf{v},1}^{(l)}, c_{\mathbf{v},2}^{(l)}, \dots, c_{\mathbf{v},k}^{(l)} \right). \quad (2b)$$

where $\mathbf{v}(i \rightarrow w)$ denotes the k -tuple resulting from replacing the i -th entry of \mathbf{v} with node w , $c_{\mathbf{v}(i \rightarrow w)}^{(l-1)}$ is the color assigned to the modified tuple at the previous iteration ($l-1$), $c_{\mathbf{v},i}^{(l)}$ denotes the multiset of colors obtained by varying the i -th element of \mathbf{v} , and $c_{\mathbf{v}}^{(l)}$ is the updated color of tuple \mathbf{v} at iteration l . The process repeats until the coloring stabilizes.

We then describe the **k -FWL** test, a variant of the k -WL test that is often referred to the Folklore version. Given a labeled graph (G, ℓ) , each k -tuple $\mathbf{v} = (v_1, \dots, v_k) \in V^k$ is initially colored according to its isomorphism type and node

labels. At each iteration $l > 0$, the color of a k -tuple \mathbf{v} is updated by first constructing a k -dimensional color vector for each node $w \in V$, and then mapping these vectors together with the previous color into a new color. Specifically,

$$c_{\mathbf{v},w}^{(l)} \leftarrow \left(c_{\mathbf{v}(1 \rightarrow w)}^{(l-1)}, \dots, c_{\mathbf{v}(k \rightarrow w)}^{(l-1)} \right), \quad (3a)$$

$$c_{\mathbf{v}}^{(l)} = \text{HASH} \left(c_{\mathbf{v}}^{(l-1)}, \{ \{ c_{\mathbf{v},w}^{(l)} \mid w \in V \} \} \right). \quad (3b)$$

Here, $c_{\mathbf{v},w}^{(l)}$ denotes the k -dimensional color vector obtained by collecting the colors $c_{\mathbf{v}(i \rightarrow w)}^{(l-1)}$ of modified tuples $\mathbf{v}(i \rightarrow w)$, and $c_{\mathbf{v}}^{(l)}$ is the updated color of tuple \mathbf{v} at iteration l , computed by hashing its previous color with the multiset of these vectors. The iteration continues until convergence.

Note that, for all $k \geq 2$, the expressive power of the k -FWL test matches that of the $(k + 1)$ -WL test (Grohe and Otto 2015; Grohe 2017, 2022).

2.3 The Line Graph

Given an undirected graph $G(V, E)$, its **line graph** $L(G)$ is a dual graph whose node set corresponds to the edge set of G , i.e., $V(L(G)) = E(G)$. For $L(G)$, a pair of nodes $e_1 = (u_1, v_1)$ and $e_2 = (u_2, v_2)$ are connected by an edge, if and only if the corresponding edges in G share at least one common endpoint. The line graph has some important properties that illustrate its relationship to the original graph. In particular, if the original graph G is connected, then its line graph $L(G)$ is also connected. And if two graphs G and H are isomorphic, i.e., $G \cong H$, then their line graphs are also isomorphic, i.e., $L(G) \cong L(H)$. Moreover, an important classical result from the graph theory (Whitney 1992) states the following isomorphism property of the line graph.

Theorem 1 (Whitney’s Isomorphism Theorem). *Let G_1 and G_2 be two connected graphs that are neither a triangle (K_3) nor a claw graph ($K_{1,3}$). If their line graphs are isomorphic, i.e., $L(G_1) \cong L(G_2)$, then the original graphs are also isomorphic, i.e., $G_1 \cong G_2$.*

The above theorem serves as an important theoretical foundation for the design of our work. Specifically, this theorem implies that, except for the special case K_3 and $K_{1,3}$, the structure of a graph can be uniquely reconstructed from its line graph. Our model will handle these exceptions via a tailored aggregation design, thus ensuring the theorem’s necessity and sufficiency within our framework.

3 Related Works

3.1 The k -WL-based GNNs

In recent years, several higher-order methods have been proposed to overcome the expressivity limitations of MPNNs. Among them, the k -WL-based GNNs are most relevant to our work. Notably, Morris et al. (2021) introduced the k -dimensional GNNs (k-GNNs), that improve the scalability by simulating the set-based k -WL, but lose fine-grained structural information. Maron et al. (2020) proposed Provably Powerful Graph Networks (PPGN), which bypass the combinatorial k -WL simulation by directly applying global matrix multiplications to encode pairwise interactions, but

lack localized message passing to explicitly model substructures such as triangles. Morris, Rattan, and Mutzel (2020) proposed a localized variant of the k -WL (δ - k -LGNN) to reduce the computational overhead, yet it remains mainly effective on sparse graphs. Bodnar et al. (2022) introduced CW Networks, which transform the original graph into a cell complex—a topological abstraction that is less intuitive than structures like line graphs. Overall, these models show that k -order invariant or equivariant GNNs can achieve expressivity equivalent to the k -WL isomorphism test.

Although these k -WL-based models capture higher-order structural information, their computational cost grows rapidly with increasing k . Even for moderate values such as $k=3$, the memory and runtime costs become intractable on real-world graphs. Moreover, models defined on node tuples or sets tend to dilute node-level semantics, limiting interpretability in practice. To address these issues, we will propose a novel LGAN model, which performs line-graph-based local aggregation to emulate the k -WL while preserving effective node-level representations.

3.2 The Line-Graph-Based Neural Networks

Line-graph-based neural networks have gained increasing attention for leveraging line graphs to model edge relations (Cai et al. 2021; Zhang et al. 2023; Liang et al. 2025). These methods typically reformulate the edge prediction on the original graph as a node prediction task on its line graph. Beyond encoding the edge-centric relationships, the line graphs possess several well-established theoretical properties that make them particularly suitable for simulating higher-order isomorphism tests such as the k -WL. For example, the line graphs are tightly connected to the fundamental graph properties like connectivity and isomorphism (Whitney 1992), and they support the linear-time complexity algorithms for reconstructing the original graph (Roussopoulos 1973; Lehot 1974). In addition, several studies in hypergraph modeling have shown that the line-graph-based representations are advantageous for capturing complex relational structures (Bai, Ren, and Hancock 2014; Bai, Escolano, and Hancock 2016).

However, existing line-graph-based networks mainly focus on edge-level tasks, without fully leveraging the structural properties. In contrast, our LGAN targets graph-level classification with a tailored aggregation mechanism to capture the structural information. Furthermore, we will bridge the gap between line graphs and the k -WL tests by formally analyzing the expressivity of the proposed framework.

4 The Proposed LGAN Model

In this section, we give the detailed definition of the proposed LGAN model. As illustrated by Figure 2, for each target node, the proposed LGAN performs the localized message passing, by first extracting the 1-hop subgraph around each target node and constructing a line graph over this subgraph. Then the features through both target-neighbor and neighbor-neighbor interactions are aggregated to update the representation of the node. Finally, a multi-layer stacking scheme and global readout function are also employed for graph-level tasks. Below, we define these core components.

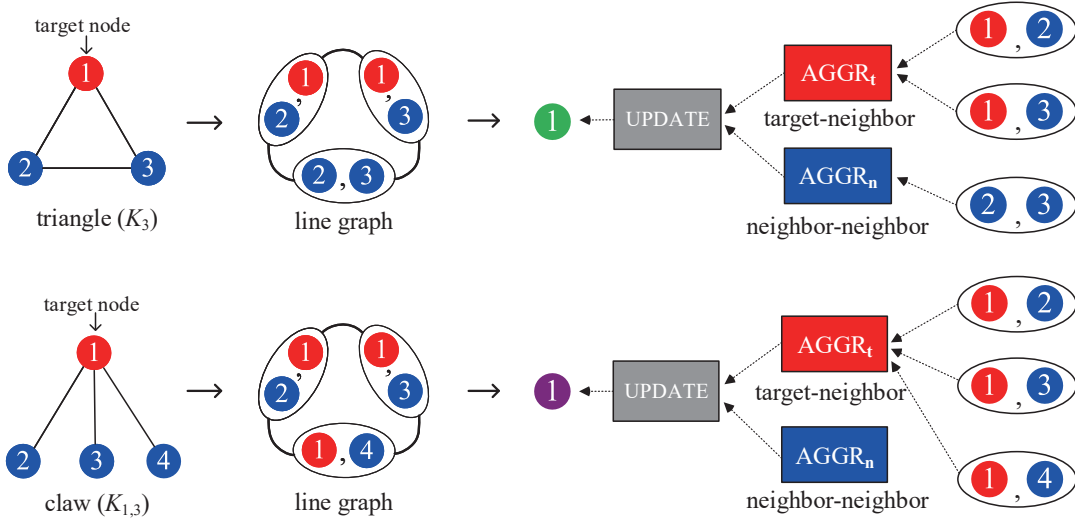


Figure 2: Overview of the Line Graph Aggregation Network (LGAN). For each target node, the LGAN constructs the induced subgraph, transforms it into a line graph, and performs two relation-specific aggregations: AGGR_t (red) for target–neighbor node pairs (i.e., edges sharing the same target node) and AGGR_n (blue) for neighbor–neighbor node pairs, followed by an UPDATE fusion. The examples on K_3 and $K_{1,3}$ illustrate how LGAN resolves the exceptional cases in Whitney’s **Theorem 1**.

4.1 The Line Graph Construction

Given an input graph $G(V, E)$, the LGAN updates the hidden representation of each target node $t \in V$ at layer l , denoted by $\mathbf{h}_t^{(l)} \in \mathbb{R}^d$. It first constructs an induced subgraph $G_t = (V_t, E_t)$ centered at t , where $V_t = \{t\} \cup \mathcal{N}(t)$ and $E_t = \{(u, v) \in E \mid u, v \in V_t\}$. Then, the LGAN builds the line graph $L(G_t) = (E_t, \mathcal{E}_t)$, where each node represents an edge in G_t , and two nodes $e_1, e_2 \in E_t$ are connected in \mathcal{E}_t if they share a common endpoint in G_t .

4.2 The Relation-Specific Aggregation

Let $\mathbf{h}_{\{u,v\}}^{(l-1)}$ denote the representation of a node in the line graph $L(G_t)$ at layer $l-1$, corresponding to the node pair $\{u, v\} \in E(G_t)$ in the original graph. It is computed as

$$\mathbf{h}_{\{u,v\}}^{(l-1)} = \text{COMB}_{\text{pair}}(\mathbf{h}_u^{(l-1)}, \mathbf{h}_v^{(l-1)}), \quad (4)$$

where $\text{COMB}_{\text{pair}}$ denotes a symmetric combination function (e.g., element-wise summation) for mapping unordered node pairs to line graph node representation.

Then, the LGAN defines two relation-specific aggregations over the line graph $L(G_t)$ as follows.

- **Target–Neighbor Aggregation:** Aggregates features of node pairs $\{t, p\}$ incident to the target node t .
- **Neighbor–Neighbor Aggregation:** Aggregates features of node pairs $\{p, q\}$ among the neighbors of t .

These aggregations correspond to the red and blue rectangles in Figure 2, labeled as AGGR_t and AGGR_n , respectively. To summarize the LGAN, we formally express the updated representation of the target node t at layer l as

$$\mathbf{h}_t^{(l)} = \phi(\text{AGGR}_t(\{\{\mathbf{h}_{\{t,p\}}^{(l-1)}\}\}), \text{AGGR}_n(\{\{\mathbf{h}_{\{p,q\}}^{(l-1)}\}\})), \quad (5)$$

where ϕ is a learnable update function fusing the two relation-specific aggregated features. Its inputs are the outputs of AGGR_t and AGGR_n , which are aggregation functions (e.g., sum) operating over multisets of features incident to the target-neighbor pairs and neighbor-neighbor pairs, respectively. Specifically, we adopt concatenation to retain maximal information from both branches. Other fusion strategies (e.g., the element-wise sum, mean, or attention-based gating) can also be applied within the framework.

To further enhance the information preservation and gradient flow, we also introduce a residual variant, termed as **LGAN-res**, which incorporates the previous node representation via a residual connection (He et al. 2015). Specifically, we add a linearly transformed residual from the previous layer to the fused aggregation result. For the **LGAN-res**, the updated representation of the target node t is given by

$$\mathbf{z}_t^{(l)} = \psi(\text{AGGR}_t(\{\{\mathbf{h}_{\{t,p\}}^{(l-1)}\}\}), \text{AGGR}_n(\{\{\mathbf{h}_{\{p,q\}}^{(l-1)}\}\})), \quad (6)$$

$$\mathbf{h}_t^{(l)} = \phi'(W^{(l)} \cdot \mathbf{h}_t^{(l-1)} + \mathbf{z}_t^{(l)}), \quad (7)$$

where ϕ' is an additional multi-layer perceptron (MLP) applied after the residual summation. The residual path includes a linear projection to match dimensions if necessary. ψ , AGGR_t , and AGGR_n retain the same roles as in Eq. (5), maintaining permutation invariance and injectivity.

For the cases where the input graph G contains isolated nodes or nodes with no incident edges (i.e., $E(G_t) = \emptyset$), the corresponding line graph $L(G_t)$ becomes empty, resulting in no available edge features for aggregation. In such scenarios, the fused message \mathbf{z}_t can be set to the zero vector. Consequently, the proposed **LGAN-res** naturally reduces to a variant of GIN (Xu et al. 2019), where the target node is

updated solely based on a learnable projection of its previous representation, i.e., $\mathbf{h}_t^{(l)} = \phi'(W^{(l)} \cdot \mathbf{h}_t^{(l-1)})$. This fallback behavior preserves stability and ensures that isolated nodes still receive meaningful updates.

4.3 The Multi-layer LGAN and Readout

While the above derivation focuses on a single LGAN layer, we stack L such layers to build a deep architecture as

$$\mathbf{H}^{(0)} = \mathbf{X}, \quad \mathbf{H}^{(l)} = \text{LGANLayer}(\mathbf{H}^{(l-1)}), \quad l = 1, \dots, L,$$

where $\mathbf{H}^{(l)}$ denotes the node representation matrix at layer l . Then we adopt a skip-cat (Xu et al. 2018) strategy, where the representations from all layers are concatenated before the final readout. The process is given by

$$\mathbf{H}^{\text{skip}} = \left[\mathbf{H}^{(1)} \parallel \dots \parallel \mathbf{H}^{(L)} \right], \quad \mathbf{h}_G = \text{READOUT}(\mathbf{H}^{\text{skip}}),$$

where \parallel denotes concatenation along the feature dimension, \mathbf{H}^{skip} is the concatenated node representation matrix from all L layers, and \mathbf{h}_G is the final graph-level representation obtained by applying a READOUT function over all nodes.

Note that, the aggregation, update and readout functions of the LGAN follow the DeepSets framework (Zaheer et al. 2018), using the sum-based aggregation and MLPs to achieve the permutation invariance and injectivity. This ensures the maximal expressive power (Xu et al. 2019), while maintaining simplicity and efficiency.

5 Theoretical Properties of the LGAN

We analyze the expressive power and computational complexity of the proposed LGAN with general cases, where the line graph aggregation is feasible. We prove that it surpasses the 2-WL test under standard injectivity assumptions, while implicitly simulating the 2-FWL behavior through the localized aggregation. We also demonstrate that the LGAN achieves the linear time complexity on sparse graphs.

5.1 The Expressive Power of the LGAN

To formalize the expressive advantage of the LGAN, we compare it with the set-based 2-WL test, a widely used and canonical variant operating on unordered node pairs with multiset aggregations. The following theorem shows that the LGAN is at least as powerful as the set-based 2-WL and can even distinguish some graphs that 2-WL cannot.

Theorem 2 (The LGAN is more expressive than the set-based 2-WL). *Let $A : \mathcal{G} \rightarrow \mathbb{R}^d$ be an L -layer LGAN, the following components are injective at every layer; i.e., 1) the node-pair encoder $\text{COMB}_{\text{pair}}$ mapping unordered node pairs to the representation of the corresponding line graph node, 2) the multiset aggregators AGGR_t and AGGR_n (e.g., sum-based DeepSets), 3) the fusion and update functions ψ , ϕ and ϕ' , and 4) the graph-level readout function applied to the multiset of node representations. For any graphs G and H such that the induced subgraph G_t (resp. H_t) for every node t contains at least one edge (i.e., the corresponding line graphs are non-empty), the following holds.*

- If the set-based 2-WL test distinguishes G and H , then the LGAN also maps them to different graph representations, i.e., $A(G) \neq A(H)$.
- There exist graphs G and H such that the set-based 2-WL test fails to distinguish them, but the LGAN produces different representations, i.e., $A(G) \neq A(H)$.

Proof. We first prove by induction that any distinction made by the set-based 2-WL is reflected in the node representations of the LGAN, so that the final injective readout function further leads to different graph-level representations. We then demonstrate that the LGAN succeeds in distinguishing certain graphs where the 2-WL test fails through a counterexample. To commence, let $c_{\{a,b\}}^{(l)}$ denote the color assigned by the 2-WL algorithm to the unordered node pair $\{a, b\}$ at layer l , and let $\mathbf{h}_t^{(l)}$ denote the hidden representation of a target node t produced by the LGAN at the same layer. We assume the inductive hypothesis that for some layer $l \geq 0$, if there exists a node pair $\{a, b\}$ such that $c_{\{a,b\}}^{(l)}(G) \neq c_{\{a,b\}}^{(l)}(H)$, then there exists a node $t \in \{a, b\}$ for which $\mathbf{h}_t^{(l+1)}(G) \neq \mathbf{h}_t^{(l+1)}(H)$.

We first consider the base case $l = 0$. Suppose the initial 2-WL color $c_{\{a,b\}}^{(0)}$ encodes both the initial features of nodes a and b (i.e., \mathbf{x}_a and \mathbf{x}_b), and their connectivity. If $c_{\{a,b\}}^{(0)}(G) \neq c_{\{a,b\}}^{(0)}(H)$, two scenarios are detailed below.

Scenario 1 represents the structural difference, such as when a and b are connected in G but not in H . Thus, the node-pair representation $\mathbf{h}_{\{a,b\}}^{(0)}$ of G exists and is included in the multiset as given by

$$\mathbf{h}_{\{a,b\}}^{(0)} \in \left\{ \left\{ \mathbf{h}_{\{a,w\}}^{(0)} \mid w \in \mathcal{N}_G(a) \right\} \right\} = M_a(G), \quad (8)$$

where $M_a(G)$ is the multiset input to AGGR_t for the target node a (cf. Eq. 5). Since $\mathbf{h}_{\{a,b\}}^{(0)} \notin M_a(H)$ and AGGR_t is injective, it follows that

$$\text{AGGR}_t(M_a(G)) \neq \text{AGGR}_t(M_a(H)). \quad (9)$$

This difference propagates to the node representations updated by the LGAN at the next layer (layer 1), given by

$$\phi(\text{AGGR}_t(M_a(G)), \cdot) \neq \phi(\text{AGGR}_t(M_a(H)), \cdot), \quad (10)$$

where ϕ denotes the update function using concatenation-based MLP, which is injective with respect to its first argument. Hence, $\mathbf{h}_a^{(1)}(G) \neq \mathbf{h}_a^{(1)}(H)$.

Scenario 2 represents the feature difference, such as when $(a, b) \in E(G) \cap E(H)$, but $\mathbf{x}_a(G) \neq \mathbf{x}_a(H)$ or $\mathbf{x}_b(G) \neq \mathbf{x}_b(H)$. In this scenario, the initial node-pair representations are constructed via the injective function $\text{COMB}_{\text{pair}}$ (cf. Eq. 4) ensuring that

$$\text{COMB}_{\text{pair}}(\mathbf{x}_a(G), \mathbf{x}_b(G)) \neq \text{COMB}_{\text{pair}}(\mathbf{x}_a(H), \mathbf{x}_b(H)). \quad (11)$$

i.e., $\mathbf{h}_{\{a,b\}}^{(0)}(G) \neq \mathbf{h}_{\{a,b\}}^{(0)}(H)$. This difference is captured by AGGR_t for target nodes a or b , leading to $\mathbf{h}_a^{(1)}(G) \neq \mathbf{h}_a^{(1)}(H)$ or $\mathbf{h}_b^{(1)}(G) \neq \mathbf{h}_b^{(1)}(H)$. Thus the base case holds.

Now suppose the inductive hypothesis holds for layer $l-1$ and consider the update at layer l . Suppose $c_{\{a,b\}}^{(l)}(G) \neq c_{\{a,b\}}^{(l)}(H)$, by the 2-WL update rule (cf. Eq. 2), this difference must arise from a mismatch in the below multisets

$$\{\{c_{\{w,b\}}^{(l-1)} \mid w \in V\}\} \quad \text{or} \quad \{\{c_{\{a,w\}}^{(l-1)} \mid w \in V\}\}.$$

Without loss of generality, assume the former differs between G and H . The difference in the multiset $\{\{c_{\{w,b\}}^{(l-1)} \mid w \in V\}\}$ implies the existence of at least one node w^* such that $c_{\{w^*,b\}}^{(l-1)}(G) \neq c_{\{w^*,b\}}^{(l-1)}(H)$. By the inductive hypothesis, this implies $\mathbf{h}_s^{(l)}(G) \neq \mathbf{h}_s^{(l)}(H)$ for some $s \in \{w^*, b\}$. By injectivity of $\text{COMB}_{\text{pair}}$ and AGGR_t , this difference further propagates to a target node t where $t \in \mathcal{N}(s)$, leading to $\mathbf{h}_t^{(l+1)}(G) \neq \mathbf{h}_t^{(l+1)}(H)$. Thus the induction holds.

Having established the general result, we now use an example to show that the LGAN can distinguish graphs that the set-based 2-WL cannot. For graphs G and H in Figure 1 with identical initial node labels, any connected node pair (e.g., $\{1, 3\}$) has the same local statistics under the set-based 2-WL, that is, node 1 connects to node 3 and one other node, and is unconnected to the remaining nodes, while node 3 is symmetric. The multisets $\{\{\{w, 3\}\}\}$ and $\{\{\{1, w\}\}\}$ therefore contain the same number of connected and unconnected pairs in both graphs. As a result, the identical node labels and matching multiset statistics lead to the same node-pair representations. The same process applies to unconnected node pairs. Thus, the set-based 2-WL cannot distinguish G and H , but the LGAN can (Figure 1), proving more expressive. \square

5.2 The Theoretical Connection to 2-FWL

The standard k -FWL algorithm colors ordered k -tuples and replaces one coordinate of the tuple by all possible nodes (cf. Eq. (3)), which incurs prohibitive computational costs. To mitigate this, we follow the common practice of converting ordered tuples into unordered sets, obtaining the set-based k -FWL update as

$$c_{\mathbf{v},w}^{(l)} \leftarrow \{\{c_{\mathbf{v}(1 \rightarrow w)}^{(l-1)}, \dots, c_{\mathbf{v}(k \rightarrow w)}^{(l-1)}\}\} \quad (12a)$$

$$c_{\mathbf{v}}^{(l)} = \text{HASH}\left(c_{\mathbf{v}}^{(l-1)}, \{\{c_{\mathbf{v},w}^{(l-1)} \mid w \in V\}\}\right) \quad (12b)$$

Fix the $k = 2$ case and choose an unordered node pair $\mathbf{v} = \{p, q\}$. In Eq. (12a), the 2-FWL traverses every $w \in V$ to form the multiset $c_{\mathbf{v},w}^{(l)} = \{\{c_{\{w,q\}}^{(l-1)}, c_{\{p,w\}}^{(l-1)}\}\}$. Observe that

- The collection $\{\{c_{\{w,q\}}^{(l-1)}, c_{\{p,w\}}^{(l-1)}\}\}$ exactly mirrors the set of node pairs incident to w and $\{p, q\}$ in the LGAN.
- $\text{AGGR}_n(\{\{\mathbf{h}_{\{p,q\}}^{(l-1)}\}\})$ plays the same role as the $c_{\{p,q\}}^{(l-1)}$ inside the 2-FWL hash.
- When we regard the auxiliary node w as the target t in the LGAN, the multiset $\{\{\mathbf{h}_{\{w,q\}}^{(l-1)}, \mathbf{h}_{\{p,w\}}^{(l-1)}\}\}$ becomes precisely the input to AGGR_t .

Hence, each set-based 2-FWL update on the pair $\{p, q\}$ can be locally simulated by a single LGAN update on the

1-hop induced subgraph around the target node $t = w$. Intuitively, the LGAN can be viewed as a sparse and local variant of the set-based 2-FWL, replacing global tuple updates with line graph aggregation over ego networks.

5.3 The Time Complexity Analysis

The LGAN attains the expressive benefits of 2-FWL **without its cubic cost**. Let d_v be the degree of node v . A single LGAN layer aggregates d_v target-neighbor edges and at most $\binom{d_v}{2} = O(d_v^2)$ neighbor-neighbor edges, giving a per-node cost of $O(d_v + d_v^2)$. Summing over all nodes yields a total complexity of

$$O\left(\sum_v d_v + \sum_v d_v^2\right) = O\left(m + \sum_v d_v^2\right), \quad (13)$$

ranging from $O(n)$ on sparse graphs to $O(n^3)$ in the fully connected worst case—yet with a much smaller constant factor than the global 2-FWL due to the locality of the LGAN.

In contrast, higher-order GNNs such as the k -GNNs (Morris et al. 2021) and tensor-based models (Maron et al. 2019, 2020) incur $O(n^k)$ complexity. In practice, the LGAN achieves linear runtime on sparse graphs, compared to the $O(n^3)$ cost of the 2-WL with lower expressivity.

6 Experiments

In this section, we evaluate the LGAN on six graph classification benchmarks from social networks, bioinformatics, and chemistry (Yanardag and Vishwanathan 2015). We also demonstrate its interpretability via edge attribution using Integrated Gradients (Sundararajan, Taly, and Yan 2017).

6.1 Experimental Setups

Setup. We adopt a 10-fold cross-validation strategy with stratified splits that preserve the label distribution across folds. We report our results following the evaluation protocol described in (Xu et al. 2019). Node degree one-hot encodings are used for social networks, and provided node labels or attributes for bioinformatics and chemical datasets.

Hyper-parameters. We tune the number of layers, hidden dimension, dropout, and learning rate for each dataset.

Baselines. We compare the proposed LGAN against three representative categories of models: graph kernels, standard GNNs, and k -WL-based methods. Specifically, the graph kernel baselines include the random walk kernel (RW) (Vishwanathan et al. 2008), shortest path kernel (SP) (Borgwardt and Kriegel 2005), propagation kernel (PK) (Neumann et al. 2016), and 2-WL kernel (Morris et al. 2021). The GNN baselines include the DCNN (Atwood and Towsley 2016), PATCHY-SAN (Niepert, Ahmed, and Kutzkov 2016), DGCNN (Zhang et al. 2018) and GIN (Xu et al. 2019). The k -WL-based methods include the 1-2-3 GNN (Morris et al. 2021), PPGN (Maron et al. 2020), δ - k -LWL (Morris, Rattan, and Mutzel 2020), and CW Networks (Bodnar et al. 2022).

		MUTAG	PTC(MR)	PROTEINS	IMDB-B	IMDB-M	COLLAB
Datasets	# graphs	188	344	1113	1000	1500	5000
	# classes	2	2	2	2	3	3
	Avg. # nodes	18	26	39	20	13	74
Kernels	RW	79.2±2.1	55.9±0.3	59.6±0.1	N/A	N/A	N/A
	SP	81.7±	58.9±	76.4±	59.2±	40.5±	N/A
	PK	76.0±2.7	59.5±2.4	73.7±0.7	N/A	N/A	N/A
	2-WL	77.0±	61.9±	75.2±	72.6±	50.6±	N/A
GNNs	DCNN	N/A	N/A	61.3±1.6	49.1±1.4	33.5±1.4	52.1±0.7
	PATCHYSAN	<u>92.6±4.2</u>	60.0±4.8	75.9±2.8	71.0±2.2	45.2±2.8	72.6±2.2
	DGCNN	85.8±1.7	58.6±2.5	75.5±0.9	70.0±0.9	47.8±0.9	73.8±0.5
	GIN	89.4±5.6	64.6±7.0	76.2±2.8	75.1±5.1	52.3±2.8	80.2±1.9
k -WL-based	1-2-3 GNN	86.1±	60.9±	75.5±	74.2±	49.5±	N/A
	PPGN	90.6±8.7	66.2±6.5	<u>77.2±4.7</u>	73.0±5.8	50.5±3.6	81.4±1.4
	δ -2-LWL	N/A	N/A	<u>75.1±0.3</u>	73.3±0.5	50.2±0.6	N/A
	CW Networks	92.7±6.1	68.2±5.6	77.0±4.3	75.6±3.7	52.7±3.1	N/A
Ours	LGAN	92.5±6.3	67.4±6.2	77.3±3.7	76.7±4.0	<u>53.3±3.2</u>	82.8±1.5
	LGAN- <i>res</i>	92.0±5.8	66.6±7.0	76.8±3.6	<u>76.6±3.1</u>	53.5±2.7	<u>82.7±1.7</u>

Table 1: Test set classification accuracies (%). The mean accuracy and standard deviation are reported. Best performances are highlighted in bold, and second-best performances are underlined. N/A means Not Available.

6.2 Experimental Results and Analysis

Table 1 presents classification accuracies on six benchmark datasets. The proposed LGAN and LGAN-*res* achieve superior or comparable performance across most datasets. Compared to kernel-based methods and standard GNNs, our models offer consistently better accuracy. While several k -WL-based models attain strong results, they face scalability issues on large graphs (e.g., COLLAB). Overall, the LGAN strikes a favorable balance between accuracy and efficiency.

6.3 Interpretability via Edge Attribution

We assess the interpretability of the LGAN using edge attribution (Integrated Gradients) to determine whether it can identify critical substructures contributing to classification.

Figures 3(a) and (b) present two synthetic graphs G' and H' , both of which are indistinguishable by the 2-WL due to having identical initial node features and neighborhood distributions. Similar to the pair in Figure 1, they differ only in the presence or absence of a triangle, which defines their class label. The LGAN successfully assigns high importance to the triangle edges, highlighting its ability to capture essential neighborhood interactions.

Figures 3(c) and (d) show two structurally similar molecules from the Mutagenicity dataset: 2-nitroanisole (mutagen) and 2-nitrobenzyl alcohol (non-mutagen). Both contain a nitro group (NO₂), which is often associated with mutagenic activity. However, substituent difference (-OCH₃ vs. -CH₂OH) determines mutagenicity. The LGAN correctly emphasizes edges near critical functional groups, revealing chemically meaningful patterns. In summary, the LGAN offers fine-grained interpretability beyond the k -WL-based models that lack localized attribution.

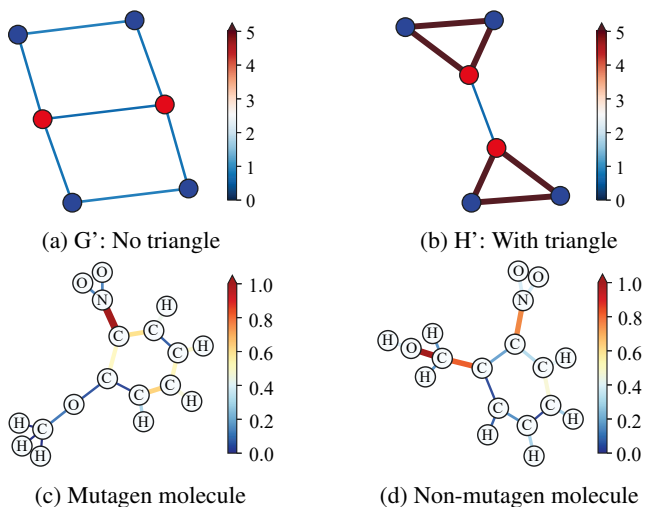


Figure 3: Edge visualizations with Integrated Gradients (redder and thicker edges represent higher importance).

7 Conclusion

In this paper, we introduced the LGAN, a novel GNN that surpasses 2-WL expressivity through localized line graph aggregation. The LGAN is both efficient and interpretable, and performs competitively on graph classification tasks against state-of-the-art methods, including the k -WL-based models. Moreover, the LGAN can be naturally applied to node- and edge-level tasks, as its target-aware aggregation mechanism is general. In future works, we plan to extend the LGAN to more diverse applications, such as complex structural alignment and relational modeling (Bai et al. 2025).

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62576371, T2122020, 62576198, and 62471288).

References

- Atwood, J.; and Towsley, D. 2016. Diffusion-Convolutional Neural Networks. arXiv:1511.02136.
- Bai, L.; Cui, L.; Li, M.; Ren, P.; Wang, Y.; Zhang, L.; Yu, P. S.; and Hancock, E. R. 2025. AEGK: Aligned Entropic Graph Kernels Through Continuous-Time Quantum Walks. *IEEE Transactions on Knowledge and Data Engineering*, 37(3): 1064–1078.
- Bai, L.; Escolano, F.; and Hancock, E. R. 2016. Depth-Based Hypergraph Complexity Traces from Directed Line Graphs. *Pattern Recognition*, 54: 229–240.
- Bai, L.; Jiao, Y.; Cui, L.; Rossi, L.; Wang, Y.; Yu, P. S.; and Hancock, E. R. 2023. Learning Graph Convolutional Networks Based on Quantum Vertex Information Propagation. *IEEE Transactions on Knowledge and Data Engineering*, 35(2): 1747–1760.
- Bai, L.; Ren, P.; and Hancock, E. R. 2014. A Hypergraph Kernel from Isomorphism Tests. In *Proceedings of International Conference on Pattern Recognition*, 3880–3885.
- Bodnar, C.; Frasca, F.; Otter, N.; Wang, Y. G.; Liò, P.; Montúfar, G.; and Bronstein, M. 2022. Weisfeiler and Lehman Go Cellular: CW Networks. arXiv:2106.12575.
- Borgwardt, K. M.; and Kriegel, H.-P. 2005. Shortest-Path Kernels on Graphs. In *Proceedings of the IEEE International Conference on Data Mining*, 74–81.
- Cai, J.-Y.; Fürer, M.; and Immerman, N. 1992. An Optimal Lower Bound on the Number of Variables for Graph Identification. *Combinatorica*, 12(4): 389–410.
- Cai, L.; Li, J.; Wang, J.; and Ji, S. 2021. Line Graph Neural Networks for Link Prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9): 5103–5113.
- Chen, Z.; Li, X.; and Bruna, J. 2020. Supervised Community Detection with Line Graph Neural Networks. arXiv:1705.08415.
- Cui, L.; Bai, L.; Bai, X.; Wang, Y.; and Hancock, E. R. 2024. Learning Aligned Vertex Convolutional Networks for Graph Classification. *IEEE Transactions on Neural Networks and Learning Systems*, 35(4): 4423–4437.
- Grohe, M. 2017. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Cambridge University Press.
- Grohe, M. 2022. The Logic of Graph Neural Networks. arXiv:2104.14624.
- Grohe, M.; and Otto, M. 2015. Pebble Games and Linear Equations. arXiv:1204.1990.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep Residual Learning for Image Recognition. arXiv:1512.03385.
- Huang, N. T.; and Villar, S. 2021. A Short Tutorial on the Weisfeiler-Lehman Test and Its Variants. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 8533–8537.
- Lehot, P. G. 1974. An Optimal Algorithm to Detect a Line Graph and Output Its Root Graph. *Journal of the Association for Computing Machinery*, 21(4): 569–575.
- Liang, J.; Pu, C.; Shu, X.; Xia, Y.; and Xia, C. 2025. Line Graph Neural Networks for Link Weight Prediction. *Physica A: Statistical Mechanics and its Applications*, 661: 130406.
- Maron, H.; Ben-Hamu, H.; Serviansky, H.; and Lipman, Y. 2020. Provably Powerful Graph Networks. arXiv:1905.11136.
- Maron, H.; Fetaya, E.; Segol, N.; and Lipman, Y. 2019. On the Universality of Invariant Networks. arXiv:1901.09342.
- Morris, C.; Rattan, G.; and Mutzel, P. 2020. Weisfeiler and Lemman Go Sparse: Towards Scalable Higher-Order Graph Embeddings. arXiv:1904.01543.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2021. Weisfeiler and Lemman Go Neural: Higher-Order Graph Neural Networks. arXiv:1810.02244.
- Neumann, M.; Garnett, R.; Bauckhage, C.; and Kersting, K. 2016. Propagation Kernels: Efficient Graph Kernels from Propagated Information. *Machine Learning*, 102(2): 209–245.
- Niepert, M.; Ahmed, M.; and Kutzkov, K. 2016. Learning Convolutional Neural Networks for Graphs. arXiv:1605.05273.
- Qin, X.; Bai, L.; Cui, L.; Li, M.; Du, H.; Wang, Y.; and Hancock, E. R. 2025. HA-SCN: Learning Hierarchical Aligned Subtree Convolutional Networks for Graph Classification. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 3245–3253.
- Roussopoulos, N. D. 1973. A Max $\{m, n\}$ Algorithm for Determining the Graph H from Its Line Graph G. *Information Processing Letters*, 2(4): 108–112.
- Sato, R. 2020. A Survey on the Expressive Power of Graph Neural Networks. arXiv:2003.04078.
- Shervashidze, N.; Schweitzer, P.; van Leeuwen, E. J.; Mehlhorn, K.; and Borgwardt, K. M. 2011. Weisfeiler-Lehman Graph Kernels. *Journal of Machine Learning Research*, 12(9): 2539–2561.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic Attribution for Deep Networks. In *Proceedings of the International Conference on Machine Learning*, 3319–3328.
- Vishwanathan, S. V. N.; Borgwardt, K. M.; Kondor, I. R.; and Schraudolph, N. N. 2008. Graph Kernels. arXiv:0807.0093.
- Weisfeiler, B.; and Lemman, A. 1968. The Reduction of a Graph to Canonical Form and the Algebra Which Appears Therein. *Nauchno-Tekhnicheskaya Informatsia*, 2(9): 12–16.
- Whitney, H. 1992. Congruent Graphs and the Connectivity of Graphs. *Hassler Whitney Collected Papers*, 61–79.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful Are Graph Neural Networks? arXiv:1810.00826.
- Xu, K.; Li, C.; Tian, Y.; Sonobe, T.; ichi Kawarabayashi, K.; and Jegelka, S. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. arXiv:1806.03536.

- Yanardag, P.; and Vishwanathan, S. 2015. Deep Graph Kernels. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1365–1374.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Poczos, B.; Salakhutdinov, R.; and Smola, A. 2018. Deep Sets. arXiv:1703.06114.
- Zhang, M.; Cui, Z.; Neumann, M.; and Chen, Y. 2018. An End-to-End Deep Learning Architecture for Graph Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4438–4445.
- Zhang, Z.; Sun, S.; Ma, G.; and Zhong, C. 2023. Line Graph Contrastive Learning for Link Prediction. *Pattern Recognition*, 140: 109537.