



LPFormer: An Adaptive Graph Transformer for Link Prediction

Harry Shomer
shomerha@msu.edu
Michigan State University
East Lansing, USA

Juanhui Li
lijuanh1@msu.edu
Michigan State University
East Lansing, USA

Yao Ma
may13@rpi.edu
Rensselaer Polytechnic Institute
Troy, USA

Bo Wu
bwu@mines.edu
Colorado School of Mines
Golden, USA

Haitao Mao
haitaoma@msu.edu
Michigan State University
East Lansing, USA

Jiliang Tang
tangjili@msu.edu
Michigan State University
East Lansing, USA

ABSTRACT

Link prediction is a common task on graph-structured data that has seen applications in a variety of domains. Classically, hand-crafted heuristics were used for this task. Heuristic measures are chosen such that they correlate well with the underlying factors related to link formation. In recent years, a new class of methods has emerged that combines the advantages of message-passing neural networks (MPNN) and heuristics methods. These methods perform predictions by using the output of an MPNN in conjunction with a “pairwise encoding” that captures the relationship between nodes in the candidate link. They have been shown to achieve strong performance on numerous datasets. However, current pairwise encodings often contain a strong inductive bias, using the same underlying factors to classify all links. This limits the ability of existing methods to learn how to properly classify a variety of different links that may form from different factors. To address this limitation, we propose a new method, **LPFormer**, which attempts to adaptively learn the pairwise encodings for each link. LPFormer models the link factors via an attention module that learns the pairwise encoding that exists between nodes by modeling multiple factors integral to link prediction. Extensive experiments demonstrate that LPFormer can achieve SOTA performance on numerous datasets while maintaining efficiency. The code is available at [The code is available at https://github.com/HarryShomer/LPFormer](https://github.com/HarryShomer/LPFormer).

CCS CONCEPTS

- Computing methodologies → Machine learning

KEYWORDS

link prediction, graph transformer

ACM Reference Format:

Harry Shomer, Yao Ma, Haitao Mao, Juanhui Li, Bo Wu, and Jiliang Tang. 2024. LPFormer: An Adaptive Graph Transformer for Link Prediction. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), August 25–29, 2024, Barcelona, Spain*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3672025>



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

KDD '24, August 25–29, 2024, Barcelona, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08.
<https://doi.org/10.1145/3637528.3672025>

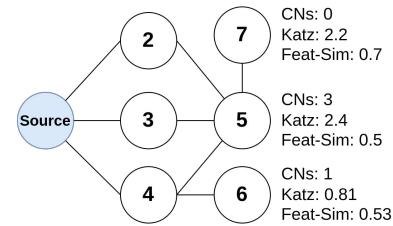


Figure 1: Example of multiple heuristic scores for the candidate links (source, 5), (source, 6), and (source, 7). Each heuristic corresponds to a different LP factor – local (CNs), global (Katz), and feature proximity (Feat-Sim).

1 INTRODUCTION

Link prediction (LP) attempts to predict unseen edges in a graph. It has been adopted in many applications including recommender systems [21], social networks [14], and drug discovery [1]. Traditionally, hand-crafted heuristics were used to identify new links in the graph [2, 38, 55]. Heuristics are often chosen based on factors that typically correlate well with the formation of new links. For example, a popular heuristic is common neighbors (CNs), which assume that the links are more likely to exist between node pairs with more shared neighbors. It has been found that these factors, which we refer to as “LP Factors”, often stem from the local and global structural information and feature proximity [32]. We give an example in Figure 1 that demonstrates different heuristic scores for multiple candidate links. Each heuristic score corresponds to one of the LP factors: CNs for local information, Katz for global, and Feat-Sim for feature proximity. We can observe that the pair (source, 5) has the highest CN and Katz score of the candidate links, indicating an abundance of local and global structural information between the pair. On the other hand, the feature similarity for (source, 5) is the lowest among the candidate links. This indicates that *different LP factors and heuristics have distinct assumptions about why links are formed*.

More recently, message passing neural networks (MPNNs) [16], which are able to learn effective node representations via message passing, have been widely adopted for LP tasks. They predict the existence of a link by combining the node representations of both nodes in the link. However, such a node-centric view is unable to incorporate the pairwise information between the nodes in the link. Because of this, conventional MPNNs have been demonstrated

to be poor link predictors due to their limited capability to learn effective and expressive link representations [43, 52]. To address this issue, recent efforts [51, 56] have attempted to move beyond the node-centric view of traditional MPNNs by equipping them with pairwise information specific to the link being predicted (i.e. the “target link”) [51, 56]. This is done by customizing the message passing process to each target link. However, a concern with this approach is that it can be prohibitively expensive [9], as message passing needs to be run for each individual target link. This is as opposed to traditional MPNNs which only run message passing once for all target links.

To overcome these inefficiencies, recent methods [9, 45, 50] have instead explored ways to inject pairwise information into the model, without individualizing the message passing to each target link. This is done by decoupling the message passing and link-specific pairwise information. By doing so, the message passing only needs to be done once for all target links. To include the pairwise information, these methods, which we refer to as “Decoupled Pairwise MPNNs” (DP-MPNNs), instead learn a “pairwise encoding” to encode the pairwise relationship of the target link. The choice of pairwise encoding is often based on heuristics that correspond to common LP factors (e.g., common neighbors). DP-MPNNs have gained attention as they can achieve promising performance while being much more efficient than methods that customize the message passing mechanism.

However, DP-MPNNs are often limited in the choice of pairwise encoding, using a one-size-fits-all solution for all target links. This has two limitations. (1) The pairwise encoding may fail to consider some integral LP factors. For example, NCNC [45] only considers the 1-hop neighborhood when computing the pairwise encoding, thereby ignoring the global structural information. *This suggests the need for a pairwise encoding that considers multiple types of LP factors.* (2) The pairwise encoding uses the same LP factors for all target links. This assumes that all target links need the same factors. However, it may not necessarily be true. Recently, Mao et al. [32] have shown that different LP factors are necessary to classify different target links. It is evident that even for the same dataset, multiple LP factors are needed to properly predict all target links. This further applies to different datasets, where certain factors are more prominent than others. As such, it faces tremendous challenges when considering multiple types of LP factors. While one factor may effectively model some target links, it will fail for other target links where those patterns aren’t present. *It is therefore desired to consider different LP factors for different target links.*

These observations motivate us to ask – *can we design an efficient method that can adaptively determine which LP factors to incorporate for each individual target link?* Essentially, it requires a pairwise encoding that (a) models multiple LP factors, (b) can be tailored to fit each individual target link, and (c) is efficient to calculate. By doing so, we can flexibly adapt the pairwise information based on the existing needs of each target link. To achieve this, we propose **LPFormer** – Link Prediction TransFormer. LPFormer is a type of graph Transformer [35] designed specifically for link prediction. Given a target link (a, b) , LPFormer models the pairwise encoding via an attention module that learns how a and b relate in the context of various LP factors. This allows for a more customizable set of pairwise encodings that are specific to each target link. Extensive

experiments validate that LPFormer can achieve SOTA on a variety of benchmark datasets. We further demonstrate that LPFormer is better at modeling several types of LP factors, highlighting its adaptability, while also maintaining efficiency on denser graphs.

2 BACKGROUND

2.1 Related Work

Link prediction (LP) aims to model how links are formed in a graph. The process by which links are formed, i.e., link formation, is often governed by a set of underlying factors [5, 30]. We refer to these as “LP factors”. Two categories of methods are used for modeling these factors – heuristics and MPNNs. We describe each class of methods. We further include a discussion on existing graph transformers.

Heuristics for Link Prediction. Heuristics methods [38, 55] attempt to explicitly model the LP factors via hand-crafted measures. Recently, Mao et al. [32] have shown that there are three main factors that correlate with the existence of a link: (1) local structural information, (2) global structural information, and (3) feature proximity. **Local structural information** only considers the immediate neighborhood of the target link. Representative methods include Common Neighbors (CN) [38], Adamic Adar (AA) [2], and Resource Allocation (RA) [55]. They are predicated on the assumption that nodes that share a greater number neighbors exhibit a higher probability of forming connections. **Global structural information** further considers the global structure of the graph. Such methods include Katz [22] and Personalized PageRank (PPR) [7]. These methods posit that nodes interconnected by a higher number of paths are deemed to have larger similarity and, therefore, are more likely to form connections. Lastly, **feature proximity** assumes nodes with more similar features connect [36]. Previous work [39, 54] have shown that leveraging the node features are helpful in predicting links. Lastly, we note that Mao et al. [32] has recently shown that *to properly predict a wide variety of links, it’s integral to incorporate all three of these factors.*

MPNNs for Link Prediction. Message Passing Neural Networks (MPNNs) [16] aim to learn node representations via the message passing mechanism. Traditional MPNNs have been used for LP including GCN [24], SAGE [18], and GAE [25]. However, they have been shown to be suboptimal for LP as they aren’t expressive enough to capture important pairwise patterns [43, 53]. SEAL [51] and NBFNet [56] try to address this by customizing the message passing process to each target link. This allows for the message passing to learn pairwise information specific to the target link. However, these methods have been shown to be unduly expensive as they require a separate round of message passing for each target link. As such, recent methods have been proposed to instead decouple the message passing and pairwise information [9, 45, 50], reducing the time needed to do message passing. Such methods include NCN/NCNC [45] which exploit the common neighbor information and BUDDY [9] and Neo-GNN [50] which consider the global structural information.

Graph Transformers. Recent work has attempted to extend the original Transformer [44] architecture to graph-structured data. Graphomer [48] learns node representations by attending all nodes to each other. To properly model the structural information, they

propose to use multiple types of structural encodings (i.e., structural, centrality, and edge). SAN [27] further considers the use of the Laplacian positional encodings (LPEs) to enhance the learnt structural information. Alternatively, TokenGT [23] considers all nodes and edges as tokens in the sequence when performing attention. Due to the large complexity of these models, they are unable to scale to larger graphs. To address this, several graph transformers [10, 46] have been proposed for node classification that attempt to efficiently attend to the graph. However, while some work [11, 40] have formulated transformers for knowledge graph completion, to our knowledge, **there are no graph transformers designed specifically for LP on uni-relational graphs**.

2.2 Preliminaries

We denote a graph as $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} and \mathcal{E} are the sets of nodes and edges in \mathcal{G} , respectively. The adjacency matrix is represented as $A \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$. The d -dimensional node features are represented by the matrix $X \in \mathbb{R}^{|\mathcal{V}| \times d}$. The set of neighbors for a node v is given by $\mathcal{N}(v)$. The set of overlapping neighbors between two nodes a and b , i.e., the common neighbors (CNs), is expressed by $\mathcal{N}_{(a,b)}^{\text{CN}}$. We further denote the set of nodes that are 1-hop neighbors of only one of a or b as $\mathcal{N}_{(a,b)}^1$ and the nodes that are >1 -hop from both nodes as $\mathcal{N}_{(a,b)}^{>1}$. Lastly, the personalized pagerank (PPR) score for a root node v and an arbitrary node u is given by $\text{ppr}(v, u)$.

3 THE PROPOSED FRAMEWORK

In Section 1, we highlighted the importance of adaptively modeling multiple types of LP factors. However, current methods that use pairwise encodings, i.e., DP-MPNNs, struggle to appropriately achieve this goal. This is due to two issues: (1) They only attempt to model a subset of the potential LP factors (e.g., only local structural information), limiting their ability to model multiple factors. (2) They use a one-size-fits-all approach in regard to pairwise encoding, using the same combination of LP factors for each target link. These issues strongly limit the potential of such methods to properly model a variety of different target links. To overcome these problems, we propose **LPFormer**, a new transformer-based method that can adaptively customize the pairwise information for each target link by considering a variety of different LP factors in an efficient manner.

3.1 A General View of Pairwise Encodings

Recent MPNNs for LP use a decoupled strategy to include the pairwise information [9, 45, 50]. These methods, DP-MPNNs, predict the existence of a link (a, b) via both the node representations and a pairwise encoding $s(a, b)$. They follow the formulation below:

$$H = \text{MPNN}(A, X),$$

$$p(a, b) = \sigma \left(\text{MLP} \left(h_a \odot h_b \parallel s(a, b) \right) \right), \quad (1)$$

where h_i is the representation of node i encoded by the MPNN. Various DP-MPNNs adopt different ways to model the pairwise encoding. For example, NCN [45] models the pairwise encoding $s(a, b)$ as the summation of the node representations of the CNs. The pairwise encodings in these existing methods are typically manually selected or extracted from the graph, which limits the LP

factors they can cover. For example, $s(a, b)$ in NCN and NCNC only capture the local structural information. BUDDY [9] ignores the node features when computing the pairwise encoding. To flexibly model multiple types of LP factors, we propose a general formulation for pairwise encodings as follows,

$$s(a, b) = \sum_{u \in \mathcal{V}} w(a, b, u) \odot h(a, b, u), \quad (2)$$

where $w(a, b, u)$ measures the importance of node u to (a, b) , and $h(a, b, u)$ is the encoding of node u relative to (a, b) . By considering which nodes should be considered for (a, b) and how they are related to the node pair, Eq. (2) can model different LP factors by manually defining $w(a, b, u)$ and $h(a, b, u)$. In particular, we demonstrate how the heuristic methods corresponding to different LP factors can fit into this framework.

Common Neighbors (CNs) [38]: CNs considers the local structural information and is defined for a pair of nodes (a, b) as $\mathcal{N}_{(a,b)}^{\text{CN}} = \mathcal{N}(a) \cap \mathcal{N}(b)$. Eq. (2) is equal to the CNs when $h(a, b, u) = 1$ and:

$$w(a, b, u) = \begin{cases} 1, & \text{when } u \in \mathcal{N}(a) \cap \mathcal{N}(b) \\ 0, & \text{else} \end{cases}. \quad (3)$$

Katz Index [22]: The Katz index models the global structural information. It is defined as weighted summation of the number of paths of different lengths connecting a and b and a decay weight $\beta \in [0, 1]$,

$$\text{Katz}(a, b) = \sum_{l=1}^{\infty} \beta^l A_{a,b}^l.$$

This is equivalent to Eq. (2) where $w(a, b, u) = \sum_{l=1}^{\infty} \beta^l e_a^T A^l$ and

$$h(a, b, u) = \begin{cases} e_b^T, & \text{when } u = b \\ 0, & \text{else} \end{cases},$$

where $e_i \in \mathbb{B}^{|\mathcal{V}|}$ is a one-hot vector for a node i .

Feature Similarity: The feature similarity of the pair of nodes (a, b) is expressed by $\text{dis}(\mathbf{x}_a, \mathbf{x}_b)$ where \mathbf{x}_a are the node features of node a and $\text{dis}(\cdot)$ is a distance function (e.g., euclidean distance). This can be rewritten as Eq. (2) by substituting $w(a, b, u) = \text{dis}(\mathbf{x}_a, \mathbf{x}_u)$ and $h(a, b, u) = e_b^T$.

These examples demonstrate that the general formulation can indeed model many different LP factors including local and global structural information and feature proximity. We further show in Appendix A that Eq. (2) can model a variety of additional LP factors including RA [55], the pairwise encodings used in NCN/NCNC [45] and Neo-GNN [50]. However, fitting these methods into the formulation in Eq. (2) requires manually defining both $w(a, b, u)$ and $h(a, b, u)$. This constrains the information represented by $s(a, b)$ based on the choice of design. Motivated by this, in the next section we introduce our method that does not rely on a handcrafting both $w(a, b, u)$ and $h(a, b, u)$.

3.2 Modeling Pairwise Encodings via Attention

In Section 3.1, we introduced a general formulation for pairwise encodings in Eq. (2), which is able to capture a variety of different LP factors. However, it requires manually defining both terms in the equation. This limits our ability to customize the pairwise information to each target link. As such, we further aim to move beyond a one-size-fits-all pairwise encoding, and enable the model

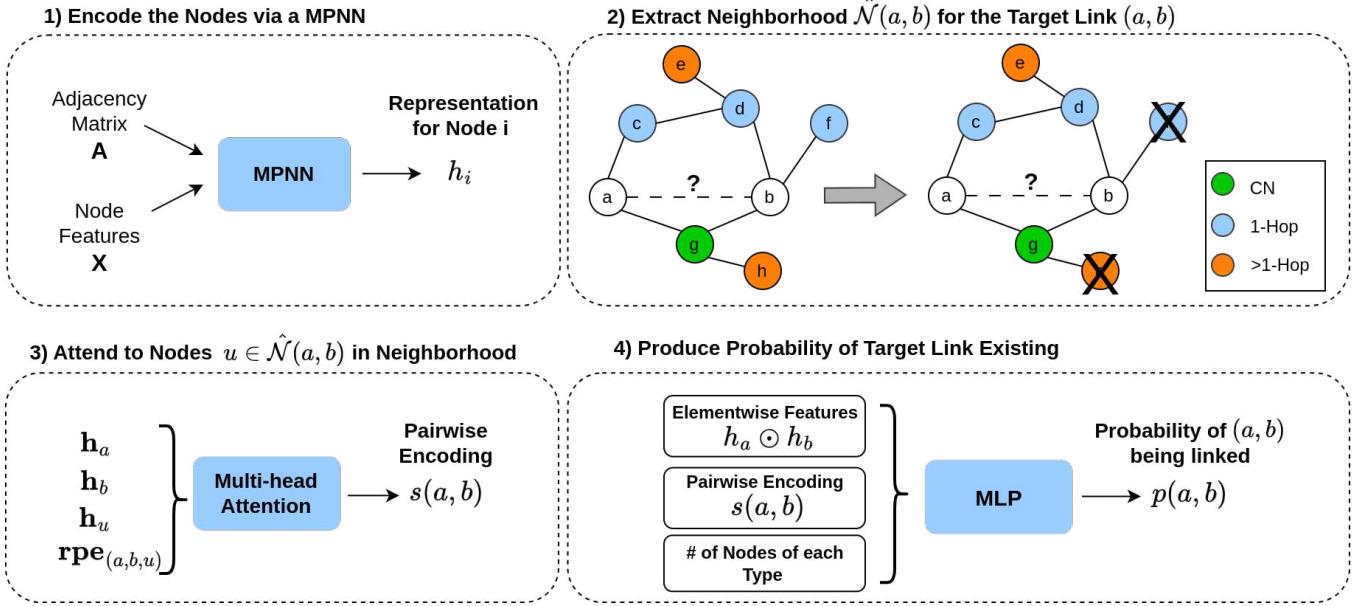


Figure 2: An overview of LPFormer. (1) Encode the nodes via a MPNN. (2) For a given target link, we determine which nodes to attend to ($\hat{\mathcal{N}}(a, b)$) via the PPR-based thresholding technique in Eq. (10). (3) The pairwise encoding is computed by attending to each node, $u \in \hat{\mathcal{N}}(a, b)$ using the feature and relative positional encoding $rpe_{(a,b,u)}$. (4) The pairwise encoding, node representations, and counts of different node types are concatenated and used to compute the final probability of the target link existing.

to produce customized pairwise encoding for each target link. This allows the model to handle more realistic graphs that often contain multiple prominent LP factors for different target links as shown in [32].

In particular, we consider the following question: *How can we model Eq (2) such that it can customize the used LP factors to each target link?* We consider parameterizing both $w(a, b, u)$ and $h(a, b, u)$. This allows us to learn how to personalize them to each target link. To achieve this, we leverage softmax attention [4]. This is due to its ability to dynamically learn the relevance of different nodes to the target link. As such, for multiple target links, it can emphasize the contributions of different nodes, thereby flexibly modeling different LP factors. We note that since the attention is between different sequences (i.e., a target link and nodes), it can be considered a form of cross attention [44].

To enhance the adaptability of the pairwise encoding for various links, it is essential to incorporate various types of information. This allows the attention mechanism to discern and prioritize relevant information for each target link, facilitating the effective modeling of diverse LP factors. In particular, we consider two types of information. The first is the **feature information**. This includes the feature representation of both nodes in the target link and the node being attended to. The node features are included due to their role in link formation and relationship to structural information [36]. Second, we consider the **relative positional information**. The relative positional information reflects the relative position in the graph of a node u to the target link (a, b) in the local and global

structural context. Due to the importance of local and global structural information [15, 20], it is vital to properly encode both. By including both the structural and feature information, we are able to cover the space of potential LP factors (see Section 2.1).

We denote the feature representation of a node u as h_u and the relative positional encoding (RPE) as $rpe_{(a,b,u)}$. The node importance $w(a, b, u)$ is modeled via attention as follows:

$$\begin{aligned} \tilde{w}(a, b, u) &= \phi(h_a, h_b, h_u, rpe_{(a,b,u)}), \\ w(a, b, u) &= \frac{\exp(\tilde{w}(a, b, u))}{\sum_{v \in \mathcal{V}(a,b)} \exp(\tilde{w}(a, b, v))}, \end{aligned} \quad (4)$$

where $\mathcal{V}(a, b) = \mathcal{V} \setminus \{a, b\}$. The attention weight $w(a, b, u)$ can be considered as the impact of a node u on (a, b) relative to all nodes in \mathcal{G} . This allows the model to emphasize different LP factors for each target link. The node encoding $h(a, b, u)$ includes the features of node u in conjunction with the RPE and is defined as:

$$h(a, b, u) = W [h_u \parallel rpe_{(a,b,u)}]. \quad (5)$$

By substituting Eq. (4) and Eq. (5) into Eq. (2) we can compute the pairwise information $s(a, b)$. We further define $\phi(\cdot)$ in Eq. (4) as the GATv2 [8] attention mechanism. The detailed formulation is given in Appendix C. The feature representations h_i are computed via a MPNN. We use GCN [26] in this work. However, it is unclear how to properly encode the RPE of a node u relative to (a, b) , $rpe_{(a,b,u)}$. We aim to design the RPE to capture both the local and global structural relationship between the node and target link while also

being efficient to calculate. In the next section, we discuss our solution for modeling $\text{rpe}_{(a,b,u)}$.

3.3 PPR-Based Relative Positional Encodings

In this section, we introduce our strategy for computing the RPE of a node u relative to a target link (a, b) . Intuitively, we want the RPE to reflect the positional relationship between u and (a, b) such that different types of information (i.e., local vs. global) are encoded differently. Using Figure 1 as an example, since node 3 is a CN of (source, 5) we expect it to have a much different relationship to the target link than node 6, which is a 2-hop neighbor of both nodes. An enticing option is to use the double radius node labeling (DRNL) trick introduced by Zhang and Chen [51]. However, Chamberlain et al. [9] have shown it to be prohibitively expensive to calculate for larger graphs. Furthermore, existing RPEs are typically infeasible to calculate on larger graphs as they often rely on pairwise distances or the eigenvectors of the Laplacian [41].

As such, we seek an RPE that can both distinguish the relationship of different nodes to the target link while also being efficient to calculate. To motivate our RPE design, we draw inspiration from the following Proposition.

Proposition 1. Consider a target link (a, b) and a node $u \in \mathcal{V} \setminus \{a, b\}$. The PPR [7] score of a root node i and target node j with teleportation probability α is denoted by $\text{ppr}(i, j)$. Let $r_a^k(u)$ be the probability of a walk of length k beginning at node a and terminating at u . We define $r_{a,b}^k(u) := r_a^k(u) + r_b^k(u)$. We also define a weight $\gamma^k := \alpha(1 - \alpha)^k$ for all walks of length k . The PPR scores, $\text{ppr}(a, u)$ and $\text{ppr}(b, u)$, along with the random walk probabilities of disparate lengths, are interconnected through the following relationship.

$$\Gamma(a, b, u) = \text{ppr}(a, u) + \text{ppr}(b, u) = \sum_{k=0}^{\infty} \gamma^k r_{a,b}^k(u). \quad (6)$$

The detailed proof is given in Appendix B. From Proposition 1, we can make the following observations: (1) The PPR scores encode the weighted sum of the probabilities of different length random walks connecting two nodes. (2) Walks of shorter length are given higher importance, as evidenced by the dampening factor $\gamma^k = \alpha(1 - \alpha)^k$ which decays with the increase in k . These observations imply that **– a larger value of $\Gamma(a, b, u)$ correlates with the existence of many shorter walks connecting node u to the both nodes in the target link (a, b) .**

Therefore, the PPR scores can be used as an intuitive and useful method to understand the structural relationship between node u and both nodes in the target link (a, b) . If both scores, $\text{ppr}(a, u)$ and $\text{ppr}(b, u)$, are high, there exists a high probability that many shorter walks connect u to both nodes in the target link. This implies that node u has a stronger impact on the nodes in the target link. On the other hand, if both PPR scores are low, there is likely very little relationship between u and the target link. This allows for a convenient way of differentiating how a node structurally relates to the target link. Furthermore, we note that the PPR matrix can be efficiently pre-computed using the algorithm introduced by Andersen et al. [3], allowing for easy computation and use.

Following this idea, to calculate the RPE of a node u , we use the PPR scores of a node u relative to both nodes in the target link (a, b) .

Instead of considering the sum of PPR scores as in Proposition 1, we further parameterize $\Gamma(\cdot)$ via an MLP,

$$\text{rpe}_{(a,b,u)} = \text{MLP}(\text{ppr}(a, u), \text{ppr}(b, u)). \quad (7)$$

By introducing learnable parameters to $\Gamma(\cdot)$, it allows for the model learn the importance of individual PPR scores and how they interact with each other. To ensure that Eq. (7) is invariant to the order of the nodes in the target link, i.e., (a, b) and (b, a) , we further set the RPE to be equal to the summation of the representations given by both (a, b) and (b, a) :

$$\overline{\text{rpe}}_{(a,b,u)} = \text{rpe}_{(a,b,u)} + \text{rpe}_{(b,a,u)}. \quad (8)$$

However, a concern with Eq. (8) is that it is not guaranteed to be able to distinguish certain types of nodes from each other. For example, it is necessary to clearly distinguish CNs from other nodes due to their important role in link formation [38]. To overcome this issue, we fit three separate MLPs for when u is a: CN of (a, b) , a 1-hop neighbor of either a and b , and a >1 -hop neighbor of both a and b . This ensures that we can properly distinguish between these three types of nodes. We verify the effectiveness of this design in Section 4.4. Lastly, we note that while other work [29, 34] has considered the use of random-walk based positional encodings, they are only designed for use on the node-level and are unable to be used for link-level tasks like LP.

3.4 Efficiently Attending to the Graph Context

The proposed attention mechanism in Section 3.2 attends to all nodes in the graph, sans those in the link itself. This makes it difficult to scale to large graphs. Motivated by selective [33] and sparse [13] attention, we opt to attend to only a small portion of the nodes.

At a high level, we are interested in determining a subset of nodes $\hat{\mathcal{N}}(a, b) \in \mathcal{V}$ to attend to for the target link (a, b) . Our goal is to choose the set of nodes $\hat{\mathcal{N}}(a, b)$ such that they are **(a)** few in number to improve scalability and **(b)** provide important contextual information to the pair (a, b) to best learn the pairwise information. This can be achieved by only considering all nodes where the importance of the node u to the target link (a, b) is considered high. Formally, we can write this as the following where $\mathcal{I}(a, b, u)$ is a function that denotes the importance of a node u to the target link (a, b) :

$$\hat{\mathcal{N}}(a, b) = \{u \in \mathcal{V} \setminus \{a, b\} \mid \mathcal{I}(a, b, u) > \eta\}. \quad (9)$$

The threshold η allows us to distinguish those nodes that are sufficiently important to the target link. This allows for a simple and efficient way of determining the set $\hat{\mathcal{N}}(a, b)$. However, *what do we use to model the importance $\mathcal{I}(a, b, u)$?* For ease of optimization and better efficiency, we avoid parameterizing the function $\mathcal{I}(a, b, u)$. Instead, we want to choose a metric such that can properly serve as a proxy for the importance of a node u to (a, b) while also being concentrated in a small subset of nodes. Such a metric will allow Eq. (9) to choose a small but influential set of nodes to attend to.

A measure that satisfies both criteria is Personalized Pagerank (PPR) [7]. In Section 3.3 we discussed that the PPR score can serve as a good tool to model the influence of a one node on another. Furthermore, existing work [3, 17, 37] shows that the PPR scores tend to be highly localized in a small subset of nodes. Therefore by

making $\mathcal{I}(a, b, u)$ contingent on the PPR scores of (a, u) and (b, u) we can extract a small but important set of nodes to attend to for the target link.

Following this idea, for a target link (a, b) , we keep all nodes whose PPR score is above some threshold η relative to both nodes in the target link. As such, we only keep a node u if it is related in some capacity to at least one of the nodes in the target link. Similarly to Section 3.3, we treat CN, 1-Hop, and >1 -Hop nodes differently by applying a different threshold for them. The filtered node set for each category of nodes is given by:

$$\hat{\mathcal{N}}_{(a,b)}^\pi = \{u \in \mathcal{N}_{(a,b)}^\pi \mid \text{ppr}(a, u) > \eta^\pi, \text{ppr}(b, u) > \eta^\pi\}, \quad (10)$$

where $\hat{\mathcal{N}}_{(a,b)}^\pi$ is the filtered node set for all nodes of the type $\pi \in \{\text{CN}, \text{1-Hop}, \text{}>1\text{-Hop}\}$ and η^π is the corresponding PPR threshold. We note that while other work [6, 49] has used PPR to filter the nodes on the *node-level*, no existing work has done so on the *link-level*.

We corroborate this design by demonstrating that LPFormer can achieve SOTA performance in LP (Section 4.2) while achieving a faster runtime than the second-best method, NCNC [45], on denser graphs (Section 4.8). This is despite the fact that LPFormer can attend to a wider variety of nodes. We further show in Section 4.5 that the performance is stable with regards to the values of η chosen, allowing us to easily choose a proper threshold on any dataset.

3.5 LPFormer

We now define the overall framework – LPFormer. The overall procedure is given in Figure 2: (1) We first learn node representations from the input adjacency and node features via an MPNN. We note that this step is agnostic to the target link. (2) For a target link (a, b) we extract the nodes to attend to, i.e. $\hat{\mathcal{N}}(a, b)$. This is done via the PPR thresholding technique defined in Section 3.4. (3) We apply L layers of attention, using the mechanism defined in Section 3.2. The output is the pairwise encoding $s(a, b)$. (4) We generate the prediction of the target link using three types of information: the element-wise product of the node representation, the pairwise encoding, and the number of CN, 1-Hop, and >1 -Hop nodes identified by Eq. (10). The score function is given by:

$$p(a, b) = \sigma \left(\text{MLP} \left(\mathbf{h}_a \odot \mathbf{h}_b \parallel s(a, b) \parallel |\hat{\mathcal{N}}_{(a,b)}^{\text{CN}}| \parallel |\hat{\mathcal{N}}_{(a,b)}^1| \parallel |\hat{\mathcal{N}}_{(a,b)}^{>1}| \right) \right) \quad (11)$$

We demonstrate in Section 4.4 that the inclusion of the node counts is helpful, as it provides complementary information to the pairwise encoding.

4 EXPERIMENTS

In this section, we conduct extensive experiments to validate the effectiveness of LPFormer. Specifically, we attempt to answer the following questions: **(RQ1)** Can LPFormer consistently outperform baseline methods on a variety of different benchmark datasets? **(RQ2)** Is LPFormer able to model a variety of different LP factors? **(RQ3)** Can LPFormer be run efficiently on large dense graphs? We further conduct studies ablating each component of our model and analyzing the effect of the PPR-based threshold on performance.

4.1 Experimental Settings

Datasets. We include Cora, Citeseer, and Pubmed [47] and ogbl-collab, ogbl-ppa, ogbl-ddi, and ogbl-citation2 [19]. Furthermore, for Cora, Citeseer, and Pubmed we experiment under a single fixed split as in Li et al. [28]. The detailed statistics for each dataset are shown in Table 1.

Baseline Models. We compare LPFormer against a wide variety of baselines including: CN [38], AA [2], RA [55], GCN [26], SAGE [18], GAE [25], SEAL [51], NBFNet [56], Neo-GNN [50], BUDDY [9], NCN [45], and NCNC [45]. Results on Cora, Citeseer, and Pubmed are taken from Li et al. [28]. Results for the heuristic methods are from Hu et al. [19]. All other results are either from their respective study or Chamberlain et al. [9].

Hyperparameters: The learning rate is tuned from $\{1e^{-3}, 5e^{-3}\}$, the decay from $\{0.95, 0.975, 1\}$, and the dropout from $[0, 0.7]$, and the weight decay from $\{0, 1e^{-4}, 1e^{-7}\}$. The size of the hidden dimension is set to 64 for ogbl-ppa and ogbl-citation2, 128 for Cora, Pubmed, and ogbl-collab, and 256 for Citeseer. Lastly, the PPR threshold is tuned from $\{1e^{-2}, 1e^{-3}, 1e^{-4}\}$.

Evaluation Metrics. Each positive target link is evaluated against a set of given negative links. The rank of the positive link among the negatives is used to evaluate performance. The two types of metrics that are used to evaluate this ranking are Hits@K and MRR. For the OGB datasets we use the metric used in the original study. This includes Hits@50 for ogbl-collab, Hits@100 for ogbl-ppa and MRR for ogbl-citation2. For Cora, Citeseer, Pubmed we follow Li et al. [28] and use MRR. Lastly, the same set of negative links is used for all positive links except on ogbl-citation2, where [19] provides a customized set of 1000 negatives for each individual positive link.

4.2 Main Results

We present the results of LPFormer compared with baselines on multiple benchmark datasets. Note that we omit ogbl-ddi from the main results due to recent issues discovered by Li et al. [28]. The results are shown in Table 2. We observe that LPFormer can achieve SOTA performance on 5/6 datasets, significantly outperforming other baselines. Moreover, LPFormer is also the most consistent of all the methods, achieving strong performance on all datasets. This is as opposed to previous SOTA methods, NCNC and BUDDY, which tend to struggle on Cora and Pubmed. We attribute the consistency of LPFormer to the flexibility of our model, allowing it to customize the LP factors needed to each link and dataset.

4.3 Performance by LP Factor

In this section, we measure the ability of LPFormer to capture a variety of different LP factors. To measure this, we identify all positive target links **when there is only one dominant LP factor**. For example, one group would contain all target links where the only dominant factor is the local structural information. We focus on links that correspond to one of the three groups identified in [32]: local structural information, global structural information, and feature proximity.

We identify these groups by using popular heuristics as proxies for each factor. For local structural information, we use CNs [38], for global structural information we use PPR [7] as it's the most computationally efficient of all global methods, and for feature

Table 1: Dataset statistics. The split ratio is the % of samples for train/validation/test.

	Cora	Citeseer	Pubmed	ogbl-collab	ogbl-ddi	ogbl-ppa	ogbl-citation2
#Nodes	2,708	3,327	18,717	235,868	4,267	576,289	2,927,963
#Edges	5,278	4,676	44,327	1,285,465	1,334,889	30,326,273	30,561,187
Split Ratio	85/5/10	85/5/10	85/5/10	92/4/4	80/10/10	70/20/10	98/1/1

Table 2: Results on benchmark datasets. OOM is an out of memory error. Colored are the results ranked **first, **second**, and **third**.**

Metric	Cora	Citeseer	Pubmed	ogbl-collab	ogbl-ppa	ogbl-citation2	Mean Rank
	MRR	MRR	MRR	H@50	H@100	MRR	
CN	20.99 \pm 0.00	28.34 \pm 0.00	14.02 \pm 0.00	56.44 \pm 0.00	27.65 \pm 0.00	51.47 \pm 0.00	11.0
AA	31.87 \pm 0.00	29.37 \pm 0.00	16.66 \pm 0.00	64.35 \pm 0.00	32.45 \pm 0.00	51.89 \pm 0.00	8.5
RA	30.79 \pm 0.00	27.61 \pm 0.00	15.63 \pm 0.00	64.00 \pm 0.00	49.33 \pm 0.00	51.98 \pm 0.00	8.7
GCN	32.50 \pm 6.87	50.01 \pm 6.04	19.94 \pm 4.24	44.75 \pm 1.07	18.67 \pm 1.32	84.74 \pm 0.21	8.0
SAGE	37.83\pm7.75	47.84 \pm 6.39	22.74 \pm 5.47	48.10 \pm 0.81	16.55 \pm 2.40	82.60 \pm 0.36	7.7
GAE	29.98 \pm 3.21	63.33\pm3.14	16.67 \pm 0.19	OOM	OOM	OOM	NA
SEAL	26.69 \pm 5.89	39.36 \pm 4.99	38.06\pm5.18	64.74 \pm 0.43	48.80 \pm 3.16	87.67 \pm 0.32	6.2
NBFNet	37.69\pm3.97	38.17 \pm 3.06	44.73\pm2.12	OOM	OOM	OOM	NA
Neo-GNN	22.65 \pm 2.60	53.97 \pm 5.88	31.45 \pm 3.17	57.52 \pm 0.37	49.13 \pm 0.60	87.26 \pm 0.84	7.0
BUDDY	26.40 \pm 4.40	59.48 \pm 8.96	23.98 \pm 5.11	65.94\pm0.58	49.85 \pm 0.20	87.56 \pm 0.11	5.7
NCN	32.93 \pm 3.80	54.97 \pm 6.03	35.65 \pm 4.60	64.76 \pm 0.87	61.19\pm0.85	88.09\pm0.06	3.8
NCNC	29.01 \pm 3.83	64.03\pm3.67	25.70 \pm 4.48	66.61\pm0.71	61.42\pm0.73	89.12\pm0.40	3.8
LPFormer	39.42\pm5.78	65.42\pm4.65	40.17\pm1.92	68.14\pm0.51	63.32\pm0.63	89.81\pm0.13	1.2

proximity, we use the cosine similarity of the features. Using these heuristics, we determine if only one factor is dominant by comparing the relative score of each heuristic. This is done by first computing the score for each factor i for the target link $(a, b) - s^i(a, b)$. For each factor, we then compute the score corresponding to the p -th percentile among all links, \hat{s}^i . We choose a larger value of p (i.e. 90%) such that a score $\geq \hat{s}^i$ indicates that a significant amount of pairwise information exists for that factor. For a single target link, we then compare the score of each factor $s^i(a, b)$ to \hat{s}^i . If $s^i(a, b) \geq \hat{s}^i$ is true **for only one factor**, this implies that the score for only one factor is “high”. Therefore there is a notable amount of pairwise information existing for only one factor for the link (a, b) . This ensures that only one factor is strongly expressed. If this is true, we then assign the target link (a, b) to factor i . Please see Appendix D.2 for a more detailed explanation.

We demonstrate the results on Cora, Citeseer, and ogbl-collab in Figure 3. We observe that LPFormer typically performs best for each individual LP factor on all datasets. Furthermore, it is also the most consistently well-performing on each factor as compared to other methods. For example, on Cora the other methods struggle for links that correspond to the feature proximity factor. LPFormer, on the other hand, is able to significantly outperform them on those target links, performing around 33% better than the second best method. Lastly, we note that most methods tend to perform well on the links corresponding to the global factor, even if they don’t explicitly model such information. This is caused by a strong correlation that tends to exist between local and global structural

information, often resulting in considerable overlap between both factors [32]. These results show that LPFormer can indeed adapt to multiple types of LP factors, as it can consistently perform well on samples belonging to a variety of different LP factors. Additional results are given in Appendix E.

4.4 Ablation Study

We further include an ablation study to verify the effectiveness of the proposed components in LPFormer. In particular, we introduce 6 variants of LPFormer. (a) **w/o Learnable Att**: No attention is learned. As such, we set all attention weights to 1 and remove the RPE. (b) **w/o Features in Att**: We remove the node feature information from the attention mechanism. (c) **w/o RPE in Att**: We remove the RPE from the attention mechanism. (d) **w/o PPR RPE**: We replace the PPR-based RPE with a learnable embedding for each of CN, 1-Hop, and >1-Hop nodes. (e) **w/o PPR RPE by Node Type**: We don’t fit a separate function for each node type when determining the PPR RPE (see Section 3.3). Instead we use one for all nodes. (f) **w/o Counts**: We remove the counts of different nodes from the scoring function.

The results are shown in Table 3. We include ogbl-collab, ogbl-ppa, and Citeseer. We observe that ablating a component always decreases the performance. However, the magnitude of the decrease is dataset-dependent. For example, on ogbl-collab, ablating the feature information in the attention marginally affects the performance. However, on ogbl-ppa and Citeseer, removing the feature information results in a large decrease in performance. On the other hand,

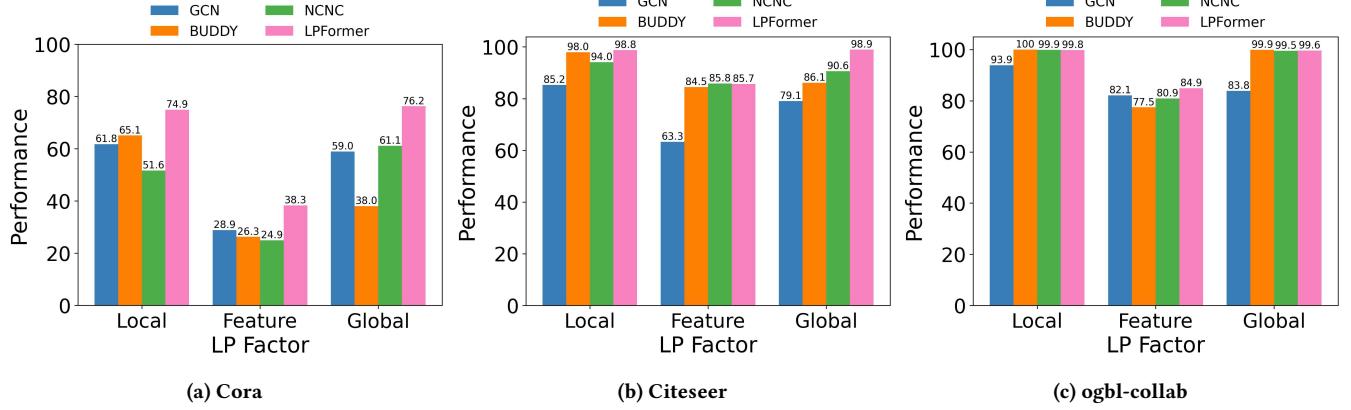


Figure 3: Performance on links that contain one dominant LP factor. Results are on (a) Cora, (b) Citeseer, and (c) ogbl-collab.

while removing learnable attention results in a modest decrease on ogbl-ppa, for the other two datasets we see a large drop. This highlights the importance of each component of our framework, as they are each necessary for consistently strong performance across multiple datasets.

Table 3: Ablation Study on LPFormer

Method	ogbl-collab	ogbl-ppa	Citeseer
w/o Learnable Att	65.05 \pm 0.50	62.77 \pm 1.03	56.23 \pm 1.75
w/o Features in Att	68.04 \pm 0.79	56.98 \pm 1.55	53.40 \pm 9.30
w/o RPE in Att	65.26 \pm 0.56	61.20 \pm 0.69	56.70 \pm 3.79
w/o PPR RPE	67.09 \pm 0.51	61.91 \pm 1.22	51.96 \pm 15.2
w/o PPR RPE by Node Type	67.95 \pm 0.54	62.92 \pm 1.06	57.40 \pm 5.71
w/o Counts	67.75 \pm 0.41	44.37 \pm 1.89	54.39 \pm 5.30
LPFormer	68.14 \pm 0.51	63.32 \pm 0.63	65.42 \pm 4.65

Table 4: Effect of Varying the PPR Thresholds

Threshold	ogbl-collab		ogbl-citation2	
	1-Hop	>1-Hop	1-Hop	>1-Hop
1e-4	68.24 \pm 0.25	67.73 \pm 0.65	89.81 \pm 0.13	89.14 \pm 0.22
1e-2	67.60 \pm 0.31	68.24 \pm 0.25	89.49 \pm 0.18	89.81 \pm 0.13
1	67.08 \pm 0.65	68.14 \pm 0.51	89.49 \pm 0.16	89.26 \pm 0.39

4.5 Effect of the PPR Thresholds

We examine the effect of varying the PPR threshold for both 1-Hop and >1-Hop nodes as described in Eq. (10). The results for ogbl-collab and ogbl-citation2 are shown in Table 4. When varying the 1-Hop threshold, we fix the value of the >1-Hop threshold to 1e-2 for both datasets. When varying the >1-Hop threshold, we fix the value of the 1-Hop threshold to 1e-4 for both datasets.

We can observe that modifying the threshold has little effect on the underlying performance of the model. For both datasets, a value of 1e-2 works well for the >1-Hop threshold and 1e-4 works well for the 1-Hop threshold. We typically find that setting both

values to 1e-2 provides a good trade-off between performance and efficiency.

4.6 Performance on HeaRT Setting

We further test the performance of our method on the HeaRT [28] evaluation setting, which considers a more realistic and difficult evaluation setting for link prediction. This is done by introducing a much harder and more realistic set of negative samples during evaluation. Li et al. [28] observe that this results in a large decrease in performance on all datasets. Furthermore, compared to the original evaluation setting, MPNNs designed specifically for link prediction are often outperformed by heuristics or other MPNNs.

The full results can be found in Table 5. We observe that LPFormer performs considerably better than all other models. For instance, the mean rank of LPFormer is 3.1x better than the 2nd best-performing model, NCN. This indeed shows the advantage of LPFormer, as it can consistently achieve extraordinary performance across all datasets under the much more challenging HeaRT evaluation setting. This is as opposed to other LP-specific methods that often perform similarly to standard MPNN methods.

4.7 Performance on Heterophilic Datasets

In this section we evaluate LPFormer on multiple heterophilic datasets. Heterophily refers to the tendency of dissimilar nodes to be connected. This is as opposed to homophily, in which nodes with similar attributes are more likely to be connected. Since most graphs used for benchmark datasets tend to contain homophilic patterns, heterophilic graphs present an interesting challenge regarding the effectiveness of graph-based methods. For a more detailed discussion on heterophilic graphs, please see [31].

We test on two prominent heterophilic datasets, Squirrel and Chameleon [42]. The statistics for each are in Table 6. We limit our comparison to those LP methods that tend to achieve the best results, including GCN, BUDDY, and NCNC. In Table 7, we report the MRR over five random seeds. Note that we test under the **original evaluation setting** and not HeaRT. We observe that LPFormer can achieve a large increase over other methods, with a 14% and 9% increase in performance on Squirrel and Chameleon, respectively.

Table 5: Results (MRR) under HeaRT. Highlighted are the results ranked **first, **second**, and **third**.**

Models	Cora	Citeseer	Pubmed	ogbl-collab	ogbl-ddi	ogbl-ppa	ogbl-citation2	Mean Rank
CN	9.78	8.42	2.28	4.20	6.71	25.70	17.11	11.1
AA	11.91	10.82	2.63	5.07	6.97	26.85	17.83	9.6
RA	11.81	10.84	2.47	6.29	8.70	28.34	17.79	8.1
GCN	16.61 ± 0.30	21.09 ± 0.88	7.13 ± 0.27	6.09 ± 0.38	13.46 ± 0.34	26.94 ± 0.48	19.98 ± 0.35	4.7
SAGE	14.74 ± 0.69	21.09 ± 1.15	9.40 ± 0.70	5.53 ± 0.5	12.60 ± 0.72	27.27 ± 0.30	22.05 ± 0.12	4.7
GAE	18.32 ± 0.41	25.25 ± 0.82	5.27 ± 0.25	OOM	3.49 ± 1.73	OOM	OOM	NA
SEAL	10.67 ± 3.46	13.16 ± 1.66	5.88 ± 0.53	6.43 ± 0.32	9.99 ± 0.90	29.71 ± 0.71	20.60 ± 1.28	6.4
NBFNet	13.56 ± 0.58	14.29 ± 0.80	>24h	OOM	>24h	OOM	OOM	NA
BUDDY	13.71 ± 0.59	22.84 ± 0.36	7.56 ± 0.18	5.67 ± 0.36	12.43 ± 0.50	27.70 ± 0.33	19.17 ± 0.20	5.9
Neo-GNN	13.95 ± 0.39	17.34 ± 0.84	7.74 ± 0.30	5.23 ± 0.9	10.86 ± 2.16	21.68 ± 1.14	16.12 ± 0.25	7.4
NCN	14.66 ± 0.95	28.65 ± 1.21	5.84 ± 0.22	5.09 ± 0.38	12.86 ± 0.78	35.06 ± 0.26	23.35 ± 0.28	4.4
NCNC	14.98 ± 1.00	24.10 ± 0.65	8.58 ± 0.59	4.73 ± 0.86	>24h	33.52 ± 0.26	19.61 ± 0.54	4.8
LPFormer	16.80 ± 0.52	26.34 ± 0.67	9.99 ± 0.52	7.62 ± 0.26	13.20 ± 0.54	40.25 ± 0.24	24.70 ± 0.55	1.4

These results indicate the superior ability of LPFormer to accurately model LP on heterophilic graphs, as compared to other methods.

Table 6: Heterophilic Dataset Statistics.

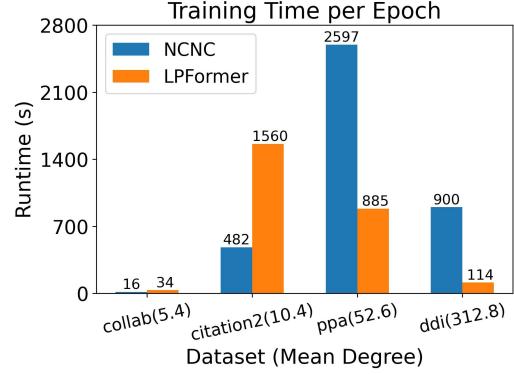
	Squirrel	Chameleon
#Nodes	5201	2277
#Edges	198,353	31,371
Split Ratio	85/5/10	85/5/10

Table 7: Results on Heterophilic Datasets.

Method	Squirrel	Chameleon
GCN	22.77 ± 4.54	20.74 ± 8.08
BUDDY	9.69 ± 0.99	6.30 ± 2.40
NCNC	32.37 ± 5.46	26.24 ± 3.37
LPFormer	36.77 ± 2.77	28.61 ± 6.68
% Improvement	14%	9%

4.8 Runtime Analysis

In this section, we compare the runtime of LPFormer against NCNC, which is the strongest performing baseline. The results are shown in Figure 4 on all four OGB datasets. We further include the mean degree of each dataset in parentheses. We observe that LPFormer shines on denser datasets, taking significantly less time to train one epoch. This is despite that LPFormer can attend to nodes beyond the 1-hop radius of the target link. This underscores the importance of the PPR thresholding technique introduced in Section 3.4, as it allows for efficient attention to a wider variety of nodes. Lastly, we note that LPFormer struggles on the ogbl-citation2 dataset due to the large number of nodes in the dataset (i.e., 2,927,963), which requires the sparse PPR matrix to be quite large. For future work we plan on exploring pre-computing the necessary PPR scores as an efficient pre-processing step, thereby removing the need to store the costly PPR matrix.

**Figure 4: Comparison of training time of 1 epoch between LPFormer and NCNC. The mean degree is in parentheses.**

5 CONCLUSION

In this paper we introduce a new framework, LPFormer, that aims to integrate a wider variety of pairwise information for link prediction. LPFormer does this via a specially designed graph transformer, which adaptively considers how a node pair relate to each other in the context of the graph. Extensive experiments demonstrate that LPFormer can achieve SOTA performance on a wide variety of benchmark datasets while retaining efficiency. We further demonstrate LPFormer’s supremacy at modeling multiple types of LP factors. For future work, we plan on exploring other methods of incorporating multiple LP factors with an emphasis on global structural information.

ACKNOWLEDGMENTS

This research is supported by the National Science Foundation (NSF) under grant numbers CNS 2246050, IIS1845081, IIS2212032, IIS2212144, IOS2107215, DUE 2234015, DRL 2025244 and IOS2035472, the Army Research Office (ARO) under grant number W911NF-21-1-0198, the Home Depot, Cisco Systems Inc, Amazon Faculty Award, Johnson&Johnson, JP Morgan Faculty Award and SNAP.

REFERENCES

[1] Khushnood Abbas, Alireza Abbasi, Shi Dong, Ling Niu, Laihang Yu, Bolun Chen, Shi-Min Cai, and Qambar Hasan. 2021. Application of network link prediction in drug discovery. *BMC bioinformatics* 22 (2021), 1–21.

[2] Lada A Adamic and Eytan Adar. 2003. Friends and neighbors on the web. *Social networks* 25, 3 (2003), 211–230.

[3] Reid Andersen, Fan Chung, and Kevin Lang. 2006. Local graph partitioning using pagerank vectors. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. IEEE, 475–486.

[4] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.

[5] Albert-László Barabási, Hawoong Jeong, Zoltan Néda, Erzsébet Ravasz, Andras Schubert, and Tamas Vicsek. 2002. Evolution of the social network of scientific collaborations. *Physica A: Statistical mechanics and its applications* 311, 3-4 (2002), 590–614.

[6] Aleksandar Bojchevski, Johannes Gasteiger, Bryan Perozzi, Amol Kapoor, Martin Blaïs, Benedek Rözemberczki, Michal Lukasik, and Stephan Günnemann. 2020. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2464–2473.

[7] Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems* 30, 1-7 (1998), 107–117.

[8] Shaked Brody, Uri Alon, and Eran Yahav. 2022. How Attentive are Graph Attention Networks? In *International Conference on Learning Representations*. <https://openreview.net/forum?id=F72ximxs7C1>

[9] Benjamin Paul Chamberlain, Sergey Shirokov, Emanuele Rossi, Fabrizio Frasca, Thomas Markovich, Niki Hammerla, Michael M Bronstein, and Max Hansmire. 2022. Graph Neural Networks for Link Prediction with Subgraph Sketching. *arXiv preprint arXiv:2209.15486* (2022).

[10] Jinsong Chen, Kaiyuan Gao, Gaichao Li, and Kun He. 2022. NAGphomer: A tokenized graph transformer for node classification in large graphs. In *The Eleventh International Conference on Learning Representations*.

[11] Sanxing Chen, Xiaodong Liu, Jianfeng Gao, Jian Jiao, Ruofei Zhang, and Yangfeng Ji. 2021. Hitter: Hierarchical Transformers for Knowledge Graph Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 10395–10407.

[12] Fan Chung. 2007. The heat kernel as the pagerank of a graph. *Proceedings of the National Academy of Sciences* 104, 50 (2007), 19735–19740.

[13] Gonçalo M Correia, Vlad Niculae, and André FT Martins. 2019. Adaptively Sparse Transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 2174–2184.

[14] Nur Nasuha Daud, Siti Hafizah Ab Hamid, Muntadher Saadoon, Firdaus Sahran, and Nor Badrul Anuar. 2020. Applications of link prediction in social networks: A review. *Journal of Network and Computer Applications* 166 (2020), 102716.

[15] Yuxiao Dong, Reid A Johnson, Jian Xu, and Nitesh V Chawla. 2017. Structural diversity and homophily: A study across more than one hundred big networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 807–816.

[16] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*. PMLR, 1263–1272.

[17] David F Gleich, Kyle Kloster, and Huda Nassar. 2015. Localization in seeded pagerank. *arXiv preprint arXiv:1509.00016* (2015).

[18] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).

[19] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems* 33 (2020), 22118–22133.

[20] Hong Huang, Jie Tang, Lu Liu, JarDer Luo, and Xiaoming Fu. 2015. Triadic closure pattern analysis and prediction in social networks. *IEEE Transactions on Knowledge and Data Engineering* 27, 12 (2015), 3374–3389.

[21] Zan Huang, Xin Li, and Hsinchun Chen. 2005. Link prediction approach to collaborative filtering. In *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*. 141–142.

[22] Leo Katz. 1953. A new status index derived from sociometric analysis. *Psychometrika* 18, 1 (1953), 39–43.

[23] Jinwoo Kim, Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. 2022. Pure transformers are powerful graph learners. *Advances in Neural Information Processing Systems* 35 (2022), 14582–14595.

[24] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).

[25] Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016).

[26] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.

[27] Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou. 2021. Rethinking graph transformers with spectral attention. *Advances in Neural Information Processing Systems* 34 (2021), 21618–21629.

[28] Juanhui Li, Harry Shomer, Haitao Mao, Shenglai Zeng, Yao Ma, Neil Shah, Jiliang Tang, and Dawei Yin. 2023. Evaluating Graph Neural Networks for Link Prediction: Current Pitfalls and New Benchmarking. *arXiv preprint arXiv:2306.10453* (2023).

[29] Pan Li, Yanbang Wang, Hongwei Wang, and Jure Leskovec. 2020. Distance encoding: Design provably more powerful neural networks for graph representation learning. *Advances in Neural Information Processing Systems* 33 (2020), 4465–4478.

[30] David Liben-Nowell and Jon Kleinberg. 2003. The link prediction problem for social networks. In *Proceedings of the twelfth international conference on Information and knowledge management*. 556–559.

[31] Haitao Mao, Zhikai Chen, Wei Jin, Haoyu Han, Yao Ma, Tong Zhao, Neil Shah, and Jiliang Tang. 2024. Demystifying Structural Disparity in Graph Neural Networks: Can One Size Fit All? *Advances in Neural Information Processing Systems* 36 (2024).

[32] Haitao Mao, Juanhui Li, Harry Shomer, Bingheng Li, Wensi Fan, Yao Ma, Tong Zhao, Neil Shah, and Jiliang Tang. 2023. Revisiting Link Prediction: A Data Perspective. *arXiv:2310.00793 [cs.SI]*

[33] Sameen Maruf, André FT Martins, and Gholamreza Haffari. 2019. Selective Attention for Context-aware Neural Machine Translation. In *Proceedings of NAACL-HLT*. 3092–3102.

[34] Grégoire Mialon, Dexiong Chen, Margot Selosse, and Julien Mairal. 2021. Graphit: Encoding graph structure in transformers. *arXiv preprint arXiv:2106.05667* (2021).

[35] Lui Müller, Mikhail Galkin, Christopher Morris, and Ladislav Rampášek. 2023. Attending to graph transformers. *arXiv preprint arXiv:2302.04181* (2023).

[36] Yohsuke Murase, Hang-Hyun Jo, János Török, János Kertész, and Kimmo Kaski. 2019. Structural transition in social networks: The role of homophily. *Scientific reports* 9, 1 (2019), 4310.

[37] Huda Nassar, Kyle Kloster, and David F Gleich. 2015. Strong Localization in Personalized PageRank Vectors. In *Proceedings of the 12th International Workshop on Algorithms and Models for the Web Graph-Volume 9479*. 190–202.

[38] Mark EJ Newman. 2001. Clustering and preferential attachment in growing networks. *Physical review E* 64, 2 (2001), 025102.

[39] Maximilian Nickel, Xueyan Jiang, and Volker Tresp. 2014. Reducing the rank in relational factorization models by including observable patterns. *Advances in Neural Information Processing Systems* 27 (2014).

[40] Vardaan Pahuja, Boshi Wang, Hugo Latapie, Jayanth Srinivasa, and Yu Su. 2023. A retrieve-and-read framework for knowledge graph link prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 1992–2002.

[41] Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. 2022. Recipe for a general, powerful, scalable graph transformer. *Advances in Neural Information Processing Systems* 35 (2022), 14501–14515.

[42] Benedek Rozemberczki, Carl Allen, and Rik Sarkar. 2021. Multi-scale attributed node embedding. *Journal of Complex Networks* 9, 2 (2021), cnab014.

[43] Balasubramanian Srinivasan and Bruno Ribeiro. 2019. On the Equivalence between Positional Node Embeddings and Structural Graph Representations. In *International Conference on Learning Representations*.

[44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[45] Xiyuan Wang, Haotong Yang, and Muhan Zhang. 2023. Neural Common Neighbor with Completion for Link Prediction. *arXiv preprint arXiv:2302.00890* (2023).

[46] Qitian Wu, Wentao Zhao, Zhenan Li, David P Wipf, and Junchi Yan. 2022. Nodeformer: A scalable graph structure learning transformer for node classification. *Advances in Neural Information Processing Systems* 35 (2022), 27387–27401.

[47] Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.

[48] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems* 34 (2021), 28877–28888.

[49] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 974–983.

[50] Seongjun Yun, Seoyoon Kim, Junhyun Lee, Jaewoo Kang, and Hyunwoo J Kim. 2021. Neo-gnns: Neighborhood overlap-aware graph neural networks for link prediction. *Advances in Neural Information Processing Systems* 34 (2021), 13683–13694.

[51] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. *Advances in neural information processing systems* 31 (2018).

- [52] Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. 2021. Labeling trick: A theory of using graph neural networks for multi-node representation learning. *Advances in Neural Information Processing Systems* 34 (2021), 9061–9073.
- [53] Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. 2021. Labeling trick: A theory of using graph neural networks for multi-node representation learning. *Advances in Neural Information Processing Systems* 34 (2021), 9061–9073.
- [54] He Zhao, Lan Du, and Wray Buntine. 2017. Leveraging node attributes for incomplete relational data. In *International conference on machine learning*. PMLR, 4072–4081.
- [55] Tao Zhou, Linyuan Lü, and Yi-Cheng Zhang. 2009. Predicting missing links via local information. *The European Physical Journal B* 71 (2009), 623–630.
- [56] Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. 2021. Neural bellman-ford networks: A general graph neural network framework for link prediction. *Advances in Neural Information Processing Systems* 34 (2021), 29476–29490.

A SPECIAL CASES OF THE GENERAL PAIRWISE ENCODING

In this section we demonstrate that multiple popular heuristics and pairwise encodings can be formulated as special cases of the general pairwise encoding given in Eq. (2).

Common Neighbors (CNs) [38]: The CNs of a pair of nodes (a, b) is defined the overlapping 1-hop neighbors of both nodes, $\mathcal{N}_{(a,b)}^{\text{CN}} = \mathcal{N}(a) \cap \mathcal{N}(b)$. Eq. (2) is equal to the CNs when $h(a, b, u) = 1$ and $w(a, b, u)$ is:

$$w(a, b, u) = \begin{cases} 1, & \text{when } u \in \mathcal{N}(a) \cap \mathcal{N}(b) \\ 0, & \text{else} \end{cases}. \quad (12)$$

Adamic-Adar (AA) [2]: AA further weights each common neighbor by the reciprocal of its log-degree, i.e., $1/\log(d_u)$ for node u , allowing us to rewrite Eq. (2) as $h(a, b, u) = 1/\log(d_u)$ where $w(a, b, u)$ is equal to Eq. (12).

Resource Allocation (RA) [55]: RA is similar to AA except that it is omits the log, resulting in $1/d_u$. As before, Eq. (2) can be rewritten as the RA when $h(a, b, u) = 1/d_u$ and $w(a, b, u)$ is equal to Eq. (12).

Katz Index [22]: The Katz index is a global structural measure. It is defined as weighted summation of the number of paths of different lengths connecting a and b . It is given by the following where the decay weight $\beta \in [0, 1]$,

$$\text{Katz}(a, b) = \sum_{l=1}^{\infty} \beta^l A_{a,b}^l. \quad (13)$$

This is equivalent to Eq. (2) when:

$$w(a, b, u) = \sum_{l=1}^{\infty} \beta^l e_a^T A^l, \quad (14)$$

where $e_i \in \mathbb{B}^{|\mathcal{V}|}$ is a one-hot vector for a node i . We further set,

$$h(a, b, u) = \begin{cases} e_b^T, & \text{when } u = b \\ 0, & \text{else} \end{cases}. \quad (15)$$

Personalized Pagerank (PPR) Score [7]: The personalized pagerank score is the pagerank score localized to a root node u . Eq. (2) can be rewritten as the PPR score when setting $h(a, b, u)$ equal to (15) and, following Chung [12], setting $w(a, b, u)$ to:

$$w(a, b, u) = \alpha \sum_{l=0}^{\infty} (1 - \alpha)^l e_a^T (D^{-1} A)^l. \quad (16)$$

Feature Similarity: The feature similarity of the pair of nodes (a, b) is expressed by $\text{dis}(\mathbf{x}_a, \mathbf{x}_b)$ where \mathbf{x}_a are the node features

of node a and $\text{dis}(\cdot)$ is a distance function (e.g., euclidean distance). This can be rewritten as Eq. (2) by substituting $w(a, b, u) = \text{dis}(\mathbf{x}_a, \mathbf{x}_u)$ and $h(a, b, u) = e_b^T$ where $e_i \in \mathbb{B}^{|\mathcal{V}|}$ is a one-hot vector for a node i .

NCN [45]: The pairwise encoding used in NCN is defined as the summation of the representations for the CNs of a link. Eq. (2) can be rewritten as NCN when $w(a, b, u)$ is equal to Eq. (12). $h(a, b, u)$ is equal to the node representation u encoded by a MPNN, i.e., $h(a, b, u) = \mathbf{h}_u$ where $H = \text{MPNN}(A, X)$.

NCNC [45]: NCNC extends NCNC by further weighting the 1-hop (non-CN) by their probability of linking to the other nodes. Given Eq. (2), the weight $w(a, b, u)$ is equal to following where 1-hop neighbors are weighted by their probability of linking with the other node:

$$w(a, b, u) = \begin{cases} 1, & \text{when } u \in \mathcal{N}_{(a,b)}^{\text{CN}} \\ \text{NCN}(A, X, b, u), & \text{when } u \in \mathcal{N}(a) \\ \text{NCN}(A, X, a, u), & \text{when } u \in \mathcal{N}(b) \\ 0, & \text{else} \end{cases}. \quad (17)$$

$\text{NCN}(A, X, a, u)$ is the probability of a and u being linked using the NCN model. We further define $h(a, b, u) = \mathbf{h}_u$.

Neo-GNN [50]: The pairwise encoding used in Neo-GNN considers the higher-order neighborhood overlap between two nodes. The formulation is given in Section A. When $l = 1$, it can be expressed using Eq. (2) by setting:

$$h(a, b, u) = f_1 \left(\sum_{v \in \mathcal{N}(u)} f_2(A_{uv}) \right)^2, \quad (18)$$

and $w(a, b, u)$ as equal to Eq. (12).

B PROOF OF PROPOSITION 1

Proposition 1. Consider a target link (a, b) and a node $u \in \mathcal{V} \setminus \{a, b\}$. The PPR [7] score of a root node i and target node j with teleportation probability α is denoted by $\text{ppr}(i, j)$. Let $r_a^k(u)$ be the probability of a walk of length k beginning at node a and terminating at u . We define $r_{a,b}^k(u) := r_a^k(u) + r_b^k(u)$. We also define a weight $\gamma^k := \alpha(1 - \alpha)^k$ for all walks of length k . The PPR scores, $\text{ppr}(a, u)$ and $\text{ppr}(b, u)$, along with the random walk probabilities of disparate lengths, are interconnected through the following relationship.

$$\Gamma(a, b, u) = \text{ppr}(a, u) + \text{ppr}(b, u) = \sum_{k=0}^{\infty} \gamma^k r_{a,b}^k(u). \quad (6)$$

PROOF. Per Chung [12], the PPR vector for a root node s , pr_s , is equivalent to:

$$\text{pr}_s = \alpha \sum_{k=0}^{\infty} (1 - \alpha)^k W^k x_s, \quad (19)$$

where W is a the random walk matrix and x_s is a preference vector that is a one-hot vector for element s . We note that $\text{pr}_s(t)$ represents the landing probability of node t given the root node s . As such, by definition, $\text{pr}_s(t) = \text{ppr}(s, t)$. Furthermore, it is clear that $r_s^k = W^k x_s \in \mathbb{R}^{|\mathcal{V}|}$ represents the probability of a walk of length k beginning at node s and stop all other nodes, individually. Also, the

probabilities of all walks of length k are weighted by $\gamma^k = \alpha(1-\alpha)^k$. $\Gamma(a, b, u)$ can be obtained by first taking the sum of the PPR vectors for nodes a and b ,

$$\begin{aligned} \text{pr}_a + \text{pr}_b &= \alpha \sum_{k=0}^{\infty} (1-\alpha)^k W^k x_a + \alpha \sum_{k=0}^{\infty} (1-\alpha)^k W^k x_b, \\ \text{pr}_{a,b} &= \alpha \sum_{k=0}^{\infty} (1-\alpha)^k W^k (x_a + x_b), \end{aligned} \quad (20)$$

where $\text{pr}_{a,b} = \text{pr}_a + \text{pr}_b$. From this, we can express $\Gamma(a, b, u)$ as:

$$\begin{aligned} \Gamma(a, b, u) &= \text{ppr}(a, u) + \text{ppr}(b, u), \\ &= \text{pr}_{a,b}(u), \\ &= \text{pr}_a(u) + \text{pr}_b(u), \end{aligned} \quad (21)$$

which as shown in Eq. (20) is equivalent to the probability of a walk that originates from either node a or b and terminates at node u . This completes the proof. \square

C ATTENTION FORMULATION

For a target link (a, b) , LPFormer attends to the nodes in the set $\bar{\mathcal{V}}(a, b)$. The attention mechanism used in LPFormer is defined in Section 3 as follows where $w(a, b, u)$ is the attention weight of u to the target link and $\bar{\mathcal{V}}(a, b) = \mathcal{V} \setminus \{a, b\}$:

$$\begin{aligned} \tilde{w}(a, b, u) &= \phi \left(\mathbf{h}_a, \mathbf{h}_b, \mathbf{h}_u, \text{rpe}_{(a,b,u)} \right), \\ w(a, b, u) &= \frac{\exp(\tilde{w}(a, b, u))}{\sum_{v \in \bar{\mathcal{V}}(a,b)} \exp(\tilde{w}(a, b, v))}. \end{aligned} \quad (22)$$

The function $\phi(\cdot)$ is modeled via the attention mechanism defined in GATv2 [8]. We define $a \in \mathbb{R}^{2d'}$ and $W \in \mathbb{R}^{d \times d'}$. The raw attention weights are then given by:

$$\tilde{w}(a, b, u) = \mathbf{a}^T \text{LeakyReLU} \left[W \mathbf{h}_a \parallel W \mathbf{h}_b \parallel W \mathbf{h}_u \parallel \text{rpe}_{(a,b,u)} \right]. \quad (23)$$

The final attention weights, $w(a, b, u)$, are given by passing $\tilde{w}(a, b, u)$ through a softmax activation layer.

D ADDITIONAL EXPERIMENTAL DETAILS

D.1 Computation of the PPR Matrix

We compute the PPR matrix via the efficient approximation algorithm introduced by Andersen et al. [3]. The estimation is controlled by a tolerance parameter ϵ . We use: $\epsilon = 1e^{-7}$ for Cora and Citeseer, $\epsilon = 5e^{-5}$ for ogbl-collab and ogbl-ppa, $\epsilon = 1e-5$ for Pubmed, and $\epsilon = 5e^{-3}$ for ogbl-Citation2. The value of ϵ is chosen as a trade-off between accuracy and sparsity to allow for ease of storage in GPU memory.

D.2 Splitting Target Links by LP Factor

In Section 4.3 we demonstrate the performance on samples that correspond to a single LP factor. In this section we further detail the algorithm used to determine the set of samples corresponding to each factor. We consider the three main factors: local structural information, global structural information, and feature proximity. We measure each using a single representative heuristic: CNs [38] for local information, PPR [7] for global information, and cosine

feature similarity for feature proximity. For each sample, we check if the score is only high in **one** heuristic. In this way, it tells us that there is a dominant factor present in the pairwise information. The detailed algorithm is given in Algorithm 1.

We note that each target link may not belong to a category. This can be due to there being no or many dominant LP factor. We further set the percentile equal to 90% on all datasets except for ogbl-collab for which we use 80%. These values were chosen as we wanted the percentile to be suitably high such that we are confident that the corresponding factor is relevant to the target link. Furthermore, we use a lower value for ogbl-collab as we found it produced a more even distribution of links by factor.

Algorithm 1 Determining Samples by LP Factor

Require:

$\text{CN}(\cdot)$ = Maps (i, j) to # of CNs of the pair

$\text{PPR}(\cdot)$ = Maps (i, j) to PPR score of the pair

$\text{FS}(\cdot)$ = Maps (i, j) to feature cosine similarity of the pair

p = Percentile used to determine whether a factor is present

$\mathcal{E}^{\text{test}}$ = Positive test links

- 1: Define the p -th percentile for each heuristic: \hat{s}^{CN} , \hat{s}^{FS} , \hat{s}^{PPR}
- 2: Create empty lists L^{CN} , L^{PPR} , and L^{FS}
- 3: **for** $(i, j) \in \mathcal{E}^{\text{test}}$ **do**
- 4: link-cn = $\text{CN}(i, j)$
- 5: link-fs = $\text{FS}(i, j)$
- 6: link-ppr = $\text{PPR}(i, j)$
- 7: // Assign sample to corresponding list based on scores
- 8: **if** link-cn $\geq \hat{s}^{\text{CN}}$ **and** link-fs $< \hat{s}^{\text{FS}}$ **and** link-ppr $< \hat{s}^{\text{PPR}}$ **then**
- 9: Append(L^{CN} , (i, j))
- 10: **else if** link-cn $< \hat{s}^{\text{CN}}$ **and** link-fs $\geq \hat{s}^{\text{FS}}$ **and** link-ppr $< \hat{s}^{\text{PPR}}$ **then**
- 11: Append(L^{FS} , (i, j))
- 12: **else if** link-cn $< \hat{s}^{\text{CN}}$ **and** link-fs $< \hat{s}^{\text{FS}}$ **and** link-ppr $\geq \hat{s}^{\text{PPR}}$ **then**
- 13: Append(L^{PPR} , (i, j))
- 14: **end if**
- 15: **end for**
- 16: **return** L^{CN} , L^{PPR} , L^{FS}

E ADDITIONAL LP FACTOR EXPERIMENTS

Additional results by LP factor are given in Figure 5 for ogbl-ppa and Pubmed. We note that for ogbl-ppa, since the feature are one-hot encodings, the feature similarity is not useful and is this omitted.

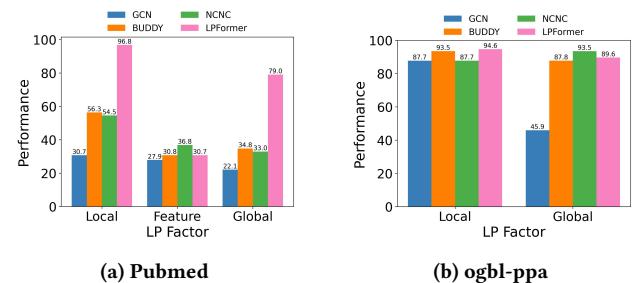


Figure 5: Additional LP Factor results on (a) Pubmed and (b) ogbl-ppa.