Graph Neural Networks for Link Prediction: A Theoretical Perspective

Abstract

In this paper, we theoretically characterize graph neural network's representation power for high-order node set prediction problems (where a prediction is made over a set of more than 1 node). In particular, we focus on one most important second-order task—link prediction. There are two representative classes of GNN methods for link prediction: GAE and SEAL. GAE (Graph Autoencoder) first applies a GNN to the whole graph, and then aggregates the representations of the source and target nodes as their link representation. SEAL extracts a subgraph around the source and target nodes, **labels** the nodes in the subgraph, and then uses a GNN to learn a link representation from the labeled subgraph. At first glance, both GAE and SEAL use a GNN. However, their performance gap can be very large. On the recent Open Graph Benchmark datasets, SEAL achieved 3 first places out of 4 datasets, outperforming the best GAE method by up to 195% in Hits@100. In this paper, by studying this performance gap between GAE and SEAL, we first point out a key limitation of GAE caused by directly aggregating two node representations as a link representation. To address this limitation, we propose the *labeling trick*. Labeling trick unifies several recent successes to improve GNNs' representation power, such as SEAL, Distance Encoding, and Identity-aware GNN, into a single and most basic form. We prove that with labeling trick a sufficiently expressive GNN can learn the most expressive structural representations for node sets. Our work establishes a theoretical foundation for using GNNs for high-order node set prediction.

1 Introduction

Graph neural networks (GNNs) [1–10] have achieved great success in recent years. While graph neural networks have been well studied for single-node tasks (such as node classification) and whole-graph tasks (such as graph classification), using GNNs to predict over a set of several nodes is less studied and less understood. Among such high-order node set prediction problems, link prediction (predicting the link existence between a set of two nodes) is perhaps the most important one. Link prediction is to predict missing or future links between nodes in a network. It has wide applications, such as friend recommendation in social networks [11], movie recommendation in Netflix [12], protein interaction prediction [13], and drug response prediction [14]. In this paper, we use link prediction as a medium to study GNN's high-order node set representation ability. Note that although our examples and experiments are all around link prediction, our main technique and theorem apply to **all** high-order node set representation problems.

Traditional link prediction methods can be categorized into three classes: heuristic methods, embedding methods, and feature-based methods. Heuristic methods compute some heuristic graph structure features as the likelihood of links [15]. Examples include common neighbors, preferential attachment [16], and Katz index [17]. Embedding methods, including matrix factorization (MF)

and Node2vec [18], learn free-parameter node embeddings in a transductive manner, thus do not generalize to unseen nodes and networks. Feature-based methods leverage explicit node features to predict links. By learning from graph structure and node features in a unified way, GNNs have shown superior link prediction performance than traditional methods in recent years [19–23].

There are two main types of GNN-based link prediction methods. One is Graph Autoencoder (**GAE**) [19], where a GNN is first applied to the entire network to compute a representation for each node. Then the representations of the source and target nodes are aggregated to predict the target link. Its variational version is called VGAE. The second type is **SEAL** [20, 23], where a local enclosing subgraph is extracted around each target link. Then the nodes in each enclosing subgraph are *labeled differently* according to their distances to the source and target nodes. Finally a GNN is applied to each enclosing subgraph to learn a link representation for link prediction. At first glance, GAE and SEAL look very similar—they both use a GNN to learn neighborhood structures and features around target links for link prediction. However, as we will see, SEAL has **fundamentally better** link representation learning ability than GAE. The key lies in SEAL's **node labeling** step.

Figure 1: In this graph, nodes v_2 and v_3 are isomorphic; links (v_1, v_2) and (v_4, v_3) are isomorphic; link (v_1, v_2) and link (v_1, v_3) are **not** isomorphic. However, if we aggregate two node representations learned by a GNN as the link representation, we will give (v_1, v_2) and (v_1, v_3) the same prediction.

We first give a simple example to show when GAE fails. In Figure 1, v_2 and v_3 have symmetric positions in the graph—from their perspective points of view, they have exactly the **same h-hop neighborhood**, no matter how large h is. Thus, without node features, a GAE will learn the **same** representation for v_2 and v_3. However, when we want to predict which one of v_2 and v_3 is more likely to form a link with v_1, GAE will aggregate the representations of v_1 and v_2 as the link representation of (v_1, v_2), and aggregate the representations of v_1 and v_3 to represent (v_1, v_3), thus giving (v_1, v_2) and (v_1, v_3) the same representation and prediction. The failure to distinguish links with apparently different structural roles reflects one key limitation of GAE methods—they cannot capture the association between nodes: v_1 is much closer to v_2 than v_3 in the graph, and shares more common neighbors with v_2. By learning v_1 and v_2's representations independently, GAE fails to detect these.

One way to alleviate the above failure is to use one-hot encoding of node indices or random initialization as the input node features [24, 25]. With such node-discriminating features, v_2 and v_3 will have different node representations, thus (v_1, v_2) and (v_1, v_3) will also have different link representations after aggregation. However, using node-discriminating features loses GNN's inductive learning ability to map nodes and links with identical neighborhoods (such as nodes v_2 and v_3, and links (v_1, v_2) and (v_4, v_3)) to the same representation, which results in some loss of the generalization ability. Is there a way to make GNNs discriminate the cases like (v_1, v_2) and (v_1, v_3), while still maintaining their inductive learning ability, so that links (v_1, v_2) and (v_4, v_3) are mapped to the same representation (so are nodes v_2 and v_3)?

In this paper, we aim to address the above problem from a *structural link representation* learning point of view. Based on Theorem 1 of [26], the node set prediction problem on graphs ultimately requires finding a (most expressive) *structural representation* of node sets which gives two node sets the same representation if and only if they are *isomorphic* in the graph. For example, in Figure 1, link (v_1, v_2) and link (v_4, v_3) are isomorphic (i.e., symmetric in positions). A (most expressive) structural link representation should give any two isomorphic links the same representation while discriminating all non-isomorphic links (such as (v_1, v_2) and (v_1, v_3)). According to our discussion above, GAE fails to learn such a structural link representation. In this regard, we propose the *labeling trick* (Section 4.1), which gives a particular label to each node as its additional feature according to its structural relationship with the target node set to predict. We prove that combined with a labeling trick, a sufficiently expressive GNN can **learn structural link representations** by aggregating its pairwise node representations learned from the labeled graph. We further show that the techniques used in SEAL [20], Distance Encoding [23], and ID-GNN [27] are all specific implementations of the labeling trick, which explains their better link prediction performance than plain GAE.

Finally, we notice that isomorphic nodes/links might not be prevalent in real-world networks. Therefore, a more practical definition called *local isomorphism* is proposed. We argue that GNNs for high-order node set prediction should focus on local isomorphism instead of (global) isomorphism,

and extend our theory to local isomorphism. Experiments on four Open Graph Benchmark (OGB) link prediction datasets [28] demonstrate the effectiveness of the proposed labeling trick.

2 Preliminaries

In this section, we introduce some important concepts that will be used in the discussion of the paper, including *permutation*, *set isomorphism* and *structural representation*.

We consider an undirected graph $\mathcal{G} = (V, E, \mathbf{A})$, where $V = \{1, 2, \ldots, n\}$ is the set of n vertices, $E \subseteq V \times V$ is the set of edges, and $\mathbf{A} \in \mathbb{R}^{n \times n \times k}$ is a 3-dimensional tensor containing node and edge features. The diagonal components $\mathbf{A}_{i,i,:}$ denote features of node i, and the off-diagonal components $\mathbf{A}_{i,j,:}$ denote features of edge (i, j). We further use $A \in \{0, 1\}^{n \times n}$ to denote the adjacency matrix of $\mathcal{G}$ with $A_{i,j} = 1$ iff $(i, j) \in E$. If there are no node/edge features, we let $\mathbf{A} = A$. Otherwise, A can be regarded as the first slice of $\mathbf{A}$, i.e., $A = \mathbf{A}_{:,:,1}$.

A *permutation* π is a bijective mapping from $\{1, 2, \ldots, n\}$ to $\{1, 2, \ldots, n\}$. Depending on the context, $\pi(i)$ can mean assigning a new index to node $i \in V$, or mapping node i to node $\pi(i)$ of another graph. All $n!$ possible π's constitute the permutation group Π_n. For joint prediction tasks over a set of nodes, we use S to denote the **target node set**. For example, $S = \{i, j\}$ if we want to predict the link between i, j. We define $\pi(S) = \{\pi(i) | i \in S\}$. We further define the permutation of $\mathbf{A}$ as $\pi(\mathbf{A})$, where $\pi(\mathbf{A})_{\pi(i),\pi(j),:} = \mathbf{A}_{i,j,:}$.

Next, we define *set isomorphism*, which generalizes graph isomorphism to arbitrary node sets.

Definition 1. *(Set isomorphism) Given two n-node graphs $\mathcal{G} = (V, E, \mathbf{A})$, $\mathcal{G}' = (V', E', \mathbf{A}')$, and two node sets $S \subseteq V$, $S' \subseteq V'$, we say $(S, \mathbf{A})$ and $(S', \mathbf{A}')$ are isomorphic (denoted by $(S, \mathbf{A}) \simeq (S', \mathbf{A}')$) if $\exists \pi \in \Pi_n$ such that $S = \pi(S')$ and $\mathbf{A} = \pi(\mathbf{A}')$.*

When $(V, \mathbf{A}) \simeq (V', \mathbf{A}')$, we say two graphs $\mathcal{G}$ and $\mathcal{G}'$ are *isomorphic* (abbreviated as $\mathbf{A} \simeq \mathbf{A}'$ because $V = \pi(V')$ for any π). Note that set isomorphism is **more strict** than graph isomorphism, because it not only requires graph isomorphism, but also requires the permutation maps a specific node set S to another node set S'. In practice, when $S \neq V$, we are often more concerned with the case of $\mathbf{A} = \mathbf{A}'$, where we are to find isomorphic node sets **in the same graph** (automorphism). For example, when $S = \{i\}, S' = \{j\}$ and $(i, \mathbf{A}) \simeq (j, \mathbf{A})$, we say nodes i and j are isomorphic in graph $\mathbf{A}$ (or they have symmetric positions/same structural role within the graph). An example is v_2 and v_3 in Figure 1.

We say a function f defined over the space of $(S, \mathbf{A})$ is *permutation invariant* (or *invariant* for abbreviation) if $\forall \pi \in \Pi_n, f(S, \mathbf{A}) = f(\pi(S), \pi(\mathbf{A}))$. Similarly, f is *permutation equivariant* if $\forall \pi \in \Pi_n, \pi(f(S, \mathbf{A})) = f(\pi(S), \pi(\mathbf{A}))$.

Now we define (most expressive) structural representation of a node set, following [26, 23]. Basically, it assigns a unique representation to each equivalence class of isomorphic node sets.

Definition 2. *Given an invariant function $\Gamma(\cdot)$, $\Gamma(S, \mathbf{A})$ is a **most expressive structural representation** for $(S, \mathbf{A})$ if $\forall S, \mathbf{A}, S', \mathbf{A}', \Gamma(S, \mathbf{A}) = \Gamma(S', \mathbf{A}') \Leftrightarrow (S, \mathbf{A}) \simeq (S', \mathbf{A}')$.*

For simplicity, we will briefly use *structural representation* to denote most expressive structural representation in the rest of the paper. We will omit $\mathbf{A}$ if it is clear from context. We call $\Gamma(i, \mathbf{A})$ a *structural node representation* for i, and call $\Gamma(\{i, j\}, \mathbf{A})$ a *structural link representation* for (i, j).

Definition 2 requires the structural representations of two node sets to be the same if and only if they are isomorphic. That is, isomorphic node sets always have the **same** structural representation, while non-isomorphic node sets always have **different** structural representations. This is in contrast to *positional node embeddings* such as DeepWalk [29] and matrix factorization [30], where two isomorphic nodes can have different node embeddings [31]. GAE using node-discriminating features also learns positional node embeddings.

Why do we need structural representations? Formally speaking, Srinivasan and Ribeiro [26] prove that any joint prediction task over node sets only requires a structural representation of node sets. They show that for any graph prediction task, positional node embeddings carry no more information beyond that of structural representations. Intuitively speaking, it is because two isomorphic nodes in a network are perfectly symmetric and interchangeable with each other, and should be indistinguishable from any perspective. Learning a structural node representation can guarantee that isomorphic nodes

143 are always classified into the same class. Similarly, learning a structural link representation guarantees
144 isomorphic links, such as (v_1, v_2) and (v_4, v_3) in Figure 1, are always predicted the same, which is
145 not guaranteed by positional node embeddings.

3 GAE cannot learn structural link representations

147 In this section, we review GAE, and show one key limitation caused by directly aggregating node
148 representations, which prevents GAE from learning structural link representations.

3.1 GAE for link prediction

150 Given a graph $\mathbf{A}$, GAE methods [19] first use a GNN to compute a node representation z_i for each
151 node i, and then use $f(z_i, z_j)$ to predict link (i, j):

$$\hat{\mathbf{A}}_{i,j} = f(z_i, z_j), \text{ where } z_i = \text{GNN}(i, \mathbf{A}), z_j = \text{GNN}(j, \mathbf{A}).$$

152 Here $\hat{\mathbf{A}}_{i,j}$ is the predicted score for link (i, j). The model is trained to maximize the likelihood
153 of reconstructing the true adjacency matrix. The original GAE uses a two-layer GCN [5] as the
154 GNN, and let $f(z_i, z_j) := \sigma(z_i^\top z_j)$. In principle, we can replace GCN with any GNN, and replace
155 $\sigma(z_i^\top z_j)$ with an MLP over any aggregation function over $\{z_i, z_j\}$. Besides inner product, other
156 aggregation choices include bilinear product, concatenation, mean and Hadamard product. In the
157 following, we will use GAE to denote a general class of GNN-based link prediction methods.

158 GAE uses a GNN to learn node representations and then aggregate pairwise node representations
159 as link representations. Two natural questions to ask are: 1) Is the node representation learned by
160 the GNN a *structural node representation*? 2) Is the link representation aggregated from two node
161 representations a *structural link representation*? We answer them respectively in the following.

3.2 GNN and structural node representation

163 Practical GNNs [32] usually simulate the 1-dimensional Weisfeiler-Lehman (1-WL) test [33] to
164 iteratively update each node's representation by aggregating its neighbors' representations. We use
165 *1-WL-GNN* to denote a GNN with 1-WL discriminating power, such as GIN [34].

166 A 1-WL-GNN ensures that isomorphic nodes always have the same representation. But the opposite
167 direction is not guaranteed. For example, a 1-WL-GNN always gives the same representation to
168 all nodes in an r-regular graph. Despite this, 1-WL is known to **discriminate almost all non-**
169 **isomorphic nodes** [35]. This indicates that a 1-WL-GNN can always give the same representation to
170 isomorphic nodes, and can give different representations to **almost all** non-isomorphic nodes.

171 We also define a *node-most-expressive GNN*, which gives different representations to **all** non-
172 isomorphic nodes.

173 **Definition 3.** A GNN *is **node-most-expressive** if* $\forall i, \mathbf{A}, j, \mathbf{A}', \text{ GNN}(i, \mathbf{A}) = \text{GNN}(j, \mathbf{A}') \Leftrightarrow$
174 $(i, \mathbf{A}) \simeq (j, \mathbf{A}')$.

175 That is, node-most-expressive GNN learns *structural node representations*[1]. We define such a GNN
176 because we want to answer: whether GAE, even equipped with a node-most-expressive GNN (so that
177 GNN's node representation power is not a bottleneck), can learn structural link representations.

3.3 GAE cannot learn structural link representations

179 Suppose GAE is equipped with a node-most-expressive GNN which outputs structural node repre-
180 sentations. Then the question becomes: does the aggregation of structural node representations of
181 i and j result in a structural *link* representation of (i, j)? The answer is no, as shown in previous
182 works [26, 37]. We have also illustrated it in the introduction: In Figure 1, we have two isomorphic
183 nodes v_2 and v_3, thus v_2 and v_3 will have the same structural node representation. By aggregating
184 structural node representations, GAE will give (v_1, v_2) and (v_1, v_3) the same link representation and
185 prediction. However, (v_1, v_2) and (v_1, v_3) are not isomorphic in the graph. This indicates:

---

[1]Although a polynomial-time implementation is not known for node-most-expressive GNNs, many practical
softwares can discriminate all non-isomorphic nodes quite efficiently [36], which provides a promising direction.

Proposition 1. *Even with a node-most-expressive GNN, GAE **cannot** learn structural link represen-tations.*

The root cause of this problem is that GAE computes two node representations independently of each other, without considering their relative positions and **associations**. Because of this, GAE even cannot learn to count the **common neighbors** between two nodes (1 for (v_1, v_2) and 0 for (v_1, v_3)), one most important link prediction heuristic [15]. This greatly hurts GAE's link prediction ability.

4 How to learn structural representations

In this section, we first define the general form of the *labeling trick*, and then use a specific example, zero-one labeling trick, to intuitively explain why labeling trick helps a GNN learn better link representations. Next, we present our main theorem which states that labeling trick enables a node-most-expressive GNN to learn structural link representations, thus formally characterizing labeling trick's power for high-order node set representation learning. Finally, we review SEAL and show it exactly uses one labeling trick. We also show two recent successful techniques, Distance Encoding [23] and ID-GNN [27], are also specific labeling tricks.

4.1 Labeling trick

Definition 4. *(Labeling trick) Given* $(S, \mathbf{A})$*, we stack a labeling tensor* $\mathbf{L}^{(S)} \in \mathbb{R}^{n \times n \times d}$ *in the third dimension of* $\mathbf{A}$ *to get a new* $\mathbf{A}^{(S)} \in \mathbb{R}^{n \times n \times (k+d)}$*, where* $\mathbf{L}$ *satisfies:* $\forall S, \mathbf{A}, S', \mathbf{A}', \pi \in \Pi_n$,

(1) $\mathbf{L}^{(S)} = \pi(\mathbf{L}^{(S')}) \Rightarrow S = \pi(S')$*, and*

(2) $S = \pi(S'), \mathbf{A} = \pi(\mathbf{A}') \Rightarrow \mathbf{L}^{(S)} = \pi(\mathbf{L}^{(S')})$*.*

To explain a bit, labeling trick assigns a label vector to each node/edge in graph $\mathbf{A}$, which constitutes the labeling tensor $\mathbf{L}^{(S)}$. By concatenating $\mathbf{A}$ and $\mathbf{L}^{(S)}$, we get the new labeled graph $\mathbf{A}^{(S)}$. By definition we can assign labels to both nodes and edges. However, in this paper, we **only consider node labels** for simplicity, i.e., we let the off-diagonal components $\mathbf{L}_{i,j,:}^{(S)}$ be all zero.

The labeling tensor $\mathbf{L}^{(S)}$ should satisfy two conditions in Definition 4. Condition (1) requires the target nodes S to have *distinct labels* from those of the rest nodes, so that S is distinguishable from others. This is because if a permutation π preserving node labels (i.e., $\mathbf{L}^{(S)} = \pi(\mathbf{L}^{(S')})$) exists between nodes of $\mathbf{A}$ and $\mathbf{A}'$, then S and S' must have distinct labels to guarantee S' is mapped to S by π (i.e., $S = \pi(S')$). Condition (2) requires the labeling function to be *permutation equivariant*, i.e., when (S, A) and (S', A') are isomorphic under π (i.e., $S = \pi(S'), \mathbf{A} = \pi(\mathbf{A}')$), the corresponding nodes $i \in S, j \in S', i = \pi(j)$ must always have the same label (i.e., $\mathbf{L}^{(S)} = \pi(\mathbf{L}^{(S')})$). In other words, the labeling function should be *consistent* across different S.

Now we introduce a simplest labeling trick satisfying the two conditions in Definition 4, and use it to illustrate how labeling trick helps GNNs learn better link representations.

Definition 5. *(Zero-one labeling trick) Given a graph* $\mathbf{A}$ *and a set of nodes* S *to predict, we give it a diagonal labeling matrix* $\mathbf{L}^{(S)} \in \mathbb{R}^{n \times n \times 1}$ *such that* $\mathbf{L}_{i,i,1}^{(S)} = 1$ *if* $i \in S$ *otherwise 0.*

In other words, the zero-one labeling trick assigns label 1 to nodes in S, and label 0 to all nodes not in S. It is a valid labeling trick, because firstly, nodes in S get *distinct labels*, and secondly, the labeling function is *consistent* by always giving nodes in S a label 1. These node labels will be used as additional node features, which are fed to a GNN together with original node features if available.

Let's return to the example in Figure 1 to see how the zero-one labeling trick helps GNNs learn better link representations. This time, when we want to predict link (v_1, v_2), we will mark v_1, v_2 with a different label from those of the rest nodes, as shown by the different color in Figure 2 left. With the source and target nodes labeled, when the GNN is computing v_2's representation, it is also "aware" of the source node v_1, instead of the previous agnostic way that treats v_1 the same as other nodes. And when we predict link (v_1, v_3), we will again mark v_1, v_3 with a different label, as shown in Figure 2 right. This way, v_2 and v_3's node representations are no longer the same in the two differently labeled graphs (due to the presence of the labeled v_1), and we are able to predict (v_1, v_2)

Figure 2: When we predict (v_1, v_2), we will label these two nodes differently from the rest, so that a GNN is aware of the target link when learning v_1 and v_2's representations. Similarly, when predicting (v_1, v_3), nodes v_1 and v_3 will be labeled differently. This way, the representation of v_2 in the left graph will be different from that of v_3 in the right graph, enabling GNNs to distinguish the non-isomorphic links (v_1, v_2) and (v_1, v_3).

and (v_1, v_3) differently. The key difference from GAE is that the node representations are no longer learned independently, but are *conditioned* on each other in order to capture two nodes' association.

At the same time, isomorphic links, such as (v_1, v_2) and (v_4, v_3), will still have the same representation, since the zero-one labeled graph for (v_1, v_2) is still symmetric to the zero-one labeled graph for (v_4, v_3). This brings an exclusive advantage over GAE using node-discriminating features.

Now we introduce our main theorem showing that with a valid labeling trick, a node-most-expressive GNN can *learn structural link representations*.

Theorem 1. *Given a node-most-expressive* GNN *and an injective set aggregation function* AGG, *for any $S, \mathbf{A}, S', \mathbf{A}'$,* GNN$(S, \mathbf{A}^{(S)}) = $ GNN$(S', \mathbf{A}'^{(S')}) \Leftrightarrow (S, \mathbf{A}) \simeq (S', \mathbf{A}')$, *where* GNN$(S, \mathbf{A}^{(S)}) := $ AGG$(\{$GNN$(i, \mathbf{A}^{(S)})|i \in S\})$.

Theorem 1 implies that AGG$(\{$GNN$(i, \mathbf{A}^{(S)})|i \in S\})$ is a structural representation for $(S, \mathbf{A})$. The significance of Theorem 1 is that it closes the gap between GNN's node representation nature and link prediction's link representation requirement. Remember that directly aggregating structural node representations learned from the original graph $\mathbf{A}$ does not lead to structural link representations. Theorem 1 shows that aggregating over the structural node representations learned from the labeled graph $\mathbf{A}^{(S)}$, somewhat surprisingly, results in a structural representation for S.

Theorem 1 also characterizes GNNs' node set representation learning ability, i.e., with labeling trick GNNs can indeed learn (most expressive) structural representations for node sets. This answers the open question raised in [26] questioning GNNs' link prediction ability: *are structural node representations in general–and GNNs in particular–fundamentally incapable of performing link (dyadic) and multi-ary (polyadic) prediction tasks?* With Theorem 1, we argue the answer is no. Although GNNs alone can at most learn structural node representations, GNNs combined with a labeling trick can learn structural link representations too.

Theorem 1 assumes a node-most-expressive GNN. To augment Theorem 1, we give the following theorem, which demonstrates labeling trick's power for 1-WL-GNNs.

Theorem 2. *In any non-attributed graphs with n nodes, if the degree of each node in the graph is between 1 and $\mathcal{O}(\log^{\frac{1-\epsilon}{2h}} n)$ for any constant $\epsilon > 0$, then there exists $\omega(n^{2\epsilon})$ many pairs of non-isomorphic links $(u, w), (v, w)$ such that an h-layer 1-WL-GNN gives u, v the same representation, while with labeling trick the 1-WL-GNN gives u, v different representations.*

Theorem 2 shows that labeling trick helps a 1-WL-GNN discriminate a large number of non-isomorphic links which are previously not distinguishable without a labeling trick.

4.2 SEAL uses a valid labeling trick

SEAL [20] is a state-of-the-art link prediction method based on GNNs. It first extracts an *enclosing subgraph* (h-hop subgraph) around each target link to predict.

Definition 6. *(Enclosing subgraph) Given $(S, \mathbf{A})$, the h-hop enclosing subgraph $\mathbf{A}_{(S,h)}$ of S is the subgraph induced from $\mathbf{A}$ by $\cup_{j \in S}\{i \mid d(i, j) \leq h\}$, where $d(i, j)$ is the shortest path distance between nodes i and j.*

Then, SEAL applies Double Radius Node Labeling (DRNL) to give an integer label to each node in the enclosing subgraph. DRNL assigns different labels to nodes with **different distances** to both the source and target nodes. It works as follows: The source and target nodes are always labeled 1. Nodes

farther away from the source and target nodes get larger labels (starting from 2). For example, nodes with distances 1 and 1 to the source and target nodes will get label 2, and nodes with distances 1 and 2 to the source and target nodes will get label 3. So on and so forth. Finally the labeled enclosing subgraph is fed to a GNN to learn the link representation and output the probability of link existence.

Theorem 3. *DRNL is a valid labeling trick.*

Theorem 3 is easily proved by noticing: across different subgraphs, 1) nodes with label 1 are always the source and target nodes, and 2) nodes with the same distances to the source and target nodes always have the same label. Thus, SEAL exactly uses a specific labeling trick to enhance its power.

SEAL only uses a subgraph $\mathbf{A}_{(S,h)}$ within h hops from the source and target nodes instead of using the whole graph. This is not a constraint but rather a practical consideration (just like GAE typically uses less than 3 message passing layers in practice), and its benefits will be discussed in detail in Section 5. When $h \to \infty$, the subgraph becomes the entire graph, and SEAL is able to *learn structural link representations* from the labeled (entire) graph.

Proposition 2. *When $h \to \infty$, SEAL can learn structural link representations with a node-most-expressive GNN.*

Remark 1 (Distance Encoding and DRNL) Recently, SEAL's distance-based node labeling scheme is generalized to *Distance Encoding*, or DE, in [23]. Basically, DRNL is equivalent to DE-2 using shortest path distance. Instead of encoding two distances into one integer label, DE directly concatenate the two distances as a size-2 label vector. For example, nodes with distances 1 and 2 to the source and target nodes will be directly labeled by a vector $[1, 2]$. **DE is also a valid labeling trick**, as it can also distinguish the target set S and is permutation equivariant. Nevertheless, there are some subtle differences in their implementations, which are discussed in Appendix D.

Remark 2 (ID-GNN and zero-one labeling trick) Concurrent to our work, You et al. [27] propose Identity-aware GNN (ID-GNN), which assigns a unique color to the "identity" nodes (the target node set to predict) and then performs message passing for them with a different set of parameters. If the unique colors are 1 and 0, ID-GNN's coloring scheme is equivalent to the zero-one labeling trick.

In summary, labeling trick unifies recent successful practices to improve the power of GNNs based on node labeling, including SEAL, DE, and ID-GNN, into a single framework. Moreover, labeling trick gives the most *basic* and *general* form for such techniques to reach theoretically the highest node set representation learning power. We discuss its computational complexity in Appendix F.

5 Local isomorphism: a more practical view of isomorphism

In the previous sections, we establish our theory based on assigning node sets the same representation if and only if they are *isomorphic* to each other in the graph. However, exact isomorphism is not very common. For example, Babai and Kucera [35] prove that at least $(n - \log n)$ nodes in almost all n-node graphs are *non-isomorphic* to each other. In practice, 1-WL-GNN also takes up to $\mathcal{O}(n)$ message passing layers to reach its maximum power for discriminating non-isomorphic nodes, making it very hard to really target on finding exactly isomorphic nodes/links.

Lemma 1. *Given a graph with n nodes, a 1-WL-GNN takes up to $\mathcal{O}(n)$ message passing layers to discriminate all the nodes that 1-WL can discriminate.*

In this regard, we propose a more practical concept, called *local isomorphism*.

Definition 7. *(Local h-isomorphism)* $\forall S, \mathbf{A}, S', \mathbf{A}'$, *we say* $(S, \mathbf{A})$ *and* $(S', \mathbf{A}')$ *are locally h-isomorphic to each other if* $(S, \mathbf{A}_{(S,h)}) \simeq (S', \mathbf{A}'_{(S',h)})$.

Local h-isomorphism only requires the h-hop enclosing subgraphs around S and S' are isomorphic, instead of the entire graphs. We argue that this is a more useful definition than isomorphism, because: 1) Exact isomorphism is rare in real-world graphs. 2) Algorithms targeting on exact isomorphism are more likely to overfit. Only assigning the same representations to exactly isomorphic nodes/links may fail to identify a large amount of nodes/links that are not isomorphic but have very similar neighborhoods. Instead, nodes/links *locally isomorphic* to each other may better indicate that they should have the same representation. With local h-isomorphism, *all our previous conclusions based on standard isomorphism still apply*. For example, GAE (without node-discriminating features) still

cannot discriminate locally h-non-isomorphic links. And a node-most-expressive GNN with labeling trick can learn the same representation for two links if and only if they are locally h-isomorphic:

Corollary 1. *Given a node-most-expressive* GNN *and an injective set aggregation function* AGG, *then for any* $S, \mathbf{A}, S', \mathbf{A}', h,$ $\text{GNN}(S, \mathbf{A}_{(S,h)}^{(S)}) = \text{GNN}(S', \mathbf{A}_{(S',h)}'^{(S')}) \Leftrightarrow (S, \mathbf{A}_{(S,h)}) \simeq (S', \mathbf{A}_{(S',h)}')$.

Corollary 1 demonstrates labeling trick's power for local isomorphism discriminating, too. To switch to local h-isomorphism, all we need to do is to extract the h-hop enclosing subgraph around a link, and apply the labeling trick and GNN only to the extracted subgraph. This is exactly what SEAL did.

6 Related work

There is emerging interest in studying the representation power of graph neural networks recently. Xu et al. [34] and Morris et al. [38] first show that the discriminating power of GNNs performing neighbor aggregation is bounded by the 1-WL test. Many works have since been proposed to increase the power of GNNs by simulating higher-order WL tests [38–40]. However, most previous works focus on improving GNN's whole-graph representation power. Little work has been done to analyze GNN's node/link representation power. Srinivasan and Ribeiro [26] first formally studied the difference between structural representations of nodes and links. Although showing that structural node representations of GNNs cannot perform link prediction, their way to learn structural link representations is to give up GNNs and instead use Monte Carlo samples of node embeddings learned by network embedding methods. In this paper, we show that GNNs combined with a simple labeling trick can as well learn structural link representations, which reassures using GNNs for link prediction.

Many works have implicitly assumed that if a model can learn node representations well, then combining the pairwise node representations can also lead to good link representations [18, 19, 41]. However, we argue in this paper that simply aggregating node representations fails to discriminate a lot of non-isomorphic links, and with labeling trick the aggregation of structural node representations lead to structural link representations. Li et al. [23] proposed distance encoding (DE), which is shown to be a particular labeling trick. They proved that DE can improve 1-WL-GNNs' discriminating power, enabling them to differentiate almost all $(S, \mathbf{A})$ tuples sampled from r-regular graphs. Our paper's theoretical contribution is orthogonal to that of DE: 1) Our theory focuses on the gap between a GNN's node and node set representation power, instead of improving one particular GNN's discriminating power. We show even a GNN has maximum node representation power, it still fails to learn structural link representations unless combined with a labeling trick. 2) Our theory is not restricted to r-regular graphs, but applies to **any** graphs. 3) The proposed labeling trick is more general than DE. We show that labeling trick is not necessarily distance based—a valid labeling trick need only be permutation equivariant and S-discriminating. More discussion on DE's theory is given in Appendix E.

You et al. [21] also noticed that the structural node representations of GNNs cannot encode the associations and distances between nodes. To learn position-aware node embeddings, they propose P-GNN, which randomly chooses some anchor nodes and aggregates messages only from the anchor nodes. In P-GNN, nodes of similar distances to the anchor nodes, instead of nodes with similar neighborhoods, will have similar embeddings. Thus, P-GNN **cannot** learn structural node/link representations. P-GNN also cannot scale to large datasets. You et al. [27] recently proposed ID-GNN. As discussed in Remark 2, ID-GNN's node coloring is also a specific labeling trick.

7 Experiments

In this section, we empirically validate the effectiveness of labeling trick on Open Graph Benchmark (OGB) [28] datasets. We use all the four link prediction datasets in OGB: `ogbl-ppa`, `ogbl-collab`, `ogbl-ddi`, and `ogbl-citation2`. These datasets are open-sourced, large-scale (up to 2.9M nodes and 30.6M edges), adopt realistic train/validation/test splits, and have standard evaluation procedures, thus providing an ideal place to benchmark an algorithm's realistic link prediction power. The evaluation metrics include Hits@K and MRR. Hits@K counts the ratio of positive edges ranked at the K-th place or above against all the negative edges. MRR (Mean Reciprocal Rank) computes the reciprocal rank of the true target node against 1,000 negative candidates, averaged over all the true source nodes. Both metrics are higher the better. We include more details and statistics of these datasets in Appendix G. Note that knowledge graph completion datasets are not used because

Table 1: Results for `ogbl-ppa`, `ogbl-collab`, `ogbl-ddi` and `ogbl-citation2`.

Category	Method	ogbl-ppa Hits@100 (%)		ogbl-collab Hits@50 (%)		ogbl-ddi Hits@20 (%)		ogbl-citation2 MRR (%)	
		Validation	Test	Validation	Test	Validation	Test	Validation	Test
Non-GNN	CN	28.23±0.00	27.6±0.00	60.36±0.00	61.37±0.00	9.47±0.00	17.73±0.00	51.19±0.00	51.47±0.00
	AA	32.68±0.00	32.45±0.00	63.49±0.00	64.17±0.00	9.66±0.00	18.61±0.00	51.67±0.00	51.89±0.00
	MLP	0.46±0.00	0.46±0.00	24.02±1.45	19.27±1.29	–	–	29.03±0.17	29.06±0.16
	Node2vec	22.53±0.88	22.26±0.88	57.03±0.52	48.88±0.54	32.92±1.21	23.26±2.09	61.24±0.11	61.41±0.11
	MF	32.28±4.28	32.29±0.94	48.96±0.29	38.86±0.29	33.70±2.64	13.68±4.75	51.81±4.36	51.86±4.43
Plain GAE	GraphSAGE	17.24±2.64	16.55±2.40	56.88±0.77	54.63±1.12	62.62±0.37	53.90±4.74	82.63±0.23	82.60±0.36
	GCN	18.45±1.40	18.67±1.32	52.63±1.15	47.14±1.45	55.50±2.08	37.07±5.07	84.79±0.23	84.74±0.21
	GCN+LRGA	25.75±2.82	26.12±2.35	60.88±0.59	52.21±0.72	**66.75**±0.58	**62.30**±9.12	66.48±1.61	66.49±1.59
Labeling Trick	GCN+DE	36.31±3.59	36.48±3.78	64.13±0.16	**64.44**±0.29	29.85±2.25	26.63±6.82	60.17±0.63	60.30±0.61
	GCN+DRNL	46.43±3.03	45.24±3.95	**64.51**±0.42	64.40±0.45	29.47±1.54	22.81±4.93	81.07±0.30	81.27±0.31
	SEAL	**51.25**±2.52	**48.80**±3.16	63.89±0.49	63.64±0.71	28.49±2.69	30.56±3.86	**87.57**±0.31	**87.67**±0.32

knowledge graph link prediction is a different problem from the homogeneous graph link prediction problem studied in this paper. Their baselines and datasets are *not* interchangeable with each other.

Baselines. We use the following baselines for comparison. We use 5 non-GNN methods: CN (common neighbor), AA (Adamic-Adar), MLP, MF (matrix factorization) and Node2vec. Among then, CN and AA are two simple link prediction heuristics based on counting common neighbors, which are used for sanity checking. We use 3 plain GAE baselines: GraphSAGE [41], GCN [19], and GCN+LRGA [42]. These methods use the Hadamard product of pairwise node representations output by a GNN as link representations, without using a labeling trick. Finally, we compare 3 GNN methods using labeling tricks: GCN+DE [23], GCN+DRNL, and SEAL [20]. GCN+DE/GCN+DRNL enhance GCN with the DE/DRNL labeling trick. SEAL uses a GCN and the DRNL labeling trick, with an additional subgraph-level readout SortPooling [9]. More details are in Appendix H. Moreover, we test the zero-one labeling trick in our ablation experiments. Results can be found in Appendix I.

Results and discussion. We present the main results in Table 1. Firstly, we can see that GAE methods without labeling trick do not always outperform non-GNN methods. For example, on `ogbl-ppa` and `ogbl-collab`, simple heuristics CN and AA outperform plain GAE methods by large margins. This suggests that GAE methods cannot learn simple heuristics like common neighbors which require capturing the associations between nodes, matching our analysis in Section 3.3. In contrast, when GNNs are enhanced by labeling trick, they are able to beat heuristics. With labeling trick, GNN methods achieve new state-of-the-art performance on 3 out of 4 datasets. In particular, we observe that SEAL outperforms GAE and positional embedding methods, sometimes by surprisingly large margins. For example, in the challenging `ogbl-ppa` graph, SEAL achieves an Hits@100 of 48.80, which is **87%-195% higher** than GAE methods without using labeling trick. On `ogbl-ppa`, `ogbl-collab` and `ogbl-citation2`, labeling trick methods achieve new state-of-the-art results.

Despite obtaining the best results on three datasets, we observe that labeling trick methods do not perform well on `ogbl-ddi`. `ogbl-ddi` is considerably denser than the other graphs. It has 4,267 nodes and 1,334,889 edges, resulting in an average node degree of 500.5. In `ogbl-ddi`, labeling trick methods fall behind GAE methods using trainable node embeddings. One possible explanation is that `ogbl-ddi` is so dense that a practical GNN with *limited expressive power* is hard to inductively learn any meaningful structural patterns. In comparison, the transductive way of learning free-parameter node embeddings makes GAEs no longer focus on learning inductive structural patterns, but on optimizing node embeddings. The added parameters also greatly increase GAEs' model capacity. An interesting future topic is to study how to improve labeling tricks' performance on dense graphs.

We present more ablation experiments in Appendix I to study the power of different labeling tricks, the effect of subgraph pooling, and the number of hops/layers, etc.

8 Conclusions

In this paper, we theoretically characterized GNNs' representation learning ability for high-order node sets, such as links. We first pointed out a key limitation of GAE caused by directly aggregating node representations of a GNN as node set representations. To address the problem, we proposed the labeling trick and theoretically proved its power. We also showed that several existing methods, including SEAL, DE and ID-GNN, all essentially used some labeling trick, explaining their strong performance. Experiments on OGB datasets demonstrated labeling trick's effectiveness.

References

[1] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.

[2] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.

[3] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pages 2224–2232, 2015.

[4] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.

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

[6] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3837–3845, 2016.

[7] Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for structured data. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 2702–2711, 2016.

[8] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

[9] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *AAAI*, pages 4438–4445, 2018.

[10] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems*, pages 4800–4810, 2018.

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

[12] James Bennett, Stan Lanning, et al. The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35. New York, 2007.

[13] Yanjun Qi, Ziv Bar-Joseph, and Judith Klein-Seetharaman. Evaluation of different biological data and computational classification methods for use in protein interaction prediction. *Proteins: Structure, Function, and Bioinformatics*, 63(3):490–500, 2006.

[14] Zachary Stanfield, Mustafa Coşkun, and Mehmet Koyutürk. Drug response prediction as a link prediction problem. *Scientific reports*, 7(1):1–13, 2017.

[15] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019–1031, 2007.

[16] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

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

[18] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864. ACM, 2016.

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

[20] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pages 5165–5175, 2018.

[21] Jiaxuan You, Rex Ying, and Jure Leskovec. Position-aware graph neural networks. *arXiv preprint arXiv:1906.04817*, 2019.

[22] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. In *Advances in neural information processing systems*, pages 4868–4879, 2019.

[23] Pan Li, Yanbang Wang, Hongwei Wang, and Jure Leskovec. Distance encoding–design provably more powerful gnns for structural representation learning. *arXiv preprint arXiv:2009.00142*, 2020.

[24] Andreas Loukas. What graph neural networks cannot learn: depth vs width. *arXiv preprint arXiv:1907.03199*, 2019.

[25] Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Random features strengthen graph neural networks. *arXiv preprint arXiv:2002.03155*, 2020.

[26] Balasubramaniam Srinivasan and Bruno Ribeiro. On the equivalence between positional node embeddings and structural graph representations. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=SJxzFySKwH`.

[27] Jiaxuan You, Jonathan Gomes-Selman, Rex Ying, and Jure Leskovec. Identity-aware graph neural networks. *arXiv preprint arXiv:2101.10320*, 2021.

[28] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.

[29] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710. ACM, 2014.

[30] Andriy Mnih and Ruslan R Salakhutdinov. Probabilistic matrix factorization. In *Advances in neural information processing systems*, pages 1257–1264, 2008.

[31] Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. struc2vec: Learning node representations from structural identity. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 385–394. ACM, 2017.

[32] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org, 2017.

[33] Boris Weisfeiler and AA Lehman. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsia*, 2(9):12–16, 1968.

[34] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.

[35] László Babai and Ludik Kucera. Canonical labelling of graphs in linear average time. In *20th Annual Symposium on Foundations of Computer Science (sfcs 1979)*, pages 39–46. IEEE, 1979.

[36] Brendan D McKay and Adolfo Piperno. Practical graph isomorphism, ii. *Journal of Symbolic Computation*, 60:94–112, 2014.

[37] Muhan Zhang and Yixin Chen. Inductive matrix completion based on graph neural networks. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=ByxxgCEYDS`.

[38] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4602–4609, 2019.

[39] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. In *Advances in Neural Information Processing Systems*, pages 2156–2167, 2019.

[40] Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with gnns. In *Advances in Neural Information Processing Systems*, pages 15894–15902, 2019.

[41] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1025–1035, 2017.

[42] Omri Puny, Heli Ben-Hamu, and Yaron Lipman. From graph low-rank global attention to 2-fwl approximation. *arXiv preprint arXiv:2006.07846*, 2020.

[43] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.