

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 IMPROVING EXPRESSIVITY IN LINK PREDICTION WITH GNNs VIA THE SHORTEST PATH

Anonymous authors

Paper under double-blind review

## ABSTRACT

Graph Neural Networks (GNNs) often fail to capture the link-specific structural patterns essential for accurate link prediction, since their node-centric message passing might overlook the subgraph structures connecting two nodes. Prior attempts to inject such structural context either suffer from high computational cost or rely on oversimplified heuristics (e.g., common neighbor counts) that cannot capture multi-hop dependencies. We propose SP4LP (Shortest Path for Link Prediction), a new framework that integrates GNN-based node encodings with sequence modeling over shortest paths. Specifically, SP4LP first computes node representations with a GNN, then extracts the shortest path between each candidate node pair and processes the sequence of node embeddings with a sequence model. This design allows SP4LP to efficiently capture expressive multi-hop relational patterns. Theoretically, we show that SP4LP is strictly more expressive than both standard message-passing GNNs and several leading structural feature methods, positioning it as a general and principled framework for link prediction in graphs. Empirically, SP4LP sets a new state of the art on many standard link prediction benchmarks.

## 1 INTRODUCTION

Graph Neural Networks (GNNs) are widely adopted for link-level tasks such as link prediction (Zhang & Chen, 2018; Lü & Zhou, 2011; Zhou, 2021), link classification (Rossi et al., 2021; Wang et al., 2021; Cheng et al., 2025) and link regression (Liang et al., 2025; Dong et al., 2019) with applications spanning recommender systems (Ying et al., 2018), knowledge graph completion (Nickel et al., 2015), and biological interaction prediction (Jha et al., 2022).

Despite their popularity, standard GNNs struggle to accurately represent links, as they typically construct link embeddings by aggregating the representations of the two endpoint nodes. This node-centric strategy leads to a key limitation: structurally distinct links may be mapped to the same representation when their endpoints are automorphic (Srinivasan & Ribeiro, 2019; Chamberlain et al., 2023; Zhang et al., 2021). For example, in the graph of Figure 1, links  $(v, u)$  and  $(v, u')$  yield identical representations under any standard GNN, even if one pair shares a common neighbor and the other does not. This issue, known as the automorphic node problem (Chamberlain et al., 2023), highlights a fundamental expressivity bottleneck in message-passing schemes for link representation.

To address this, several methods enhance GNNs with structural features (SFs), which can be broadly classified into three paradigms (Wang et al., 2024): SF-then-GNN, which injects structural context into the graph before message passing (e.g., SEAL (Zhang et al., 2021), NBFNet (Zhu et al., 2021)); SF-and-GNN, which computes SFs and node embeddings in parallel (e.g., Neo-GNN (Yun et al., 2021), BUDDY (Chamberlain et al., 2023)); and GNN-then-SF, which applies message passing once to compute node representations and then combines them using task-specific structural context (e.g., NCN and NCNC (Wang et al., 2024)).

While SF-then-GNN methods are expressive, they are computationally inefficient, often requiring subgraph extraction or retraining per link. SF-and-GNN models are efficient but rely on predefined heuristics, limiting their ability to capture rich relational patterns. GNN-then-SF approaches offer a compelling trade-off between expressivity and scalability, but current methods in this class, i.e., NCN and NCNC, still rely on overlap between the neighborhoods of the endpoints: NCN uses

054 common neighbors, while NCNC also leverages predicted second-order common neighbors. In the  
 055 absence of such overlapping neighbors, they revert to standard GNN behavior and lose expressive  
 056 power.

057 In this paper we propose SP4LP, a novel method in the GNN-  
 058 then-SF paradigm that combines high expressiveness with compu-  
 059 tational efficiency. SP4LP constructs a path-aware representation  
 060 by incorporating the embeddings of all nodes along the shortest path  
 061 connecting the two endpoints. These node embeddings, obtained  
 062 via a base GNN, are then processed as a sequence using a dedi-  
 063 cated sequence model, such as a Transformer (Vaswani et al., 2017),  
 064 LSTM (Hochreiter & Schmidhuber, 1997), or an injective summa-  
 065 tion function (Xu et al., 2019).

066 The first key advantage of SP4LP is its **expressiveness**. Unlike  
 067 structural features such as common neighbors, which may not ex-  
 068 ist for many node pairs and can lead to degenerate cases in sparse  
 069 graphs, the shortest path is always defined between any two nodes if  
 070 the graph is connected. Shortest paths are more broadly defined, as  
 071 common neighbors imply a path, but not vice versa. Moreover, since  
 072 the embeddings of the nodes along the path are generated through  
 073 message passing, they implicitly encode the broader local structure  
 074 surrounding the link. This richer structural context enables SP4LP to distinguish non-automorphic  
 075 links even when their endpoints are automorphic, thereby overcoming the automorphic node prob-  
 076 lem. We formally prove that SP4LP is strictly more expressive than [some existing approaches](#).

077 SP4LP is **efficient** and **scalable**. The message-passing step is performed only once on the entire  
 078 graph, and the shortest path computation is a preprocessing step. Unlike SF-then-GNN methods  
 079 such as SEAL or NBFNet, SP4LP avoids costly per-link subgraph extraction or online traversal  
 080 during inference, allowing it to scale to large graphs and high-throughput settings.

081 Moreover, SP4LP is a **general and flexible framework**. It can be instantiated with any GNN  
 082 architecture to compute node embeddings (e.g., GCN (Kipf & Welling, 2016a), GAT (Velickovic  
 083 et al., 2017), GraphSAGE (Hamilton et al., 2017b)), and supports a range of sequence models for  
 084 encoding the path structure, from lightweight aggregators to fully expressive recurrent or attention-  
 085 based models.

086 Finally, SP4LP achieves **state-of-the-art performance** across several benchmark datasets. Under  
 087 the challenging HeaRT evaluation protocol (Li et al., 2023), it consistently outperforms existing link  
 088 prediction methods while maintaining competitive inference speed and low memory usage.

089 These properties make SP4LP a principled and practical solution for learning expressive link rep-  
 090 resentations in real-world graph learning scenarios.

## 092 2 PRELIMINARIES

095 **Definition 2.1 (graph).** A **graph** is a tuple  $G = (V, E, \mathbf{X}^0)$  where  $V = \{1, \dots, n\}$  is a set of  
 096 nodes,  $E \subseteq V \times V$  is a set of edges and  $\mathbf{X}^0 \in \mathbb{R}^{n \times f}$  is the node features matrix. To each graph  
 097 is associated an adjacency matrix  $\mathbf{A} \in \{0, 1\}^{n \times n}$  with  $\mathbf{A}_{i,j} = 1$  if and only if  $(i, j) \in E$ . In this  
 098 work, we consider simple, finite and undirected graphs.

099 **Definition 2.2 (message passing).** Let  $G = (V, E, \mathbf{X}^0)$  be a graph. In **message passing** scheme,  
 100 representation of nodes  $v \in V$  is iteratively updated as follows:

$$101 \quad \mathbf{x}_v^0 = \mathbf{X}_{[v,:]}^0 \quad (1)$$

$$103 \quad \mathbf{x}_v^l = \text{UPDATE}(\mathbf{x}_v^{l-1}, \text{AGGREGATE}(\{\mathbf{x}_u^{l-1} \mid u \in N(v)\})) \quad (2)$$

105 where  $N(v)$  is the first-order neighborhood of node  $v$ .

106 Graph Neural Networks (GNNs) are a class of neural architectures that operate on graphs by it-  
 107 eratively updating node representations through the message passing scheme. It has been proven

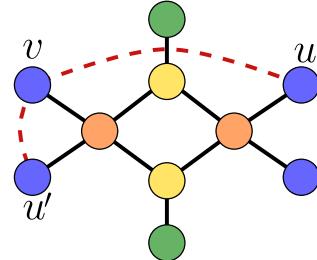


Figure 1: Links  $(v, u)$  and  $(v, u')$  have different structural roles within the graph, yet a GNN assigns them identical representations.

108 that GNNs are at most as effective as the Weisfeiler–Lehman (WL) test in distinguishing between  
 109 graphs (Morris et al., 2019; Xu et al., 2019).

110 **Definition 2.3** (GNN link representation model). *A **GNN link representation model**  $M$  is a class  
 111 of functions*

$$F : ((u, v), G) \mapsto \mathbf{x}_{(u, v)} \in \mathbb{R}^d \quad (3)$$

114 which maps node pairs in  $(u, v) \in V \times V$  to vector representations using the message passing  
 115 scheme defined in Definition 2.2.

116 Note that the pair  $(u, v)$  belongs to  $V \times V$ , meaning that we compute a representation for any node  
 117 pair, what we refer to as a *link*, regardless of whether an edge between them exists in  $E$ . This  
 118 general definition reflects the nature of the downstream tasks we aim to address once the link repre-  
 119 sentation is available, most notably, link prediction, where the objective is to estimate the likelihood  
 120 of a connection between arbitrary node pairs. To this end, the model learns representations for all  
 121 possible pairs, not just those connected by an edge. A widely adopted approach for learning such  
 122 representations is what we refer to as a *pure GNN*, defined as follows:

123 **Definition 2.4** (pure GNN). *A **pure GNN** model calculates representation  $\mathbf{x}_{(u, v)} \in \mathbb{R}^d$  for each  
 124 pair of nodes  $(u, v)$  with  $u, v \in V$  as follows:*

$$\mathbf{x}_{(u, v)} = g(\mathbf{x}_u^L, \mathbf{x}_v^L) \quad (4)$$

126 where  $g$  is an aggregation function and  $\mathbf{x}_u^L, \mathbf{x}_v^L$  are the node representation of  $u$  and  $v$  learned by  $L$   
 127 layers of message passing as defined in Definition 2.2.

128 Pure GNNs are inherently limited in terms of expressiveness. In particular even when the base GNN  
 129 is the most powerful, they may assign the same representation to structurally different links. Con-  
 130 sider, for example, the graph in Figure 1: the colors of the nodes indicate the colors produced by  
 131 the WL algorithm; thus, using a most powerful GNN nodes  $u, u'$  will be assigned to the same rep-  
 132 resentation. As a result, no matter how expressive the aggregation function  $g$  is, the representations  
 133 of the pairs  $(v, u)$  and  $(v, u')$  will be identical. However, the links  $(v, u)$  and  $(v, u')$  have different  
 134 roles within the graph structure. We provide a formal definition of what it means for two links to be  
 135 different.

136 **Definition 2.5** (node permutation). *A **node permutation**  $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$  is a bijective  
 137 function that assigns a new index to each node of the graph. All the  $n!$  possible node permutations  
 138 constitute the permutation group  $\Pi_n$ . Given a subset of nodes  $S \subseteq V$ , we define the permutation  
 139  $\pi$  on  $S$  as  $\pi(S) := \{\pi(i) | i \in S\}$ . Additionally, we define  $\pi(\mathbf{A})$  as the matrix  $\mathbf{A}$  with rows and  
 140 columns permuted based on  $\pi$ , i.e.,  $\pi(\mathbf{A})_{\pi(i), \pi(j)} = \mathbf{A}_{i, j}$ .*

141 **Definition 2.6** (automorphism). *An **automorphism** on the graph  $G = (V, E, \mathbf{X}^0)$  is a permutation  
 142  $\sigma \in \Pi_n$  such that  $\sigma(\mathbf{A}) = \mathbf{A}$ . All the possible automorphisms on a graph constitute the automor-  
 143 phism group  $\Sigma_n^G$ .*

144 **Definition 2.7** (automorphic nodes). *Let  $G = (V, E, \mathbf{X}^0)$  be a graph and  $\Sigma_n^G$  its automorphism  
 145 group. Two nodes  $u, v \in V$  are said to be **automorphic nodes** ( $u \simeq v$ ) if:*

$$\exists \sigma \in \Sigma_n^G \quad \text{s.t.} \quad \sigma(\{u\}) = \{v\}. \quad (5)$$

146 **Definition 2.8** (automorphic links). *Let  $G = (V, E, \mathbf{X}^0)$  be a graph and  $\Sigma_n^G$  its automorphism  
 147 group. Two pairs of nodes  $(u, v), (u', v') \in V \times V$  are said to be **automorphic links** ( $(u, v) \simeq  
 148 (u', v')$ ) if:*

$$\exists \sigma \in \Sigma_n^G \quad \text{s.t.} \quad \sigma(\{u, v\}) = \{u', v'\}. \quad (6)$$

149 **Proposition 2.9.** *Pure GNN methods suffer from the automorphic node problem, i.e., for any graph  
 150  $G = (V, E, \mathbf{X}^0)$ , for pairs of links  $(u, v), (u', v') \in V \times V$  such that there exist  $\sigma_1 \in \Sigma_n^G$  and  
 151  $\sigma_2 \in \Sigma_n^G$  with  $\sigma_1(u) = u'$  and  $\sigma_2(v) = v'$ ,  $\mathbf{x}_{(u, v)} = \mathbf{x}_{(u', v')}$ , independently whether  $(u, v)$  and  
 152  $(u', v')$  are isomorphic, i.e, whether exist  $\sigma \in \Sigma_n^G$  with  $\sigma(\{u, v\}) = \{u', v'\}$ .*

153 This limitation is well-known in the literature (Chamberlain et al., 2023; Zhang et al., 2021). Import-  
 154 antly, it does not arise from the expressiveness bounds of GNNs, which are constrained by the WL  
 155 test. Even considering higher-order GNNs, i.e.,  $k$ -GNN (Morris et al., 2019), automorphic nodes  
 156 will be assigned to the same representation as the  $k$ -WL algorithm preserves graph automorphisms

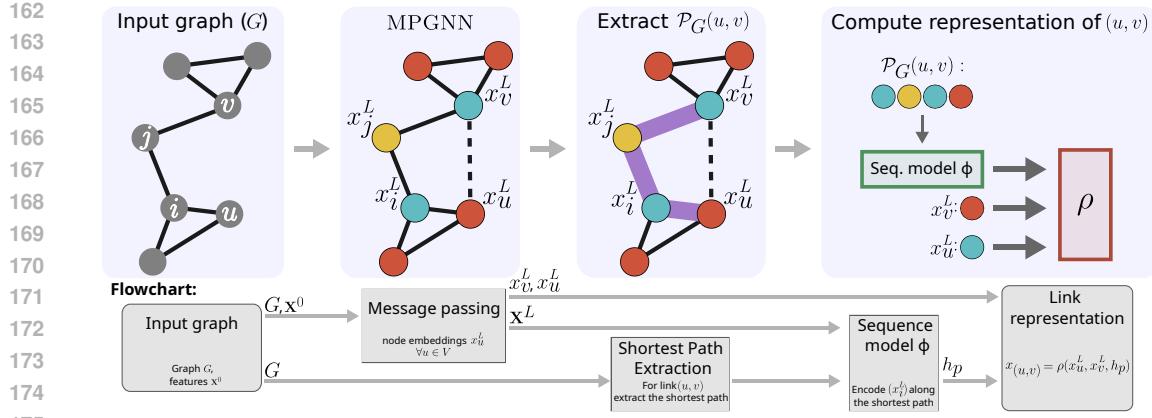


Figure 2: Overview of the SP4LP framework. First, a GNN is used to compute contextualized embeddings for all nodes in the graph. Then, for each target link, the shortest path connecting the two endpoints is extracted. The embeddings of the nodes along this path are passed to a sequence model (e.g., Transformer or LSTM) to compute a path-aware link representation. Below, we also include a flowchart summarizing the full SP4LP pipeline, highlighting each computational step and its inputs/outputs.

for every  $k$  (Lichter et al., 2025; Dawar & Vagnozzi, 2020; Cai et al., 1992). Thus, no standard GNN can distinguish non automorphic links composed by automorphic nodes.

To tackle this, several models have been proposed that enhance message passing by incorporating structural features (Wang et al., 2024; Zhu et al., 2021; Chamberlain et al., 2023; Wang et al., 2022; Zhang et al., 2021), thereby increasing the expressive power of the resulting link representations. We provide a formal definition of what it means for one link representation model to be more expressive and strictly more expressive than another.

**Definition 2.10 (more expressive).** Let  $M_1$  and  $M_2$  be two link representation models (Def. 2.3).  $M_2$  is **more expressive** than  $M_1$  ( $M_1 \preceq M_2$ ) if, for any graph  $G = (V, E, \mathbf{X}^0)$  and any pair  $(u, v), (u', v') \in V \times V$  with  $(u, v) \neq (u', v')$ :

$$\exists F_1 \in M_1 : F_1((u, v), G) \neq F_1((u', v'), G) \Rightarrow \exists F_2 \in M_2 : F_2((u, v), G) \neq F_2((u', v'), G). \quad (7)$$

**Definition 2.11 (strictly more expressive).** Let  $M_1$  and  $M_2$  be two link representation models (Def. 2.3). We say that  $M_2$  is **strictly more expressive** than  $M_1$  ( $M_1 \prec M_2$ ) if:

- $M_2$  is more expressive than  $M_1$  (Def. 2.10), and
- there exists a graph  $G = (V, E, \mathbf{X}^0)$  and a pair of links  $(u, v), (u', v') \in V \times V$  with  $(u, v) \neq (u', v')$  such that:

$$\begin{aligned} \forall F_1 \in M_1 : \quad & F_1((u, v), G) = F_1((u', v'), G) \\ \text{and} \quad & \exists F_2 \in M_2 : \quad F_2((u, v), G) \neq F_2((u', v'), G). \end{aligned}$$

In the following section, we introduce our model SP4LP and demonstrate its improved expressive power in distinguishing structurally different links.

### 3 RELATED WORK

GNNs have been extensively used for link representation tasks. In standard approaches like Graph Autoencoders (GAE) (Kipf & Welling, 2016b), node embeddings are computed via message passing, and a simple decoder (e.g., inner product followed by a sigmoid) predicts link existence. While efficient, these models exhibit limited expressiveness (Zhang et al., 2021; Chamberlain et al., 2023), primarily due to their inability to capture rich structural patterns beyond immediate neighborhoods. This limitation has motivated the integration of explicit structural information into GNN-based models.

216 **Incorporating Structural Information into GNNs.** To overcome expressiveness bottlenecks, several  
 217 methods augment GNNs with structural features. Neo-GNN (Yun et al., 2021) injects hand-  
 218 crafted features into the message-passing process, while ELPH (Chamberlain et al., 2023) and its  
 219 scalable variant BUDDY employ MinHash and HyperLogLog sketches to capture multi-hop pat-  
 220 terns, with BUDDY precomputing sketches offline. NCN (Wang et al., 2024) aggregates embed-  
 221 dings from endpoint nodes and their common neighbors, and NCNC extends this by predicting  
 222 missing neighbors before reapplying NCN. NBFNet (Zhu et al., 2021) instead aggregates informa-  
 223 tion over all paths between node pairs via Bellman-Ford-inspired recursive functions. Differently  
 224 from these approaches, our method focuses explicitly on the shortest path between node pairs, using  
 225 a sequence model to capture dependencies along this path. This results in more focused and inter-  
 226 pretable representations while avoiding the inefficiencies of modeling broader multi-hop neighbor-  
 227 hoods or exhaustive path sets. **Beyond GNN-based approaches**, PROXI (Tola et al., 2025) proposes  
 228 a hybrid proximity-based framework built on handcrafted structural indices combined with a non-  
 229 GNN predictor (XGBoost). While conceptually related through the use of structural signals, PROXI  
 230 fundamentally differs from GNN-based link representation models.

231 **Enhancing GNN Expressiveness through Positional and Structural Encoding.** Positional en-  
 232 codings further enrich GNN expressiveness. PEG (Wang et al., 2022) integrates Laplacian-based  
 233 encodings into message passing, weighting neighbors by positional distances. SEAL (Zhang et al.,  
 234 2021) extracts  $h$ -hop enclosing subgraphs and labels nodes via the DRNL scheme before applying  
 235 a GNN. While expressive, these methods struggle to scale due to subgraph extraction overhead. In  
 236 contrast, our model achieves expressiveness by operating directly on compact, informative shortest-  
 237 path sequences, enabling better scalability without sacrificing representational power.

238 **Shortest-Path Structures in Graph Learning.** Shortest-path information has also proven effective  
 239 in tasks beyond link prediction, such as graph classification (Ying et al., 2021; Airale et al., 2025)  
 240 and node classification on heterophilous graphs (Li et al., 2020). These works highlight the power of  
 241 shortest-path structures across graph learning domains. Building on this insight, our model directly  
 242 leverages shortest-path sequences for link representation, showing that this structure is particularly  
 243 effective when combined with modern sequence models.

244 **Comparison with GDGNN Geodesic GNN (GDGN)** Kong et al. (2022) is the closest method to  
 245 ours in that it also extracts shortest-path information after a single GNN run. However, our approach  
 246 differs in three key aspects. (i) Path modeling: GDGNN applies permutation-invariant pooling over  
 247 the path nodes, losing the order and direction of the geodesic. SP4LP instead treats the shortest  
 248 path as an ordered sequence and processes it with a sequence model, enabling the capture of struc-  
 249 tural patterns that pooling cannot represent. (ii) Objective: GDGNN is a broad framework aimed at  
 250 improving WL-level expressiveness; SP4LP is a simple, task-specific model for structural link repre-  
 251 sentation. (iii) Scalability: GDGNN includes additional geodesic modules and subgraph reasoning,  
 252 whereas SP4LP keeps inference lightweight by decoupling the global GNN run from per-query path  
 253 processing.

## 254 4 SP4LP: AN EXPRESSIVE GNN-THEN-SF MODEL FOR LINK 255 REPRESENTATION

256 Existing GNN link representation models that leverage Structural Features (SF) fall into three cat-  
 257 egories (Wang et al., 2024): **SF-and-GNN**, which compute SF and GNN embeddings separately and  
 258 then combine them (e.g., Neo-GNN (Yun et al., 2021), BUDDY (Chamberlain et al., 2023)); **SF-**  
 259 **then-GNN**, which augment the graph with SF before applying GNN (e.g., SEAL (Zhang & Chen,  
 260 2018), NBFNet (Zhu et al., 2021)); and **GNN-then-SF**, which compute GNN embeddings first and  
 261 then aggregate them using SF (e.g., NCN, NCNC (Wang et al., 2024)).

262 In this work, we adopt the GNN-then-SF paradigm, which combines the scalability advantages  
 263 of applying message passing only once, as in the SF-and-GNN setting, with the expressiveness  
 264 typical of SF-then-GNN approaches, due to the ability to aggregate over task-specific sets of node  
 265 representations. To date, the only existing models that follow this paradigm are NCN and NCNC.  
 266 For a more in-depth discussion on the differences between our method, SP4LP and NCN, refer to  
 267 Section 4.1. In the following, we introduce the necessary definitions, formally describe the model,  
 268 and present theoretical results characterizing its expressive power.

270 **Definition 4.1 (path).** Let  $G = (V, E, \mathbf{X}^0)$  be a graph and  $u, v \in V$  to nodes. A **path** in  $G$  from  
 271  $u$  to  $v$  is a sequence of nodes  $P = (u_0, u_1, \dots, u_k)$  with (i)  $u_i \in V$  for all  $i = 0, \dots, k - 1$ , (ii)  
 272  $u_0 = u$  and  $u_k = v$ , (iii)  $(u_i, u_{i+1}) \in E$  for all  $i = 0, \dots, k - 1$ , and (iv) all nodes in the sequence  
 273 are distinct (i.e.,  $u_i \neq u_j$  for all  $i \neq j$ ). The length of a path  $P$ ,  $\text{len}(P)$  is the number of edges it  
 274 contains.

275 **Definition 4.2 (shortest path length).** Let  $\mathcal{P}_G(u, v)$  denote the set of all paths from  $u$  to  $v$  in  
 276  $G$ . The **shortest path length**  $d_G(u, v)$  is the minimum length among all paths i.e.,  $d_G(u, v) =$   
 277  $\min_{P \in \mathcal{P}(u, v)} \text{len}(P)$ .

278 **Definition 4.3 (shortest path).** A **shortest path** between  $u$  to  $v$  in  $G$  is any path  $P^* \in \mathcal{P}_G(u, v)$  such  
 279 that  $\text{len}(P^*) = d_G(u, v)$ . The set of all the shortest path from  $u$  to  $v$  in  $G$  is denoted as  $\mathcal{P}_G^*(u, v)$ .  
 280

281 Let  $G = (V, E, \mathbf{X}^0)$  be a graph,  $u, v \in V$ . SP4LP is a GNN link representation model (see  
 282 Definition 2.3) that computes link representation as follows:

$$283 \text{SP4LP}((u, v), G) = \rho \left( \text{GNN}(u, G), \text{GNN}(v, G), \text{AGG} \left( \left\{ \phi \left( \text{GNN}(u_i, G) \right)_{i=1}^k \mid (u_i)_{i=1}^k \in \mathcal{P}_G^*(u, v) \right\} \right) \right) \quad (8)$$

285 where  $k = d_G(u, v)$ ,  $\text{GNN}(u, G) \in \mathbb{R}^d$  is the representation of node  $u \in V$  obtained at the final  
 286 layer of message passing as in Definition 2.2,  $\phi : \mathbb{R}^{k \times d} \rightarrow \mathbb{R}^d$  is a sequence model on the GNN  
 287 representations of nodes in the shortest path from  $u$  to  $v$ ,  $\text{AGG}$  is an aggregation function over  
 288 multiset of shortest paths representations and  $\rho : \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  combine the endpoint nodes  
 289 representations with the shortest paths representation to get a final link representation. Since we  
 290 consider undirected graphs, we consider  $\mathcal{P}_G^*(u, v) := \mathcal{P}_G^*(u, v) \cup \mathcal{P}_G^*(v, u)$ .  
 291

292 Although shortest paths provide a focused and efficient structural summary, they may lose information  
 293 carried by alternative routes or larger subgraphs, reducing robustness and emphasizing local  
 294 rather than global structure. We quantify this trade-off experimentally in Section 5.2, showing that  
 295 replacing full path embeddings with simple distance information leads to significant performance  
 296 degradation.

297 For a graph with  $n$  nodes and  $m$  edges, single-source BFS (Cormen et al., 2009) runs in  $O(n + m)$   
 298 time on an adjacency-list representation. Although computing all-pairs shortest paths (APSP) would  
 299 require  $O(n(n + m))$  time, SP4LP does not perform APSP. In our experiments, we compute shortest  
 300 paths only for the supervised train/validation/test links, which requires at most a small number of  
 301 BFS traversals. Let  $S$  be the number of distinct source nodes among these links; the preprocessing  
 302 cost is therefore  $O(S(n + m))$ , typically several orders of magnitude smaller than  $O(n(n + m))$ .  
 303 The resulting paths are cached and reused throughout training and inference, so no graph traversal  
 304 is performed during model training or batched evaluation. In fully-inductive settings where new  
 305 nodes may appear, one can compute shortest-path trees by running a single BFS per new node,  
 306 providing distances and parent pointers to all existing nodes at once. This preserves the same per-  
 307 link inference efficiency and avoids BFS operations per queried link. A more detailed analysis of  
 308 the overall complexity, including a comparison with prior approaches, is provided in Appendix E. In  
 309 practice, all shortest paths are precomputed before training, and SP4LP simply retrieves the stored  
 310 sequences during training and inference. No BFS operations are performed online.

311 SP4LP is a general and flexible framework: both the underlying GNN used to compute node rep-  
 312 resentations and the sequence model used to process the embeddings along the shortest path can be  
 313 chosen modularly. For instance, the GNN component can be instantiated with architectures such  
 314 as GCN (Kipf & Welling, 2016a), GAT (Velickovic et al., 2017) or GraphSAGE (Hamilton et al.,  
 315 2017b), while the sequence model can range from simple aggregation functions like injective sum-  
 316 mation as the one proposed in Xu et al. (2019), to more complex architectures such as LSTMs  
 317 (Hochreiter & Schmidhuber, 1997), GRU (Chung et al., 2014) or Transformers (Vaswani et al.,  
 318 2017). An overview of SP4LP is illustrated in Figure 2.

319 The additional structural context given by the sequence of embeddings of nodes within the shortest  
 320 path enables the model to distinguish links that are otherwise indistinguishable to standard message-  
 321 passing methods, such as those involving automorphic nodes.

322 **Proposition 4.4.** SP4LP does not suffer from the automorphic node problem.

323 The proof can be found in Appendix A. As an example of non-automorphic links composed of  
 324 automorphic nodes that SP4LP can successfully distinguish, consider the links  $(v, u)$  and  $(v, u')$

324 shown in Figure 1. While  $u' \simeq u'$  via the identity, and  $v \simeq u$  via an automorphism induced by a  
 325 vertical axis of symmetry (i.e., a mirror reflection), the links  $(v, u)$  and  $(v, u')$  are not automorphic.  
 326 This asymmetry is captured by the distinct shortest paths between the endpoints: the shortest path  
 327 from  $v$  to  $u'$  consists of  $v$ , and orange node, and  $u'$ , whereas the shortest path from  $v$  to  $u$  includes  $v$ ,  
 328 and orange node, a yellow node, another orange node, and finally  $u$ . In addition to overcoming this  
 329 limitation, SP4LP is strictly more expressive than several state-of-the-art message passing methods  
 330 for link representation learning.

331 **Theorem 4.5.** SP4LP is strictly more expressive than Pure GNNs, NCN, BUDDY, NBFnet and  
 332 Neo-GNN.

333 The proof can be found in Appendix A. In the following sections, we complement the theoretical  
 334 analysis with an extensive experimental evaluation, showing that SP4LP also achieves state-of-the-  
 335 art performance on standard link prediction benchmarks.

#### 337 4.1 FURTHER COMPARISON WITH NCN

338 NCN computes the representation of a link by aggregating the GNN em-  
 339 beddings of its endpoints and their common neighbors. However, its re-  
 340 liance on the common neighbor structure makes it particularly vulnera-  
 341 ble to graph incompleteness. Moreover, the use of common neighbors  
 342 as the sole source of structural features results in a critical failure mode:  
 343 when two nodes share no neighbors, NCN reduces to a pure GNN and  
 344 can no longer leverage structural information. This situation is far from  
 345 rare in practice (see Table 1), where a substantial fraction of positive links  
 346 involve endpoints without common neighbors. In contrast, we propose  
 347 SP4LP, which incorporates additional pairwise information by encoding  
 348 the sequence of node embeddings along the shortest path connecting the  
 349 endpoints. Unlike common neighbors, the shortest path is always defined  
 350 in connected graphs and captures richer structural patterns, even in sparse  
 351 or incomplete settings.

352 Moreover, while NCN pools neighbors as unordered sets, SP4LP encodes paths as ordered se-  
 353 quences via LSTMs or Transformers, yielding strictly higher expressiveness (Theorem 4.5).

## 355 5 EXPERIMENTS

356 Models that compute link representations can be applied to a wide range of downstream tasks, with  
 357 link prediction being particularly impactful due to its broad applicability in domains such as recom-  
 358 mender systems (Ying et al., 2018), knowledge graph completion (Nickel et al., 2015), and biolog-  
 359 ical interaction prediction (Jha et al., 2022). The link representations produced by GNN methods  
 360 are used to estimate the probability of existence for each candidate link. To train models for link  
 361 prediction, existing edges in the graph are treated as positive examples, while negative examples  
 362 are generated through negative sampling, selecting node pairs that are not connected in the original  
 363 graph.

364 In this section, we extensively evaluate the performance of SP4LP on real-world link prediction  
 365 benchmarks against several baselines. In particular, we use three Planetoid citation networks: Cora,  
 366 Citeseer, and Pubmed (Yang et al., 2016) as well as two datasets from Open Graph Benchmark Hu  
 367 et al. (2020), i.e., ogbl-collab and ogbl-ddi. For Cora, Citeseer, and Pubmed, we use a single fixed  
 368 data split in all experiments. Table 7 in appendix C provides a summary of dataset statistics.

369 As baseline methods we consider three class of models: 1) **heuristic methods**: Common Neigh-  
 370 bors (CN) (Newman, 2001), Adamic-Adar (AA) (Adamic & Adar, 2003), Resource Allocation  
 371 (RA) (Zhou et al., 2009), Shortest Path (SP) (Liben-Nowell & Kleinberg, 2003), and Katz (Katz,  
 372 1953); 2) **Embedding-based methods**: Node2Vec (Grover & Leskovec, 2016), Matrix Factoriza-  
 373 tion (MF) (Menon & Elkan, 2011), and a Multilayer Perceptron (MLP) applied to node features; 3)  
 374 **Pure GNN methods**: Graph Convolutional Network (GCN) (Kipf & Welling, 2016a), Graph Atten-  
 375 tion Network (GAT) (Veličković et al., 2018), GraphSAGE (Hamilton et al., 2017a), and Graph Au-  
 376 toencoder (GAE) (Kipf & Welling, 2016b); 4) **Structural Features GNN methods**: SEAL (Zhang

377 Table 1: Fraction of  
 378 positive test links with-  
 379 out common neighbors.

Dataset	w/o CN
Cora	46%
Citeseer	55%
Pubmed	67%
ogbl-collab	52%
ogbl-ddi	0.04%
ogbl-ppa	10%
ogbl-citation2	37%

378 Table 2: MRR and Hits@K (%) results across all datasets, following the HeaRT evaluation setting Li  
 379 et al. (2023). The top three results for each metric are highlighted using **first**, **second**, and **third**.  
 380 *OOM* indicates that the model ran out of memory, while *>24h* denotes that the method did not  
 381 complete within 24 hours. Standard deviations over 5 runs are reported in the appendix B.

Models	Cora		Citeseer		Pubmed		Ogbl-ddi		Ogbl-collab		Ogbl-ppa		Ogbl-Citation2		
	MRR	Hits@10	MRR	Hits@10	MRR	Hits@10	MRR	Hits@20	MRR	Hits@20	MRR	Hits@20	MRR	Hits@20	
GNN	GCN	16.61	36.26	21.09	47.23	7.13	15.22	<b>13.46</b>	64.76	6.09	<b>22.48</b>	26.94	68.38	19.98	51.72
	GAT	13.84	32.89	19.58	45.30	4.95	9.99	12.92	<b>66.83</b>	4.18	18.30	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>
	SAGE	14.74	34.65	21.09	48.75	<b>9.40</b>	<b>20.54</b>	12.60	<b>67.19</b>	5.53	21.26	27.27	69.49	22.05	53.13
	GAE	<b>18.32</b>	<b>37.95</b>	25.25	49.65	5.27	10.50	3.49	17.81	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>
SFGNN	SEAL	10.67	24.27	13.16	27.37	5.88	12.47	9.99	49.74	<b>6.43</b>	<b>21.57</b>	29.71	76.77	20.60	48.62
	BUDDY	13.71	30.40	22.84	48.35	7.56	16.78	12.43	58.71	5.67	<b>23.35</b>	27.70	71.50	19.17	47.81
	Neo-GNN	13.95	31.27	17.34	41.74	7.74	17.88	10.86	51.94	5.23	21.03	21.68	64.81	16.12	43.17
	NCN	14.66	35.14	<b>28.65</b>	<b>53.41</b>	5.84	13.22	12.86	<b>65.82</b>	5.09	20.84	<b>35.06</b>	<b>81.89</b>	<b>23.35</b>	<b>53.76</b>
	NCNC	14.98	<b>36.70</b>	24.10	<b>53.72</b>	8.58	18.81	>24h	>24h	4.73	20.49	33.52	<b>82.24</b>	19.61	51.69
	NBFNet	13.56	31.12	14.29	31.39	>24h	>24h	>24h	>24h	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>
	PEG	15.73	36.03	21.01	45.56	4.40	8.70	12.05	50.12	4.83	18.29	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>	<i>OOM</i>
	LPFORMER	<b>16.80</b>	34.03	<b>26.34</b>	51.72	<b>9.99</b>	<b>21.43</b>	<b>13.20</b>	62.66	<b>7.62</b>	21.04	<b>40.25</b>	<b>84.1</b>	<b>24.70</b>	<b>57.30</b>
SP4LP (our)	<b>17.27</b>	<b>38.52</b>	<b>41.08</b>	<b>66.28</b>	<b>10.87</b>	<b>23.01</b>	<b>15.00</b>	47.96	<b>9.46</b>	20.00	<b>36.45</b>	76.90	<b>24.91</b>	<b>55.45</b>	

& Chen, 2018), BUDDY (Chamberlain et al., 2023), Neo-GNN (Yun et al., 2021), NBFNet (Zhu et al., 2021), NCN (Wang et al., 2024), NCNC (Wang et al., 2024) and PEG (Wang et al., 2022).

Importantly, as described in Section 4, SP4LP is a general framework that allows for different choices of both the underlying GNN architecture and the sequence model ( $\phi$  Equation 8). In our experimental setting, we treat the choice of GNN and the choice of  $\phi$  as hyperparameters, and perform hyperparameter tuning based on validation set performance. Specifically, we explore GCN, GAT, and GraphSAGE as GNN backbones, and LSTM, Transformer, and an injective sum Xu et al. (2019) aggregator as sequence models. Moreover, we choose an MLP for  $\rho$  and as AGG we choose to select the first shortest path retrieved by the BFS procedure for computational efficiency. Appendix D provides implementation details, including how shortest paths between node pairs are computed, as well as the hyperparameter configurations used in our experiments. Code to reproduce all experiments is available at<sup>1</sup>.

**Evaluation Setting** We evaluate model performance under the more challenging and realistic HeaRT evaluation setting Li et al. (2023). In this setting, each positive target link (i.e., an existing link) is ranked against a carefully selected set of hard negative samples (i.e., non-existing links), providing a more realistic assessment of link prediction performance in practical scenarios. We adopt two standard ranking metrics: Hits@K and Mean Reciprocal Rank (MRR). Following the HeaRT protocol, we report Hits@10 and MRR for Cora, Citeseer, and Pubmed, and Hits@20 along with MRR for ogbl-collab and ogbl-ddi. The same set of negative samples is used across all positive links, as specified in the HeaRT benchmark. The HeaRT evaluation introduces significantly harder negative samples compared to traditional evaluation settings, resulting in a more challenging and realistic benchmark. Li et al. (2023) show that this leads to a substantial performance drop across most models, with GNNs specifically designed for link prediction often being outperformed by simple heuristics or general-purpose GNNs. By adopting this challenging evaluation setting, we ensure a rigorous and meaningful comparison of model performance under conditions that closely resemble real-world applications.

## 5.1 RESULTS ON REAL-WORLD BENCHMARKS

Table 2 presents the performance of SP4LP and the baseline models in terms of MRR and Hits@10 on Cora, Citeseer, and Pubmed, and MRR and Hits@20 on the OGB datasets. SP4LP ranks first in terms of MRR on four out of five datasets and second on the remaining one. The improvements in MRR are often substantial: on Citeseer, for instance, SP4LP achieves a 43% gain over the second-best method, NCN. SP4LP also achieves the best Hits@K score on three out of five datasets. On the Ogbl-Collab dataset, SP4LP is comparable on Hits@20 to the third-best model (SEAL), when accounting for standard deviations (Appendix B). On Ogbl-ddi, where SP4LP performs worse, the

<sup>1</sup><https://anonymous.4open.science/r/sp4lp-3875/README.md>

lower score can be explained by the lack of node features. Our model benefits from the availability of node features, as it leverages nodes representations obtained via message passing. In settings where such features are absent, like in Oglb-ddi, the discriminative power of the learned representations is reduced. In addition to achieving the best performance in several datasets, SP4LP is also the most consistent model across all benchmarks. Unlike previous state-of-the-art models such as BUDDY and Neo-GNN, which tend to struggle on datasets like Cora, Citeseer and Pubmed, SP4LP maintains strong performance regardless of dataset characteristics. Overall, these results clearly demonstrate the superiority of SP4LP under the more challenging and realistic HeaRT evaluation setting, confirming its effectiveness for real-world link prediction tasks.

## 5.2 ABLATION STUDY

We perform an ablation study to evaluate the contribution of the main components of our model. In particular, we investigate two simplified variants to understand the importance of node representation learning and sequential modeling along the shortest path between target nodes. (1) **Sequence Model Only**: in this variant, the sequential model operates directly on the raw input features of the nodes along the shortest path, without incorporating node representations learned by the GNN. This setup

isolates the contribution of the sequential model in capturing relational patterns based solely on node features and structural path information. (2) **GNN + Shortest Path Length**: in this variant, the sequential model is completely removed. Link prediction is performed using only the learned node representations from the GNN, combined with the length of the shortest path between the target nodes. This evaluates the effectiveness of combining node embeddings with simple distance information, without explicitly modeling the intermediate nodes along the path.

Table 3 reports the results of these ablations, conducted on the Cora, Citeseer, and Pubmed datasets. Both variants show a clear performance drop compared to the full model, demonstrating the importance of jointly leveraging node representations and sequential model. The Sequence Model Only variant achieves reasonable results on simpler datasets such as Cora and Citeseer, but its performance degrades significantly on more complex datasets like Pubmed, highlighting the limitations of relying solely on raw node features without learned representations. The GNN + Shortest Path Length variant consistently underperforms across all datasets, indicating that simple distance information is insufficient. Indeed, replacing the embedding of the path with its length discards the information encoded in the node representations along the path. In contrast, when node embeddings are computed via message passing, they incorporate information from each node’s local neighborhood, thus implicitly encoding a broader subgraph around the path, not just the path itself. Overall, the full model SP4LP achieves the best results across all datasets, confirming the importance of combining learned node representations with sequential modeling over the shortest paths to capture both local and global structural patterns in the graph.

## 5.3 SCALABILITY ANALYSIS

We assess the scalability of SP4LP by examining how its GPU memory consumption and inference time evolve as the batch size increases, in comparison to several baseline methods. The results are presented in Figure 3. In terms of GPU memory usage, SP4LP exhibits remarkable efficiency: memory consumption remains nearly constant across small to medium batch sizes, and increases moderately only for the largest batches, starting with 0.77 GB and reaching at most 6.84 GB while consistently avoiding out-of-memory (OOM) failures.

Table 3: Ablation study results (%). MRR and Hits@K are reported for two model variants: (1) *Sequence Model Only*, using a sequential model on raw node features, and (2) *GNN + Shortest Path Length*, using GNN representations with path length. The full model SP4LP consistently outperforms both variants. Standard deviations over 5 runs are reported in Appendix B

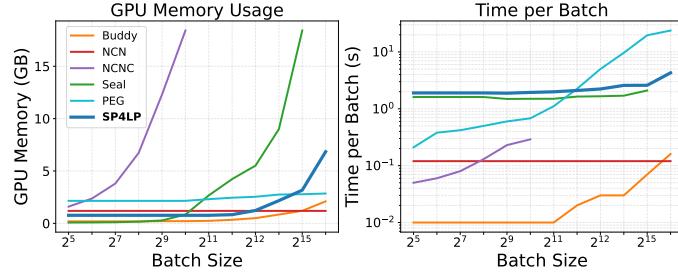
Models	Cora		Citeseer		Pubmed	
	MRR	Hits@10	MRR	Hits@10	MRR	Hits@10
<i>GNN + SP len.</i>	14.21	33.43	20.90	47.82	7.12	5.63
<i>Sequence Model</i>	16.86	36.03	27.45	54.20	8.58	12.87
SP4LP	<b>17.27</b>	<b>38.52</b>	<b>41.08</b>	<b>66.28</b>	<b>10.87</b>	<b>23.01</b>

486 PEG also maintains low memory usage; however, this advantage is undermined by its impractically slow inference, limiting its applicability in large-scale scenarios. SEAL, while competitive with SP4LP in terms of inference speed, suffers from excessive memory consumption, rapidly exceeding 18 GB and encountering OOM issues beyond a batch size of 32,768. NCNC is even more constrained, experiencing OOM failures already at relatively small batch sizes. Considering inference time, SP4LP matches the efficiency of SEAL, requiring only 2.1 seconds for a batch size of 16,384 and scaling smoothly to 4.29 seconds at 65,536. PEG, by contrast, is significantly slower, already taking 5 seconds at a batch size of 8,192 and exceeding 23 seconds at the maximum batch size tested. Although NCN and Buddy consistently achieve low inference times, this comes at the cost of substantially lower predictive performance, as they rely on simple heuristics with limited modeling capacity. This trade-off is clearly reflected in the results of Table 2. In summary, SP4LP achieves an excellent balance between low memory consumption, fast inference, high predictive accuracy, and strong model expressiveness.

## 507 6 CONCLUSION

510 We introduced SP4LP, a novel message-passing based framework for link representation that enhances the expressiveness of standard GNNs by incorporating sequential modeling over the shortest 511 path between target nodes. SP4LP follows the GNN-then-SF paradigm, thus effectively combining 512 the benefits of computing node embeddings only once with high expressive power. SP4LP explicitly 513 models multi-hop relational patterns through the use of a sequence encoder applied to the node 514 embeddings along the shortest path. We formally proved that SP4LP is strictly more expressive than 515 several state-of-the-art link representation models. Extensive experiments under the HeaRT evalua- 516 tion protocol confirm that SP4LP achieves state-of-the-art performance across diverse datasets.

517 As future work, we plan to extend SP4LP in several directions. First, we aim to incorporate mul- 518 tiple (or all) shortest paths between node pairs to improve robustness and capture richer structural 519 signals. Second, we intend to explore the application of SP4LP to heterogeneous graphs: since 520 shortest paths are well-defined in multi-relational settings and the path encoder can incorporate node 521 and edge type embeddings, SP4LP can be adapted to heterogeneous link prediction with minimal 522 architectural changes. Third, for very large graphs, we plan to investigate approximate shortest-path 523 extraction techniques (e.g., truncated BFS, landmark-based methods) to further reduce preprocess- 524 ing cost while retaining the structural benefits of path-based modeling.



540 USE OF LARGE LANGUAGE MODELS (LLMs)  
541542 Large Language Models were employed to enhance the readability of the manuscript, refine the  
543 phrasing of selected passages, and provide assistance in code debugging. All original content was  
544 produced by the authors; LLMs were used exclusively to improve clarity and presentation.  
545546 ETHICS STATEMENT  
547548 Our study does not involve human subjects or personally identifiable data. The datasets used are  
549 publicly available benchmarks or synthetically generated. We follow the ICLR Code of Ethics and  
550 note that our work raises no foreseeable ethical concerns beyond those inherent to the general study  
551 of machine learning with missing data.  
552553 REPRODUCIBILITY STATEMENT  
554555 We have made every effort to ensure reproducibility. Details of the experimental setup are pro-  
556 vided in Section 5, with dataset descriptions in Appendix C and complete training configurations  
557 in Appendix D. All proofs are included in Appendix A. Anonymous source code to reproduce our  
558 experiments is provided in the supplementary material.  
559560 REFERENCES  
561562 Lada A Adamic and Eytan Adar. Friends and neighbors on the web. *Social Networks*,  
563 25(3):211–230, 2003. ISSN 0378-8733. doi: [https://doi.org/10.1016/S0378-8733\(03\)  
564 00009-1](https://doi.org/10.1016/S0378-8733(03)00009-1). URL [https://www.sciencedirect.com/science/article/pii/  
565 S0378873303000091](https://www.sciencedirect.com/science/article/pii/S0378873303000091).  
566567 Louis Airale, Antonio Longa, Mattia Rigon, Andrea Passerini, and Roberto Passerone. Simple path  
568 structural encoding for graph transformers. *arXiv preprint arXiv:2502.09365*, 2025.  
569570 Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables  
571 for graph identification. *Combinatorica*, 12(4):389–410, 1992.  
572573 Benjamin Paul Chamberlain, Sergey Shirobokov, Emanuele Rossi, Fabrizio Frasca, Thomas  
574 Markovich, Nils Yannick Hammerla, Michael M. Bronstein, and Max Hansmire. Graph neu-  
575 ral networks for link prediction with subgraph sketching. In *The Eleventh International Confer-  
576 ence on Learning Representations*, 2023. URL <https://openreview.net/forum?id=m1oqEOAoZQU>.  
577578 Xueqi Cheng, Yu Wang, Yunchao Liu, Yuying Zhao, Charu C Aggarwal, and Tyler Derr. Edge clas-  
579 sification on graphs: New directions in topological imbalance. In *Proceedings of the Eighteenth  
580 ACM International Conference on Web Search and Data Mining*, pp. 392–400, 2025.  
581582 Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of  
583 gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.  
584585 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to  
586 Algorithms*. MIT Press, Cambridge, MA, 3rd edition, 2009.  
587588 Anuj Dawar and Danny Vagnozzi. Generalizations of k-dimensional weisfeiler–leman stabilization.  
589 *Moscow Journal of Combinatorics and Number Theory*, 9(3):229–252, 2020.  
590591 Bowen Dong, Charu C Aggarwal, and S Yu Philip. The link regression problem in graph streams.  
592 In *2019 IEEE International Conference on Big Data (Big Data)*, pp. 1088–1095. IEEE, 2019.  
593594 Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings  
595 of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*,  
596 pp. 855–864, 2016.  
597

594 Aric Hagberg, Pieter J Swart, and Daniel A Schult. Exploring network structure, dynamics, and  
 595 function using networkx. Technical report, Los Alamos National Laboratory (LANL), Los  
 596 Alamos, NM (United States), 2008.

597

598 Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs.  
 599 In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett  
 600 (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,  
 601 2017a. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf).

602

603 Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs.  
 604 *Advances in neural information processing systems*, 30, 2017b.

605

606 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):  
 607 1735–1780, 1997.

608

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

612

613 Kanchan Jha, Sriparna Saha, and Hiteshi Singh. Prediction of protein–protein interaction using  
 614 graph neural networks. *Scientific Reports*, 12(1):8360, 2022.

615

616 Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43,  
 617 1953.

618

619 Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional net-  
 620 works. *arXiv preprint arXiv:1609.02907*, 2016a.

621

622 Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint  
 623 arXiv:1611.07308*, 2016b.

624

625 Lecheng Kong, Yixin Chen, and Muhan Zhang. Geodesic graph neural network for efficient graph  
 626 representation learning. *Advances in neural information processing systems*, 35:5896–5909,  
 627 2022.

628

629 Juanhui Li, Harry Shomer, Haitao Mao, Shenglai Zeng, Yao Ma, Neil Shah, Jiliang Tang, and  
 630 Dawei Yin. Evaluating graph neural networks for link prediction: Current pitfalls and new bench-  
 631 marking. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and  
 632 Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=YdjWXrdOTH>.

633

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

637

638 Jinbi Liang, Cunlai Pu, Xiangbo Shu, Yongxiang Xia, and Chengyi Xia. Line graph neural networks  
 639 for link weight prediction. *Physica A: Statistical Mechanics and its Applications*, pp. 130406,  
 640 2025.

641

642 David Liben-Nowell and Jon Kleinberg. The link prediction problem for social networks. In *Pro-  
 643 ceedings of the Twelfth International Conference on Information and Knowledge Management,  
 644 CIKM '03*, pp. 556–559, New York, NY, USA, 2003. Association for Computing Machinery.  
 645 ISBN 1581137230. doi: 10.1145/956863.956972. URL <https://doi.org/10.1145/956863.956972>.

643

644 Moritz Licher, Simon Raßmann, and Pascal Schweitzer. Computational complexity of the  
 645 weisfeiler-leman dimension. In *33rd EACSL Annual Conference on Computer Science Logic  
 646 (CSL 2025)*, pp. 13–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.

647

Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: statistical  
 648 mechanics and its applications*, 390(6):1150–1170, 2011.

648 Aditya Krishna Menon and Charles Elkan. Link prediction via matrix factorization. In *Machine*  
 649 *Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2011,*  
 650 *Athens, Greece, September 5-9, 2011, Proceedings, Part II* 22, pp. 437–452. Springer, 2011.  
 651

652 Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav  
 653 Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks.  
 654 In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 4602–4609, 2019.  
 655

656 M. E. J. Newman. Clustering and preferential attachment in growing networks. *Phys. Rev. E*,  
 657 64:025102, Jul 2001. doi: 10.1103/PhysRevE.64.025102. URL <https://link.aps.org/doi/10.1103/PhysRevE.64.025102>.  
 658

659 Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational  
 660 machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2015.  
 661

662 Andrea Rossi, Denilson Barbosa, Donatella Firmani, Antonio Matinata, and Paolo Merialdo.  
 663 Knowledge graph embedding for link prediction: A comparative analysis. *ACM Transactions*  
 664 *on Knowledge Discovery from Data (TKDD)*, 15(2):1–49, 2021.  
 665

666 Christian Sommer. Shortest-path queries in static networks. *ACM Computing Surveys (CSUR)*, 46  
 667 (4):1–31, 2014.  
 668

669 Balasubramaniam Srinivasan and Bruno Ribeiro. On the equivalence between positional node em-  
 670 beddings and structural graph representations. *arXiv preprint arXiv:1910.00452*, 2019.  
 671

672 Astrit Tola, Jack Myrick, and Baris Coskunuzer. PROXI: Challenging the GNNs for link predic-  
 673 tion. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=u9EHndbiVw>.  
 674

675 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
 676 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*  
 677 *tion processing systems*, 30, 2017.  
 678

679 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Ben-  
 680 gio, et al. Graph attention networks. *stat*, 1050(20):10–48550, 2017.  
 681

682 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua  
 683 Bengio. Graph attention networks. In *International Conference on Learning Representations*,  
 684 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.  
 685

686 Haorui Wang, Haoteng Yin, Muhan Zhang, and Pan Li. Equivariant and stable positional encoding  
 687 for more powerful graph neural networks. In *International Conference on Learning Representa-*  
 688 *tions*, 2022. URL <https://openreview.net/forum?id=e95i1IHcWj>.  
 689

690 Xiyuan Wang, Haotong Yang, and Muhan Zhang. Neural common neighbor with completion for  
 691 link prediction. In *The Twelfth International Conference on Learning Representations*, 2024.  
 692 URL <https://openreview.net/forum?id=sNFLN3itAd>.  
 693

694 Xuhong Wang, Ding Lyu, Mengjian Li, Yang Xia, Qi Yang, Xinwen Wang, Xinguang Wang, Ping  
 695 Cui, Yupu Yang, Bowen Sun, et al. Apan: Asynchronous propagation attention network for  
 696 real-time temporal graph embedding. In *Proceedings of the 2021 international conference on*  
 697 *management of data*, pp. 2628–2638, 2021.  
 698

699 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural  
 700 networks?, 2019.  
 701

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

Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and  
 Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in neural*  
*information processing systems*, 34:28877–28888, 2021.

702 Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec.  
703 Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the*  
704 *24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 974–  
705 983, 2018.

706 Seongjun Yun, Seoyoon Kim, Junhyun Lee, Jaewoo Kang, and Hyunwoo J Kim. Neo-gnns: Neigh-  
707 borhood overlap-aware graph neural networks for link prediction. *Advances in Neural Information*  
708 *Processing Systems*, 34:13683–13694, 2021.

709

710 Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *Advances in neural*  
711 *information processing systems*, 31, 2018.

712 Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. Labeling trick: A theory of using  
713 graph neural networks for multi-node representation learning. *Advances in Neural Information*  
714 *Processing Systems*, 34:9061–9073, 2021.

715

716 Tao Zhou. Progresses and challenges in link prediction. *Iscience*, 24(11), 2021.

717 Tao Zhou, Linyuan Lü, and Yi-Cheng Zhang. Predicting missing links via local information. *The*  
718 *European Physical Journal B*, 71(4):623–630, October 2009. ISSN 1434-6036. doi: 10.1140/  
719 *epjb/e2009-00335-8*. URL <http://dx.doi.org/10.1140/EPJB/E2009-00335-8>.

720

721 Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. Neural bellman-ford net-  
722 works: A general graph neural network framework for link prediction. *Advances in neural infor-*  
723 *mation processing systems*, 34:29476–29490, 2021.

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A PROOFS  
757758 **Proposition 4.4** SP4LP does not suffer from the automorphic node problem.  
759760 *Proof.* According to Proposition 2.9, a model  $M$  suffers from the *automorphic node problem* if, for  
761 any graph  $G = (V, E, \mathbf{X}^0)$ , for any pairs of links  $(u, v), (u', v') \in V \times V$ , for any  $F \in M$  it holds  
762 that:  
763

764 
$$(u, v) \not\simeq (u', v'), \quad u \simeq u', \quad v \simeq v', \quad \text{and} \quad F((u, v), G) \neq F((u', v'), G).$$

765 In order to prove that SP4LP does not suffer from the automorphic node problem, it suffices to  
766 provide an example of a graph  $G = (V, E, X^0)$  and node pairs  $(u, v), (u', v') \in V \times V$  such that:  
767

768 
$$(u, v) \not\simeq (u', v'), \quad u \simeq u', \quad v \simeq v', \quad \text{and} \quad \text{SP4LP}((u, v), G) \neq \text{SP4LP}((u', v'), G).$$

769 Such an example is provided in Figure 1: the shortest path from  $v$  to  $u'$  consists of  $v$ , an orange  
770 node, and  $u'$ , whereas the shortest path from  $v$  to  $u$  includes  $v$ , an orange node, a yellow node,  
771 another orange node, and finally  $u$ . Thus, simply using a summation as function  $\phi$ , leads to distinct  
772 representations for links  $(v, u)$  and  $(v, u')$ .  $\square$   
773774 **Theorem 4.5** SP4LP is strictly more expressive than Pure GNN, NCN, BUDDY, NBFnet and Neo-  
775 GNN.  
776777 *Proof.* We proceed by prove each comparison separately.  
778779 • SP4LP is strictly more expressive than Pure GNN.780 We prove this by noting that SP4LP architecture (Equation 8) generalizes that of pure  
781 GNNs, meaning that SP4LP can simulate any pure GNN by simply ignoring the shortest  
782 path information. Thus, for any pair of non-automorphic links that a specific pure GNN can  
783 distinguish, there exists a configuration of SP4LP that distinguishes them as well. We can  
784 conclude that SP4LP is strictly more expressive than pure GNNs considering as example  
785 of pair of links indistinguishable by GNNs but distinguishable by SP4LP the one provided  
786 in the proof of Proposition 4.4.  
787788 • SP4LP is strictly more expressive than NCN.789 We prove this in two steps: (1) When two links share no common neighbors, NCN on  
790 them reduces to a pure GNN. As we have already proved that SP4LP is strictly more  
791 expressive than pure GNNs, it follows that SP4LP is also strictly more expressive than  
792 NCN in this case. (2) When the links have common neighbors, setting AGG as summation,  
793  $\rho$  as the Hadamard product between the endpoint representations and the concatenation  
794 with the result of the aggregation, and  $\phi$  as the identity function, SP4LP reduces exactly to  
795 NCN. Therefore, if NCN can distinguish two links under some configuration, SP4LP can  
796 as well. By definition 2.10), this implies that SP4LP is more expressive than NCN. We  
797 can conclude that SP4LP is strictly more expressive than NCN considering the example in  
798 Figure 4. The graph is regular and thus all nodes receive the same embedding from a GNN.  
799 Consider nodes  $u$  and  $v$ :  $N(u) \cap N(v) = N(u) \cap N(v') = \emptyset$ . In this case, NCN reduces to  
800 a pure GNN and is thus unable to distinguish the links  $(u, v)$  and  $(u', v')$ . Now, let  $\mathbf{x} \in \mathbb{R}^d$   
801 be the representation assigned to every node by a GNN. Then, the representation of the  
802 shortest path  $\mathcal{P}_G^*(u, v)$  is simply  $(\mathbf{x}, \mathbf{x}, \mathbf{x})$ , while  $\mathcal{P}_G^*(u', v') = (\mathbf{x}, \mathbf{x}, \mathbf{x}, \mathbf{x})$ . Even using  
803 a simple sum as aggregation function, SP4LP successfully distinguishes between the two  
804 links.  
805806 • SP4LP is strictly more expressive than Neo-GNN and BUDDY.807 We have already shown that SP4LP is strictly more expressive than NCN. In Theorem 2  
808 of the NCN paper Wang et al. (2024), it has been proved that NCN is more expressive  
809 than both Neo-GNN and BUDDY. It follows that SP4LP is strictly more expressive than  
Neo-GNN and BUDDY as well.

810  
811  
812  
813  
814

- **SP4LP** is strictly more expressive than **NBFNet**.

We prove that **NBFNet** is as expressive as a pure GNN. Therefore, since we have already proven that **SP4LP** is strictly more expressive than pure GNNs, it immediately follows that **SP4LP** is also strictly more expressive than **NBFNet**.

To complete the argument, we prove that **NBFNet** is as expressive as a pure GNN. First of all we report the formulation of **NBFNet** for simple undirected graph. Given a graph  $G = (V, E, \mathbf{X}^0)$ , **NBFNet** assigns a representation  $\mathbf{x}(u, v)$  to each edge  $(u, v) \in E$ . The iterative update rule follows a message-passing scheme:

815  
816  
817  
818

$$\begin{aligned} \mathbf{x}_{(u,v)}^{(0)} &= \text{INDICATOR}(\mathbf{x}_u^0, \mathbf{x}_v^0), \\ \mathbf{x}_{(u,v)}^{(l)} &= \text{AGGREGATE} \left( \left\{ \text{MESSAGE}(\mathbf{h}_{(i,j)}^{(l-1)}) \mid (i, j) \in N(u, v) \right\} \cup \{h_{(u,v)}^{(0)}\} \right) \end{aligned} \quad (9)$$

819  
820  
821  
822

823 where **INDICATOR** assigns an initial representation based on the nodes  $u, v \in V$  and  
824  $N(u, v)$  is the set of edges incident to  $(u, v)$ .

825 We prove that, at any layer  $l$ , the representations of the two links produced by a pure GNN  
826 are equal if and only if also the ones produced by **NBFNet** are equal, i.e.,

827  
828

$$\mathbf{x}_{(u,v)}^{\text{GNN}^l} = \mathbf{x}_{(i,j)}^{\text{GNN}^l} \Leftrightarrow \mathbf{x}_{(u,v)}^{\text{NBF}^l} = \mathbf{x}_{(i,j)}^{\text{NBF}^l} \quad \forall l \quad (10)$$

829

830 where  $\mathbf{x}_{(u,v)}^{\text{NBF}^l}$  and  $\mathbf{x}_{(i,j)}^{\text{NBF}^l}$  are calculated via Equation 9, while  $\mathbf{x}_{(u,v)}^{\text{GNN}^l}$  and  $\mathbf{x}_{(i,j)}^{\text{GNN}^l}$  are calculated  
831 following the standard message passing scheme reported below:

832  
833  
834  
835  
836

$$\begin{aligned} \mathbf{x}_v^{\text{GNN}^0} &= \mathbf{x}_v^0, \\ \mathbf{x}_v^{\text{GNN}^l} &= \text{COMB} \left( \mathbf{x}_v^{\text{GNN}^{l-1}}, \text{AGG} \left( \{\mathbf{x}_u^{\text{GNN}^{l-1}} \mid u \in N(v)\} \right) \right) \\ \mathbf{x}_{(u,v)}^{\text{GNN}^l} &= g(\mathbf{x}_u^{\text{GNN}^l}, \mathbf{x}_v^{\text{GNN}^l}) \end{aligned} \quad (11)$$

837

838 Let the functions **INDICATOR**, **AGGREGATE** and **MESSAGE** of Equation 9, as well as the  
839 functions **COMB**, **AGG** and  $g$  be injective. We prove Equation 10 by induction on the  
840 number of layer  $l$ .

841  
842

**Base Case:**  $l = 0$

843  
844  
845  
846

$$\begin{aligned} \mathbf{x}_{(u,v)}^{\text{GNN}^0} = \mathbf{x}_{(i,j)}^{\text{GNN}^0} &\stackrel{(11)}{\iff} g(\mathbf{x}_u^0, \mathbf{x}_v^0) = g(\mathbf{x}_i^0, \mathbf{x}_j^0) \stackrel{\text{inj}}{\iff} (\mathbf{x}_u^0, \mathbf{x}_v^0) = (\mathbf{x}_i^0, \mathbf{x}_j^0) \stackrel{\text{inj}}{\iff} \\ &\stackrel{\text{inj}}{\iff} \text{INDICATOR}(\mathbf{x}_u^0, \mathbf{x}_v^0) = \text{INDICATOR}(\mathbf{x}_i^0, \mathbf{x}_j^0) \stackrel{(9)}{\iff} \mathbf{x}_{(u,v)}^{\text{NBF}^0} = \mathbf{x}_{(i,j)}^{\text{NBF}^0} \end{aligned}$$

847  
848  
849

**Inductive Step** We assume Equation 10 holds for  $l - 1$  and prove it holds for  $l$ .

In particular, we want to prove

850  
851  
852

$$\mathbf{x}_{(u,v)}^{\text{GNN}^l} = \mathbf{x}_{(i,j)}^{\text{GNN}^l} \iff \mathbf{x}_{(u,v)}^{\text{NBF}^l} = \mathbf{x}_{(i,j)}^{\text{NBF}^l} \quad (12)$$

using the inductive hypothesis

853  
854

$$\mathbf{x}_{(u,v)}^{\text{GNN}^{l-1}} = \mathbf{x}_{(i,j)}^{\text{GNN}^{l-1}} \iff \mathbf{x}_{(u,v)}^{\text{NBF}^{l-1}} = \mathbf{x}_{(i,j)}^{\text{NBF}^{l-1}} \quad (13)$$

855 Applying Equation 11 to the left-hand side of Equation 12 we get

856  
857  
858  
859  
860  
861  
862  
863

$$\begin{aligned} \mathbf{x}_{(u,v)}^{\text{GNN}^l} &= \mathbf{x}_{(i,j)}^{\text{GNN}^l} \\ &\stackrel{(11)}{\iff} \\ g(\text{COMB}(\mathbf{x}_u^{\text{GNN}^{l-1}}, \text{AGG}(\{\mathbf{x}_x^{\text{GNN}^{l-1}} \mid x \in N(u)\})), \text{COMB}(\mathbf{x}_v^{\text{GNN}^{l-1}}, \text{AGG}(\{\mathbf{x}_y^{\text{GNN}^{l-1}} \mid y \in N(u)\}))) &= \\ g(\text{COMB}(\mathbf{x}_i^{\text{GNN}^{l-1}}, \text{AGG}(\{\mathbf{x}_m^{\text{GNN}^{l-1}} \mid m \in N(i)\})), \text{COMB}(\mathbf{x}_j^{\text{GNN}^{l-1}}, \text{AGG}(\{\mathbf{x}_n^{\text{GNN}^{l-1}} \mid n \in N(j)\}))) \end{aligned}$$

Given the injectivity of  $g$ , COMB and AGG, this is equivalent to

$$\begin{aligned}
& \mathbf{x}_u^{\text{GNN}^{l-1}} = \mathbf{x}_i^{\text{GNN}^{l-1}} \wedge \{\mathbf{x}_x^{\text{GNN}^{l-1}} \mid x \in N(u)\} = \{\mathbf{x}_m^{\text{GNN}^{l-1}} \mid m \in N(i)\} \wedge \\
& \wedge \mathbf{x}_v^{\text{GNN}^{l-1}} = \mathbf{x}_j^{\text{GNN}^{l-1}} \wedge \{\mathbf{x}_y^{\text{GNN}^{l-1}} \mid y \in N(v)\} = \{\mathbf{x}_n^{\text{GNN}^{l-1}} \mid n \in N(j)\} \\
& \iff \\
& \{(\mathbf{x}_u^{\text{GNN}^{l-1}}, \mathbf{x}_x^{\text{GNN}^{l-1}}) \mid x \in N(u)\} = \{(\mathbf{x}_i^{\text{GNN}^{l-1}}, \mathbf{x}_m^{\text{GNN}^{l-1}}) \mid m \in N(i)\} \\
& \wedge \\
& \{(\mathbf{x}_v^{\text{GNN}^{l-1}}, \mathbf{x}_y^{\text{GNN}^{l-1}}) \mid y \in N(v)\} = \{(\mathbf{x}_j^{\text{GNN}^{l-1}}, \mathbf{x}_n^{\text{GNN}^{l-1}}) \mid n \in N(j)\} \\
& \iff \\
& \{(\mathbf{x}_u^{\text{GNN}^{l-1}}, \mathbf{x}_x^{\text{GNN}^{l-1}}) \mid x \in N(u)\} \cup \{(\mathbf{x}_v^{\text{GNN}^{l-1}}, \mathbf{x}_y^{\text{GNN}^{l-1}}) \mid y \in N(v)\} \\
& = \\
& \{(\mathbf{x}_i^{\text{GNN}^{l-1}}, \mathbf{x}_m^{\text{GNN}^{l-1}}) \mid m \in N(i)\} \cup \{(\mathbf{x}_j^{\text{GNN}^{l-1}}, \mathbf{x}_n^{\text{GNN}^{l-1}}) \mid n \in N(j)\}.
\end{aligned}$$

By Definition of  $N(u, v)$  (Equation 9), this is equivalent to

$$\begin{aligned}
& \{(\mathbf{x}_w^{\text{GNN}^{l-1}}, \mathbf{x}_t^{\text{GNN}^{l-1}}) \mid (w, t) \in N(u, v)\} = \{(\mathbf{x}_a^{\text{GNN}^{l-1}}, \mathbf{x}_b^{\text{GNN}^{l-1}}) \mid (a, b) \in N(i, j)\} \\
& \qquad \qquad \qquad \xrightleftharpoons{\text{inj}} \\
& \{g(\mathbf{x}_w^{\text{GNN}^{l-1}}, \mathbf{x}_t^{\text{GNN}^{l-1}}) \mid (w, t) \in N(u, v)\} = \{g(\mathbf{x}_a^{\text{GNN}^{l-1}}, \mathbf{x}_b^{\text{GNN}^{l-1}}) \mid (a, b) \in N(i, j)\} \\
& \qquad \qquad \qquad \xrightleftharpoons{(11)} \\
& \{\mathbf{x}_{(w, t)}^{\text{GNN}^{l-1}} \mid (w, t) \in N(u, v)\} = \{\mathbf{x}_{(a, b)}^{\text{GNN}^{l-1}} \mid (a, b) \in N(i, j)\} \\
& \qquad \qquad \qquad \xrightleftharpoons{\text{IND. HYP.(13)}} \\
& \{\mathbf{x}_{(w, t)}^{\text{NBF}^{l-1}} \mid (w, t) \in N(u, v)\} = \{\mathbf{x}_{(a, b)}^{\text{NBF}^{l-1}} \mid (a, b) \in N(i, j)\}.
\end{aligned}$$

Using the hypotheses of injective AGG and MESSAGE, this is equivalent to:

$$\text{AGG}(\{\text{MESSAGE}(\mathbf{x}_{(w,t)}^{\text{NBF}^{l-1}}) \mid (w,t) \in N(u,v)\}) = \text{AGG}(\{\text{MESSAGE}(\mathbf{x}_{(a,b)}^{\text{NBF}^{l-1}}) \mid (a,b) \in N(i,j)\})$$

$$\iff^{(9)}$$

$$\mathbf{x}_{(u,v)}^{\text{NBF}^l} = \mathbf{x}_{(i,j)}^{\text{NBF}^l} \quad (14)$$

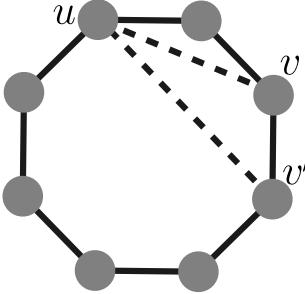
which complete the proof.  $\square$

## B ADDITIONAL RESULTS

We complement the main results of Section 5 with additional tables reporting the standard deviation computed over five different random seeds, to better assess the stability of each method.

**Real-world Datasets: Results with Standard Deviations** Table 4 and Table 5 expand the main results in Table 2 by including both the mean and standard deviation of MRR and Hits@ $K$  across runs.

**Ablation Study: Results with Standard Deviations** Similarly, Table 6 complements the ablation results in Table 3 by reporting mean and standard deviation for Cora, Citeseer, and Pubmed.

Figure 4: Links  $(u, v)$ ,  $(u, v')$  are not distinguished by NCN while are distinguished by SP4LP.

Models	Cora MRR	Citeseer MRR	Pubmed MRR	Ogbl-ddi MRR	Ogbl-collab MRR	Ogbl-ppa MRR	Ogbl-Citation2 MRR
Node2Vec	14.47 ( $\pm 0.60$ )	21.17 ( $\pm 1.01$ )	3.94 ( $\pm 0.24$ )	11.14 ( $\pm 0.95$ )	4.68 ( $\pm 0.08$ )	18.33 ( $\pm 0.10$ )	14.67 ( $\pm 0.18$ )
MF	6.20 ( $\pm 1.42$ )	7.80 ( $\pm 0.79$ )	4.46 ( $\pm 0.32$ )	13.99 ( $\pm 0.47$ )	4.89 ( $\pm 0.25$ )	22.47 ( $\pm 1.53$ )	8.72 ( $\pm 2.60$ )
MLP	13.52 ( $\pm 0.65$ )	22.62 ( $\pm 0.55$ )	6.41 ( $\pm 0.25$ )	N/A	5.37 ( $\pm 0.14$ )	0.98 ( $\pm 0.00$ )	16.32 ( $\pm 0.07$ )
GCN	16.61 ( $\pm 0.30$ )	21.09 ( $\pm 0.88$ )	7.13 ( $\pm 0.27$ )	13.46 ( $\pm 0.34$ )	6.09 ( $\pm 0.38$ )	26.94 ( $\pm 0.48$ )	19.98 ( $\pm 0.35$ )
GAT	13.84 ( $\pm 0.68$ )	19.58 ( $\pm 0.84$ )	4.95 ( $\pm 0.14$ )	12.92 ( $\pm 0.39$ )	4.18 ( $\pm 0.33$ )	OOM	OOM
SAGE	14.74 ( $\pm 0.69$ )	21.09 ( $\pm 1.15$ )	9.40 ( $\pm 0.70$ )	12.60 ( $\pm 0.72$ )	5.53 ( $\pm 0.50$ )	27.27 ( $\pm 0.30$ )	22.05 ( $\pm 0.12$ )
GAE	18.32 ( $\pm 0.41$ )	25.25 ( $\pm 0.82$ )	5.27 ( $\pm 0.25$ )	3.49 ( $\pm 1.73$ )	OOM	OOM	OOM
SEAL	10.67 ( $\pm 3.46$ )	13.16 ( $\pm 1.66$ )	5.88 ( $\pm 0.53$ )	9.99 ( $\pm 0.90$ )	6.43 ( $\pm 0.32$ )	29.71 ( $\pm 0.71$ )	20.60 ( $\pm 1.28$ )
BUDDY	13.71 ( $\pm 0.59$ )	22.84 ( $\pm 0.36$ )	7.56 ( $\pm 0.18$ )	12.43 ( $\pm 0.50$ )	5.67 ( $\pm 0.36$ )	27.70 ( $\pm 0.33$ )	19.17 ( $\pm 0.10$ )
Neo-GNN	13.95 ( $\pm 0.39$ )	17.34 ( $\pm 0.84$ )	7.74 ( $\pm 0.30$ )	10.86 ( $\pm 2.16$ )	5.23 ( $\pm 0.90$ )	21.68 ( $\pm 1.14$ )	16.12 ( $\pm 0.25$ )
NCN	14.66 ( $\pm 0.95$ )	28.65 ( $\pm 1.21$ )	5.84 ( $\pm 0.22$ )	12.86 ( $\pm 0.78$ )	5.09 ( $\pm 0.38$ )	35.06 ( $\pm 0.26$ )	23.35 ( $\pm 0.28$ )
NCNC	14.98 ( $\pm 1.00$ )	24.10 ( $\pm 0.65$ )	8.58 ( $\pm 0.59$ )	>24h	4.73 ( $\pm 0.86$ )	33.52 ( $\pm 0.26$ )	19.61 ( $\pm 0.54$ )
NBFNet	13.56 ( $\pm 0.58$ )	14.29 ( $\pm 0.80$ )	>24h	>24h	OOM	OOM	OOM
PEG	15.73 ( $\pm 0.39$ )	21.01 ( $\pm 0.77$ )	4.40 ( $\pm 0.41$ )	12.05 ( $\pm 1.14$ )	4.83 ( $\pm 0.21$ )	OOM	OOM
LPFORMER	16.80 ( $\pm 0.52$ )	26.34 ( $\pm 0.67$ )	9.99 ( $\pm 0.52$ )	13.20 ( $\pm 0.54$ )	7.62 ( $\pm 0.26$ )	40.25 ( $\pm 0.24$ )	24.70 ( $\pm 0.55$ )
SP4LP	17.27 ( $\pm 0.57$ )	41.08 ( $\pm 1.84$ )	10.87 ( $\pm 0.31$ )	15.00 ( $\pm 0.57$ )	9.46 ( $\pm 0.55$ )	36.45 ( $\pm 1.21$ )	24.91 ( $\pm 0.41$ )

Table 4: MRR results across all datasets, following the HeaRT evaluation setting Li et al. (2023). The top three results for each metric are highlighted using **first**, **second**, and **third**. **OOM** indicates that the model ran out of memory, while **>24h** denotes that the method did not complete within 24 hours.

## C DATASETS STATISTICS

Table 7 summarizes the main datasets used in our link prediction experiments. Cora, Citeseer, and Pubmed are well-known citation networks frequently used as benchmarks for graph-based learning methods. These datasets are relatively small, both in the number of nodes and edges. In contrast, the datasets from the Open Graph Benchmark (OGB), namely ogbl-collab and ogb-ddi, are substantially larger and more complex, offering challenging scenarios for evaluating model scalability and performance on large-scale graphs.

For all datasets, we adopt the splits provided by the Li et al. (2023) setting.

## D EXPERIMENTAL SETTINGS

This section outlines the experimental setup used to evaluate all models. We describe the computational resources and the hyperparameter search space. Moreover for SP4LP we include details regarding how the calculation of the shortest path is performed. Details are reported below.

**Computational Resources** All experiments were conducted on a workstation running Ubuntu 22.04 with an AMD Ryzen 9 7950X CPU (32 threads), 124GB of RAM, and two NVIDIA GeForce RTX 4090 GPUs (24GB each).

**Hyperparameters** All models are tuned using a grid search over learning rate  $\in [1 \times 10^{-2}, 1 \times 10^{-3}]$ , dropout  $\in [0, 0.7]$ , weight decay  $\in [0, 10^{-4}, 10^{-7}]$ , number of GNN layers  $\in \{1, 2, 3\}$ , hid-

972	Models	Cora Hits@10	Citeseer Hits@10	Pubmed Hits@10	Ogbl-ddi Hits@20	Ogbl-collab Hits@20	Ogbl-ppa Hits@20	Ogbl-Citation2 Hits@20
973	Node2Vec	32.77 ( $\pm$ 1.29)	45.82 ( $\pm$ 2.01)	8.51 ( $\pm$ 0.77)	63.63 ( $\pm$ 2.05)	16.84 ( $\pm$ 0.17)	53.42 ( $\pm$ 0.11)	42.68 ( $\pm$ 0.20)
974	MF	15.26 ( $\pm$ 3.39)	16.72 ( $\pm$ 1.99)	9.42 ( $\pm$ 0.80)	59.50 ( $\pm$ 1.68)	18.86 ( $\pm$ 0.40)	70.71 ( $\pm$ 4.82)	29.64 ( $\pm$ 7.30)
975	MLP	31.01 ( $\pm$ 1.71)	48.02 ( $\pm$ 1.79)	15.04 ( $\pm$ 0.67)	N/A	16.15 ( $\pm$ 0.27)	1.47 ( $\pm$ 0.00)	43.15 ( $\pm$ 0.10)
976	GCN	36.26 ( $\pm$ 1.14)	47.23 ( $\pm$ 1.88)	15.22 ( $\pm$ 0.57)	64.76 ( $\pm$ 1.45)	<b>22.48</b> ( $\pm$ 0.81)	68.38 ( $\pm$ 0.73)	51.72 ( $\pm$ 0.46)
977	GAT	32.89 ( $\pm$ 1.27)	45.30 ( $\pm$ 1.30)	9.99 ( $\pm$ 0.64)	<b>66.83</b> ( $\pm$ 2.23)	18.30 ( $\pm$ 1.42)	OOM	OOM
978	SAGE	34.65 ( $\pm$ 1.47)	48.75 ( $\pm$ 1.85)	<b>20.54</b> ( $\pm$ 1.40)	<b>67.19</b> ( $\pm$ 1.18)	21.26 ( $\pm$ 1.32)	69.49 ( $\pm$ 0.43)	53.13 ( $\pm$ 0.15)
979	GAE	<b>37.95</b> ( $\pm$ 1.24)	49.65 ( $\pm$ 1.48)	10.50 ( $\pm$ 0.46)	17.81 ( $\pm$ 9.80)	OOM	OOM	OOM
980	SEAL	24.27 ( $\pm$ 6.74)	27.37 ( $\pm$ 3.20)	12.47 ( $\pm$ 1.23)	49.74 ( $\pm$ 2.39)	21.57 ( $\pm$ 0.38)	76.77 ( $\pm$ 0.94)	48.62 ( $\pm$ 1.93)
981	BUDDY	30.40 ( $\pm$ 1.18)	48.35 ( $\pm$ 1.18)	16.78 ( $\pm$ 0.53)	58.71 ( $\pm$ 1.63)	<b>23.35</b> ( $\pm$ 0.73)	71.50 ( $\pm$ 0.68)	47.81 ( $\pm$ 0.37)
982	Neo-GNN	31.27 ( $\pm$ 0.72)	41.74 ( $\pm$ 1.18)	17.88 ( $\pm$ 0.71)	51.94 ( $\pm$ 10.33)	21.03 ( $\pm$ 3.39)	64.81 ( $\pm$ 2.26)	43.17 ( $\pm$ 0.53)
983	NCN	35.14 ( $\pm$ 1.04)	<b>53.41</b> ( $\pm$ 1.46)	13.22 ( $\pm$ 0.56)	<b>65.82</b> ( $\pm$ 2.66)	20.84 ( $\pm$ 1.31)	<b>81.89</b> ( $\pm$ 0.31)	<b>53.76</b> ( $\pm$ 0.20)
984	NCNC	<b>36.70</b> ( $\pm$ 1.57)	<b>53.72</b> ( $\pm$ 0.97)	<b>18.81</b> ( $\pm$ 1.16)	>24h	20.49 ( $\pm$ 3.97)	<b>82.24</b> ( $\pm$ 0.40)	51.69 ( $\pm$ 1.48)
985	NBFNet	31.12 ( $\pm$ 0.75)	31.39 ( $\pm$ 1.34)	>24h	>24h	OOM	OOM	OOM
986	PEG	36.03 ( $\pm$ 0.75)	45.56 ( $\pm$ 1.38)	8.70 ( $\pm$ 1.26)	50.12 ( $\pm$ 6.55)	18.29 ( $\pm$ 1.06)	OOM	OOM
987	LPFORMER	33.27 ( $\pm$ 1.33)	51.58 ( $\pm$ 1.83)	22.71 ( $\pm$ 1.30)	35.23 ( $\pm$ 0.37)	<b>23.05</b> ( $\pm$ 0.16)	<b>84.01</b> ( $\pm$ 0.10)	<b>57.30</b> ( $\pm$ 0.50)
988	SP4LP	<b>38.52</b> ( $\pm$ 1.19)	<b>66.28</b> ( $\pm$ 0.63)	<b>23.01</b> ( $\pm$ 0.39)	47.96 ( $\pm$ 3.82)	20.00 ( $\pm$ 1.20)	76.9 ( $\pm$ 1.11)	<b>55.45</b> ( $\pm$ 0.92)

Table 5: Hits@K (%) results across all datasets, following the HeaRT evaluation setting Li et al. (2023). The top three results for each metric are highlighted using **first**, **second**, and **third**. *OOM* indicates that the model ran out of memory, while *>24h* denotes that the method did not complete within 24 hours.

992 993 994 995 996 997	Models	Cora		Citeseer		Pubmed	
		MRR	Hits@10	MRR	Hits@10	MRR	Hits@10
GNN + SP len.	14.21 ( $\pm$ 1.44)	33.43 ( $\pm$ 2.69)	20.90 ( $\pm$ 0.79)	47.82 ( $\pm$ 1.11)	7.12 ( $\pm$ 0.41)	5.63 ( $\pm$ 0.52)	
Sequence Model	16.86 ( $\pm$ 1.26)	36.03 ( $\pm$ 1.75)	27.45 ( $\pm$ 1.55)	54.20 ( $\pm$ 2.35)	8.58 ( $\pm$ 0.75)	12.87 ( $\pm$ 0.85)	
SP4LP	<b>17.27</b> ( $\pm$ 0.57)	<b>38.52</b> ( $\pm$ 1.19)	<b>41.08</b> ( $\pm$ 1.84)	<b>66.28</b> ( $\pm$ 0.63)	<b>10.87</b> ( $\pm$ 0.31)	<b>23.01</b> ( $\pm$ 0.39)	

Table 6: Ablation study results (%). MRR and Hits@K with mean and std. deviations over 5 runs with different seeds.

den dimensions  $\in \{32, 64, 128, 256\}$  and prediction layers  $\in \{1, 2, 3\}$ . For large-scale datasets, we follow the reduced search space adopted in Li et al. (2023) to avoid excessive compute. For SP4LP, we additionally explore the choice of GNN component  $\in \{\text{GCN, GraphSAGE, GAT}\}$  and sequence model  $\in \{\text{LSTM, Transformer}\}$ , the best models are shown in Table 8. The best hyperparameters are selected based on validation performance. All reported metrics are averaged over 5 different seeds.

**Shortest path calculation for SP4LP** To compute shortest paths for SP4LP, we use the `networkx.shortest_path` BFS implementation (Hagberg et al., 2008). Importantly, we compute and cache shortest paths only for the train/validation/test links provided by the HeaRT benchmark splits. If multiple links share an endpoint, only one BFS is required for that source node. This keeps preprocessing time low: in all datasets, the total number of BFS sources is small, and the full preprocessing step completes in seconds to a few minutes. The cached paths are reused across all training epochs and inference batches. No all-pairs shortest path computation is performed.

In our implementation of SP4LP, the sequential encoder can be instantiated as a Transformer with learnable positional encodings. We use PyTorch’s TransformerEncoder, composed of TransformerEncoderLayer blocks with 4-head self-attention, feedforward sublayers, ReLU activation, and dropout. We tune the number of layers  $\in \{1, 2\}$ , attention heads  $\in \{2, 3, 4\}$ , and feedforward dimensionality  $\in \{32, 64, 128\}$ . A trainable positional embedding matrix is added to node embeddings to preserve path order. Variable-length paths are handled via a source key padding mask, and the output is aggregated using masked mean pooling followed by layer normalization. The resulting path representation is then combined with the GNN-derived embeddings of the source and target nodes to compute the final link score.

For the LSTM-based encoder, we implement both unidirectional and bidirectional variants using PyTorch’s `nn.LSTM`. Input sequences are packed with `pack_padded_sequence` to handle variable-length

Dataset	Cora	Citeseer	Pubmed	ogbl-collab	ogbl-ddi	ogbl-ppa	ogbl-citation2
#Nodes	2708	3327	18717	235868	4267	576289	2927963
#Edges	5278	4676	44327	1285465	1334889	30326273	30561187

Table 7: Dataset statistics.

Dataset	GNN model	Sequence model
Cora	GCN	Transformer
Citeseer	GCN	LSTM
Pubmed	SAGE	Transformer
ogbl-collab	SAGE	Transformer
ogbl-ddi	GCN	Transformer

Table 8: Best GNN and sequence models selected via hyperparameter tuning.

paths. We tune the number of layers  $\in \{1, 2\}$  and hidden size  $\in \{32, 64, 128\}$ . The final hidden state (or the concatenation of forward and backward states in the bidirectional case) is used as the path representation and combined with the GNN-based node embeddings for link prediction. Node embeddings are obtained via a GNN encoder selected from GCN, GAT, or GraphSAGE, depending on the experimental setting.

## E COMPUTATIONAL COMPLEXITY

In this section, we analyze the computational complexity of our proposed SP4LP model for link prediction, following the formalism and notation introduced in Wang et al. (2024). This framework allows a direct comparison with prior approaches such as GAE, Neo-GNN, BUDDY, PEG, SEAL, NCN, and NCNC.

Let  $n$  be the number of nodes,  $m$  the number of edges,  $\Delta$  the maximum degree,  $d$  the dimensionality of node features or embeddings,  $t$  the number of target links to predict, and  $k$  the length of the shortest path between node pairs (typically  $k \ll n$ ). We express total time complexity as  $\mathcal{O}(B + C \cdot t)$ , where  $B$  is the precomputation cost (independent of  $t$ ), and  $C$  is the per-link cost.

Our SP4LP method follows the GNN-then-Structural-Feature (SF) paradigm Wang et al. (2024), applying a Message Passing Neural Network once across the entire graph to compute node embeddings, then leveraging shortest path extraction combined with sequence modeling for each candidate link.

Following the paradigm, the precomputation cost of SP4LP is:

$$B = \mathcal{O}(n\Delta d + nd^2 + T_{\text{sp}})$$

where:  $n\Delta d$  accounts for neighborhood aggregation in the GNN,  $nd^2$  represents linear transformations in GNN layers,  $T_{\text{sp}}$  is the cost of computing shortest paths (e.g., via BFS).

For sparse graphs, single-source BFS runs in  $\mathcal{O}(m)$  time. Although all-pairs shortest paths (APSP) would cost  $\mathcal{O}(nm)$ , SP4LP never computes APSP. In all benchmark settings, we compute shortest paths only for the supervised train/validation/test links. Let  $S$  be the number of distinct endpoint nodes among these links. The total preprocessing cost is therefore

$$T_{\text{sp}} = \mathcal{O}(S \cdot m),$$

typically several orders of magnitude smaller than  $\mathcal{O}(nm)$ . Once these paths are cached, no BFS is performed during training or batched inference. In fully inductive scenarios, shortest-path trees for new nodes can be obtained via a single BFS per new node, which provides distances and parent pointers to all existing nodes at once. The per-link inference cost of SP4LP is thus

$$C = \mathcal{O}(kd + T_{\text{seq}}),$$

where  $k$  is the path length and  $d$  the embedding dimension, making it independent of the graph size  $n$ .  $T_{\text{seq}}$  denotes the cost of a sequence model (e.g., LSTM, Transformer) applied to the node embeddings along that path.

Importantly, since both the GNN embeddings and shortest paths can be precomputed,  $B$  is incurred only once. Given that  $k$  is typically small in practice,  $C$  remains low even at scale. Unlike subgraph-based methods such as SEAL, which require per-link GNN inference over extracted neighborhoods, SP4LP performs lightweight sequence modeling on compact paths, supporting efficient batched inference. Additionally, the framework is flexible:  $T_{\text{sp}}$  can be cached or approximated Sommer (2014), and  $T_{\text{seq}}$  depends on the chosen architecture and path length.

We summarize the complexities in the following table:

Method	B	C
GAE	$n\Delta d + nd^2$	$d^2$
Neo-GNN	$n\Delta d + nd^2 + n\Delta^l$	$\Delta^l + d^2$
BUDDY	$n\Delta d + nh$	$h + d^2$
SEAL	0	$\Delta^{(l+1)}d + \Delta^l d^2$
NCN	$n\Delta d + nd^2$	$\Delta d + d^2$
NCNC	$n\Delta d + nd^2$	$\Delta^2 d + \Delta d^2$
PEG	$n\Delta d + nd^2 + nD^2$	$d^2$
<b>SP4LP</b>	$n\Delta d + nd^2 + T_{\text{sp}}$	$kd + T_{\text{seq}}$

Table 9: Comparison of precomputation ( $B$ ) and per-link ( $C$ ) complexities across methods.

All methods conform to the form  $\mathcal{O}(B + C \cdot t)$ , with a one-time graph-level computation and a per-target-link cost. In the table,  $h$  is the cost of the hash function used in BUDDY,  $l$  denotes the radius of local neighborhoods in Neo-GNN and SEAL, and  $D$  denotes the number of Laplacian eigenvectors used for spectral positional encoding in PEG. GAE, NCN, and BUDDY are efficient but limited in expressiveness. NCNC enhances NCN via soft neighbor completion at a slightly higher cost. PEG incorporates spectral encoding ( $\mathcal{O}(nD^2)$ ), while SEAL is the most computationally intensive due to subgraph extraction and per-link GNN processing. SP4LP, in contrast, strikes a balance between efficiency and expressiveness by combining GNN-based embeddings with sequence modeling over shortest paths, whose cost can be mitigated via approximation techniques Sommer (2014), ensuring scalability for large graphs.

1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133