
Flashlight : Scalable Link Prediction with Effective Decoders

Anonymous Author(s)

Anonymous Affiliation

Anonymous Email

Abstract

Link prediction (LP) has been recognized as an important task in graph learning with its board practical applications. A typical application of LP is to retrieve the top scoring neighbors for a given source node, such as the friend recommendation. These services desire the high inference scalability to find the top scoring neighbors from many candidate nodes at low latencies. There are two popular decoders that the recent LP models mainly use to compute the edge scores from node embeddings: the **HadamardMLP** and **Dot Product** decoders. After theoretical and empirical analysis, we find that the HadamardMLP decoders are generally more effective for LP. However, HadamardMLP lacks the scalability for retrieving top scoring neighbors on large graphs, since to the best of our knowledge, there does not exist an algorithm to retrieve the top scoring neighbors for HadamardMLP decoders in sublinear complexity. To make HadamardMLP scalable, we propose the *Flashlight* algorithm to accelerate the top scoring neighbor retrievals for HadamardMLP: a sublinear algorithm that progressively applies approximate maximum inner product search (MIPS) techniques with adaptively adjusted query embeddings. Empirical results show that Flashlight improves the inference speed of LP by more than 100 times on the large OGBL-CITATION2 dataset without sacrificing effectiveness. Our work paves the way for large-scale LP applications with the effective HadamardMLP decoders by greatly accelerating their inference.

1 Introduction

The goal of link prediction (LP) is to predict the missing links in a graph [1]. LP is drawing increasing attention in the past decade due to its board practical applications [2]. For instance, LP can be used to recommend new friends on social media [3], and recommend attractive items to the costumers on E-commerce sites [4], so as to improve the user experience. During inference, these applications demand the LP methods to retrieve the top scoring neighbors for a source node at low latencies. This is especially challenging on large graphs because the LP methods need to search many candidate nodes to find the top scoring neighbors.

There are two main kinds of architecture followed by the recent LP models. The first uses an encoder, e.g., GCN [5], to obtain the node-level embeddings and uses a decoder, e.g., Dot Product, to get the edge scores between the paired nodes [6]. The second crops a subgraph for every edge and computes the edge score from the subgraph directly [7]. The inference speed of the second is much lower than the first, so we focus on the first kind of models to achieve fast inference on large graphs. In the last years, extensive research focuses on developing more expressive LP encoders [6, 8]. However, much less work pays attention to the essential impacts of the choice of decoders on LP's performance. In this work, we theoretically and empirically analyze two popular LP decoders: Dot Product and HadamardMLP (a MLP following the Hadamard Product), and find that the latter is generally more effective than the former.

In practical applications, we should not only consider the effectiveness of LP, but also inference efficiency. Many LP applications generally require fast retrieval of the top scoring neighbors for low-latency services [3, 9, 10]. For a Dot Product decoder, this retrieval can be approximated efficiently at the sublinear time complexity [11]. However, to the best of our knowledge, no such sublinear algorithms exist for the top scoring neighbor retrievals of the HadamardMLP decoders. This means

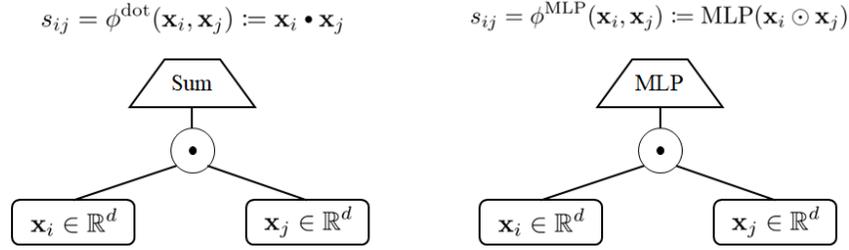


Figure 1: Two popular LP decoders: The Dot Product (left), equivalent to the element-wise summation following the Hadamard product, and the HadamardMLP decoder (right).

44 that for every source node, we have to iterate over all the nodes in the graph to compute the scores
 45 so as to find the top scoring neighbors for HadamardMLP, which is of linear complexity and cannot
 46 scale to large graphs.

47 To allow LP applications to enjoy the high effectiveness of HadamardMLP decoders while avoiding
 48 the poor inference scalability, we propose the scalable top scoring neighbor search algorithm named
 49 *Flashlight*. Our Flashlight progressively calls the well-developed approximate maximum inner
 50 product search (MIPS) techniques for a few iterations. At every iteration, we analyze the retrieved
 51 neighbors and adaptively adjust the query embedding for Flashlight to find the missed high scoring
 52 neighbors. Our Flashlight algorithm holds sublinear time complexity on finding top scoring neighbors
 53 for HadamardMLP decoders, allowing for fast and scalable inference. Empirical results show that
 54 Flashlight accelerates the inference of LP models by more than 100 times on the large OGBL-
 55 CITATION2 dataset without sacrificing the effectiveness. Overall, our work paves the way for the
 56 use of effective LP decoders in practical settings by greatly accelerating their inference.

57 2 Revisiting Link Prediction Decoders

58 In this section, we formalize the link prediction (LP) problem and the LP decoders. Typically, many
 59 LP models include an encoder that learns the node-level embeddings $\mathbf{x}_i, i \in \mathcal{V}$, where \mathcal{V} is the set of
 60 nodes, and an decoder $\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ that combines the node-level embeddings of a pair of nodes:
 61 $\mathbf{x}_i, \mathbf{x}_j$ into a single score: s_{ij} . If s_{ij} is higher, the link between nodes i and j is more likely to exist.
 62 The state-of-the-art models generally use graph neural networks as the encoders [5, 6, 8, 12, 13].
 63 From here on, we mainly focus on the decoder ϕ .

64 2.1 Dot Product Decoder

65 The most common decoder of link prediction is the Dot Product [6, 8, 10]:

$$s_{ij} = \phi^{\text{dot}}(\mathbf{x}_i, \mathbf{x}_j) := \mathbf{x}_i \bullet \mathbf{x}_j, \quad (1)$$

66 where \bullet denotes the dot product.

67 Training a link prediction model with the Dot Product decoder encourages the embeddings of the
 68 connected nodes to be close to each other. Intuitively, the score s_{ij} can be thought as a measure of the
 69 squared Eulidean distance between the node embeddings $\mathbf{x}_i, \mathbf{x}_j$, as $\|\mathbf{x}_i - \mathbf{x}_j\|^2 = \|\mathbf{x}_i\|^2 - 2\mathbf{x}_i \bullet \mathbf{x}_j +$
 70 $\|\mathbf{x}_j\|^2$, if the $\|\mathbf{x}_j\|$ is constant over the neighbors $j \in \mathcal{N}$, e.g., after normalization [14]. Because the
 71 node embeddings represent the semantic information of nodes, Dot Product assumes the homophily
 72 of graph topology, i.e., the semantically similar nodes are more likely to be connected.

73 2.2 HadamardMLP (MLP following Hadamard Product) Decoder

74 Multi layer perceptrons (MLPs) are known to be universal approximators that can approximate any
 75 continuous function on a compact set [15]. A MLP layer can be defined as a function $f : \mathbb{R}^{d_{\text{in}}} \rightarrow$
 76 $\mathbb{R}^{d_{\text{out}}}$:

$$f_{\mathbf{W}}(\mathbf{x}) = \text{ReLU}(\mathbf{W}\mathbf{x}) \quad (2)$$

77 which is parameterized by the learnable weight $\mathbf{W} \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ (the bias, if exists, can be represented
 78 by an additional column in \mathbf{W} and an additional channel in the input \mathbf{x} with the value as 1). ReLU
 79 is the activation function. In a MLP, several layers of f are stacked, e.g., a 3-layer MLP can be
 80 formalized as $f_{\mathbf{W}_3}(f_{\mathbf{W}_2}(f_{\mathbf{W}_1}(\mathbf{x})))$.

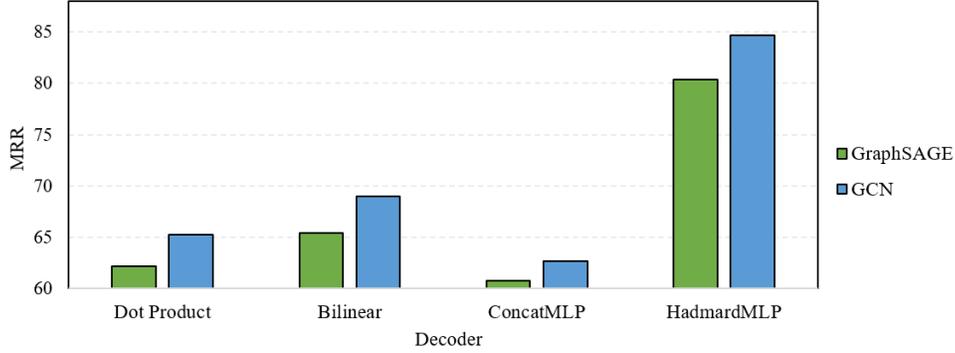


Figure 2: HadamardMLP achieves higher Mean Reciprocal Rank (MRR, higher is better) than other decoders on the OGBL-CITATION2 [16] dataset with the encoder as GraphSAGE [12] and GCN [5]. More empirical results and the detailed settings are in Sec. 6.3.

81 The state-of-the-art models widely use a MLP following the Hadamard Product between the paired
 82 nodes as the decoder (short as the HadamardMLP decoders) [6, 8, 10, 16]:

$$s_{ij} = \phi^{\text{MLP}}(\mathbf{x}_i, \mathbf{x}_j) := \text{MLP}(\mathbf{x}_i \odot \mathbf{x}_j) = \mathbf{w}_L^T (f_{\mathbf{w}_{L-1}}(\dots f_{\mathbf{w}_1}(\mathbf{x}_i \odot \mathbf{x}_j) \dots)), \quad (3)$$

83 where \odot denotes the Hadamard Product. Fig. 1 illustrates these two models the Dot Product and
 84 HadamardMLP decoders.

85 2.3 Other Link Prediction Decoders

86 In principle, every function that takes two vectors as the input and outputs a scalar can act as the
 87 decoder. For example, there are bilinear dot product decoder (short as **Bilinear decoder**) [6]:

$$s_{ij} = \mathbf{h}_i^T \mathbf{W} \mathbf{h}_j, \quad (4)$$

88 where \mathbf{W} is the learnable weight, and the MLPs following the concatenate decoder [6, 10] (short as
 89 **ConcatMLP decoder**):

$$s_{ij} = \text{MLP}(\mathbf{h}_i \parallel \mathbf{h}_j) \quad (5)$$

90 , etc. These two decoders are used much less than Dot Product and HadamardMLP in the state-of-
 91 the-art LP models possibly due to their lower effectiveness [6, 8, 10, 16].

92 2.4 HadamardMLP is Generally More Effective than Other Decoders

93 Dot Product demands the homophily of graph data to effectively infer the link between nodes. In
 94 contrast, thanks to the universal approximation capability, MLP can approximate any continuous
 95 function, and thus does not demand the homophily of graph data for effective LP. This gap in the
 96 expressiveness accounts for the performance difference of these two decoders on many datasets (see
 97 Sec. 6.3). We additionally show in Appendix. A that using a HadamardMLP is easy to learn Dot
 98 Product, which also partially accounts for the better effectiveness of the HadamardMLP decoders
 99 over the Dot Product. Existing work also finds that the effectiveness of Bilinear and ConcatMLP is
 100 generally worse than the HardmardMLP or Dot Product decoder [6, 8, 10, 16]. We confirm these
 101 findings more rigorously in the empirical results in Fig. 2 and more complete in Sec. 6.3.

102 3 Scalability of Link Prediction Decoders

103 Most academic studies focus on training runtime when discussing scalability. However, in industrial
 104 applications, the inference speed is often more important. The inference of many LP applications
 105 needs to retrieve the top scoring neighbors given a source node, e.g., recommending friends to a
 106 user for friend recommendation. Given a source node, if there are n nodes in the graph, then the
 107 inference time complexity is $\mathcal{O}(n)$ if the decoder needs to iterate over all the n nodes to compute
 108 the edge scores. For large scale applications, n is typically in the range of millions, or even larger.
 109 The empirical results show that the inference time of finding the top scoring neighbors for a source

node is longer than one second for HadamardMLP on the OGBL-CITATION2 dataset of nearly three million nodes (see Sec. 6.5).

For a Dot Product decoder, the problem of finding the top scoring neighbors can be approximated efficiently. This is a well-studied problem, known as approximate maximum inner product search (MIPS) [17, 18] (see Sec. 5.2 for a comprehensive literature review). MIPS techniques allow Dot Product’ inference to be completed in a few milliseconds, even with millions of neighbors. There exists some work that tries to extend MIPS to the ConcatMLP [19, 20]. These methods hold strict assumptions on the models’ training and are not directly applicable to the HadamardMLP. To the best of our knowledge, no such sublinear techniques exist for the top scoring neighbor retrieval with the HadamardMLP [10], which is a complex nonlinear function.

To summarize, the HadamardMLP decoder is not scalable for the real time LP services on large graphs, while the Dot Product decoder allows fast retrieval using the well established MIPS techniques.

4 Flashlight: Scalable Link Prediction with Effective Decoders

Sec. 2 has shown that the HadamardMLP decoder enjoys higher effectiveness than the Dot Product decoder, which supports the superior performance of HadamardMLP on many LP benchmarks. On the other hand, Sec. 3 has shown that the HadamardMLP is not scalable for real time LP applications on large graphs, while Dot Product supports the fast inference using the well-established MIPS techniques. In this section, we aim to devise fast inference algorithms for HadamardMLP to enable scalable LP with effective decoders.

We try to exploit the advances in the well-developed MIPS techniques to accelerate the inference of HadamardMLP. Specifically, we divide the top scoring retrievals for HadamardMLP predictors into a sequence of MIPS. Our algorithm works in a progressive manner. The query embedding in every search is adaptively adjusted to find the high scoring neighbors missed in the last search.

The challenge of retrieving the neighbors of highest scores for HadamardMLP is rooted in the unawareness of which neurons are activated, since if we know which neurons are activated, the nonlinear HadamardMLP degrades to a linear model. On the l th MLP layer, we define the mask matrix $\mathbf{M}_{\mathcal{A},l} \in \mathbb{R}^{d_l \times d_l}$ to represent the set of activated neurons \mathcal{A} as

$$M_{ij} = \begin{cases} 1, & \text{if } i = j \text{ and } i \in \mathcal{A} \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

With $\mathbf{M}_{\mathcal{A},l}$, we reformulate the HadamardMLP decoder as:

$$\begin{aligned} s_{ij} &= \phi^{\text{MLP}}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{w}_L^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{W}_{L-1} \dots \mathbf{M}_{\mathcal{A},1} \mathbf{W}_1 (\mathbf{x}_i \odot \mathbf{x}_j) \\ &= (\mathbf{W}_1^T \mathbf{M}_{\mathcal{A},1} \dots \mathbf{W}_{L-1}^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{w}_L \odot \mathbf{x}_i) \cdot \mathbf{x}_j \end{aligned} \quad (7)$$

Because the vector $\mathbf{W}_1^T \mathbf{M}_{\mathcal{A},1} \dots \mathbf{W}_{L-1}^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{w}_L$ is determined by the weights of MLP and the activated neurons \mathcal{A} , we term it as $\text{MLP}_{\mathcal{A}}(\cdot)$:

$$\text{MLP}_{\mathcal{A}}(\cdot) := \mathbf{W}_1^T \mathbf{M}_{\mathcal{A},1} \dots \mathbf{W}_{L-1}^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{w}_L \quad (8)$$

Given the source node i , because the score s_{ij} is obtained by the dot product between $(\mathbf{W}_1^T \mathbf{M}_{\mathcal{A},1} \dots \mathbf{W}_{L-1}^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{w}_L \odot \mathbf{x}_i)$ and the neighbor embedding \mathbf{x}_j , we term the former vector as the query embedding \mathbf{q} :

$$\mathbf{q} := \mathbf{W}_1^T \mathbf{M}_{\mathcal{A},1} \dots \mathbf{W}_{L-1}^T \mathbf{M}_{\mathcal{A},L-1} \mathbf{w}_L \odot \mathbf{x}_i = \text{MLP}_{\mathcal{A}}(\cdot) \odot \mathbf{x}_i \quad (9)$$

In this way, we can reformulate the output of decoder $\phi^{\text{MLP}}(\mathbf{x}_i, \mathbf{x}_j)$ as

$$s_{ij} = \phi^{\text{MLP}}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{q} \cdot \mathbf{x}_j. \quad (10)$$

In practice, we can use the \mathbf{q} as the query embedding in MIPS to retrieve the neighbors of highest inner products, which correspond to the highest scores. Here, how to get the activated neurons \mathcal{A} so as to obtain the query embedding \mathbf{q} is an issue. Different node pairs activate different neurons \mathcal{A} . Initially, without knowing which neurons are activated, we first assume all the neurons are activated, i.e., we have the initial query embedding as:

$$\mathbf{q}[1] = \left(\prod_{i=1}^{L-1} \mathbf{W}_i^T \right) \mathbf{w}_L \odot \mathbf{x}_i \quad (11)$$

Algorithm 1 *Flashlight* : progressively “illuminates” the semantic space to retrieve the high scoring neighbors for the LP HadamardMLP decoders.

Input: A trained HadamardMLP decoder ϕ^{MLP} that outputs the logit s_{ij} for the input $\mathbf{x}_i \odot \mathbf{x}_j$. The set of nodes \mathcal{V} . The node embedding set $\mathcal{X} = \{\mathbf{x}_i | i \in \mathcal{V}\}$. A source node i . The number of iterations T . The number of neighbors to retrieve at every iteration: $\mathbf{N} = [N_1, N_2, \dots, N_T]$.

Output: The recommended neighbors \mathcal{N} for the source node i .

- 1: Initialize the set of retrieved recommended neighbors $\mathcal{N} \leftarrow \emptyset$
 - 2: Initialize the set of activated neurons as $\mathcal{A}[0]$ as all the neurons in MLP.
 - 3: **for** $t \leftarrow 1$ to T **do**
 - 4: Calculate the query embedding $\mathbf{q}[t] \leftarrow \mathbf{x}_i \odot \text{MLP}_{\mathcal{A}[t-1]}(\cdot)$.
 - 5: $\mathcal{N}[t] \leftarrow N_t$ neighbors in \mathcal{X} that maximizes the inner product with $\mathbf{q}[t]$.
 - 6: $\mathcal{X} \leftarrow \mathcal{X} \setminus \{\mathbf{x}_j | j \in \mathcal{N}[t]\}$.
 - 7: $j^*[t] = \arg \max_{j \in \mathcal{N}[t]} \text{MLP}(\mathbf{x}_i \odot \mathbf{x}_j)$
 - 8: $\mathcal{A}[t] \leftarrow A(\text{MLP}(\cdot), \mathbf{x}_i \odot \mathbf{x}_{j^*[t]})$.
 - 9: $\mathcal{N} \leftarrow \mathcal{N} \cup \mathcal{N}[t]$.
 - 10: **return** \mathcal{N}
-

149 This initial design can reflect the general trends of increasing the edge scores on LP, without restricting
 150 which neurons are activated. We use $\mathbf{q}[1]$ as the query embedding to retrieve the highest inner product
 151 neighbors as $\mathcal{N}[1]$ in the first iteration. Then, given the retrieved neighbors in the t th iteration as
 152 $\mathcal{N}[t]$, we analyze the $\mathcal{N}[t]$ and adaptively adjust the query embedding $\mathbf{q}[t+1]$ that we use in the
 153 next iteration to find more high scoring neighbors. Specifically, we operate the feed-forward to MLP
 154 for $\mathcal{N}(t)$. We define the function $A(\cdot, \cdot)$ that returns the set of activated neurons for a MLP (the first
 155 input) with the input $\mathbf{x}_i \odot \mathbf{x}_j$ (the second input). Then we can use it to extract \mathcal{A} as:

$$\mathcal{A} = A(\text{MLP}(\cdot), \mathbf{x}_i \odot \mathbf{x}_j). \quad (12)$$

156 Then, we obtain the set of activated neurons of the highest scored neighbor at the t th iteration as:

$$\mathcal{A}[t] \leftarrow A(\text{MLP}(\cdot), \mathbf{x}_i \odot \mathbf{x}_{j^*[t]}), \text{ where } j^*[t] = \arg \max_{j \in \mathcal{N}[t]} \text{MLP}(\mathbf{x}_i \odot \mathbf{x}_j). \quad (13)$$

157 This implies that the neighbors activating $\mathcal{A}[t]$ can obtain the high edge scores. Then, if we take $\mathcal{A}[1]$
 158 as the set of neurons that we activate at the next query, we could find more high scoring neighbors. In
 159 this way, we set the neurons that we assume to activate in the next iteration as $\mathcal{A}[t]$. We repeat the
 160 above iterations until enough neighbors are retrieved. The algorithm is summarized in Alg. 1.

161 We name our algorithm as Flashlight because it works like a flashlight to progressively “illuminates”
 162 the semantic space to find the high scoring neighbors. The query embeddings are like the lights sent
 163 from the flashlight. And our process of adjusting the query embeddings is just like progressively
 164 adjusting the “lights” from the “flashlight” by checking the “objects” found in the last “illumination”.

165 In the experiments, we find that our Flashlight algorithm is effective to find the top scoring neighbors
 166 from the massive candidate neighbors. For example, in Fig. 3, our Flashlight is able to find the top
 167 100 scoring neighbors from nearly three million candidates by retrieving only 200 neighbors in the
 168 large OGBL-CITATION2 graph dataset for the HadamardMLP decoders.

169 **Complexity Analysis.** Using MLP decoders to compute the LP probabilities of all the neighbors
 170 holds the complexity as $\mathcal{O}(N)$, where N is the number of nodes in the whole graph. Finding the top
 171 scoring neighbors from the exact probabilities of all the neighbors also holds the linear complexity
 172 $\mathcal{O}(N)$. Overall, using MLP decoders to find the top scoring neighbors is of the time complexity
 173 $\mathcal{O}(N)$. In contrast, our Flashlight progressively calls the MIPS techniques for a constant number
 174 of times invariant to the graph data, which leads to the sublinear complexity as same as MIPS. In
 175 conclusion, our Flashlight improves the scalability and applicability of HadamardMLP decoders by
 176 reducing their inference time complexity from linear to sublinear time.

Table 1: Statistics of datasets.

Dataset	OGBL-DDI	OGBL-COLLAB	OGBL-PPA	OGBL-CITATION2
#Nodes	4,267	235,868	576,289	2,927,963
#Edges	1,334,889	1,285,465	30,326,273	30,561,187

5 Related Work

5.1 Link Prediction Models

Existing LP models can be categorized into three families: heuristic feature based [3, 9, 21–23], latent embedding based [12, 24–28], and neural network based ones. The neural network-based link prediction models are mainly developed in recent years, which explore non-linear deep structural features with neural layers. Variational graph auto-encoders [13] predict links by encoding graph with graph convolutional layer [5]. Another two state-of-the-art neural models WLN [29] and SEAL [30] use graph labeling algorithm to transfer union neighborhood of two nodes (enclosing subgraph) as meaningful matrix and employ convolutional neural layer or a novel graph neural layer DGCNN [31] for encoding. More recently, [6, 8] summarized the architectures LP models, and formally define the encoders and decoders.

Different from the previous work, we focus on analyzing the effectiveness of different LP decoders and improving the scalability of the effective LP decoders. In practice, we find that the Hadamard decoders exhibit superior effectiveness but poor scalability for inference. Our work significantly accelerates the inference of HadamardMLP decoders to make the effective LP scalable.

5.2 Maximum Inner Product Search

Finding the top scoring neighbors for the Dot Product decoder at the sublinear time complexity is a well studied research problem, known as the approximate maximum inner product search (MIPS). There are several approaches to MIPS: sampling based [11, 32, 33], LSH-based [34–37], graph based [38–40], and quantization approaches [17, 18]. MIPS is a fundamental building block in various application domains [41–46], such as information retrieval [47, 48], pattern recognition [49, 50], data mining [51, 52], machine learning [53, 54], and recommendation systems [55, 56].

With the explosive growth of datasets’ scale and the inevitable curse of dimensionality, MIPS is essential to offer the scalable services. However, the HadamardMLP decoders are nonlinear and there do not exist the well studied sublinear complexity algorithms to find the top scoring neighbors for HadamardMLP [10]. In this work, we utilize the well studied approximate MIPS techniques with the adaptively adjusted query embeddings to find the top scoring neighbors for the MLP decoders in a progressive manner. Our method supports the plug-and-play use during inference and significantly accelerates the LP inference with the effective MLP decoders.

6 Experiments

In this section, we first compare the effectiveness of different LP decoders. We find that the HadamardMLP decoders generally perform better than other decoders. Then, we implement our Flashlight algorithm with LP models to show that Flashlight effectively retrieves the top scoring neighbors for the HadamardMLP decoders. As a result, the inference efficiency and scalability of HadamardMLP decoders are improved significantly by our work.

6.1 Datasets

We evaluate the link prediction on Open Graph Benchmark (OGB) data [57]. We use four OGB datasets with different graph types, including OGBL-DDI, OGBL-COLLAB, OGBL-CITATION2, and OGBL-PPA. OGBL-DDI is a homogeneous, unweighted, undirected graph, representing the drug-drug interaction network. Each node represents a drug. Edges represent interactions between drugs. OGBL-COLLAB is an undirected graph, representing a subset of the collaboration network between authors indexed by MAG. Each node represents an author and edges indicate the collaboration between authors. All nodes come with 128-dimensional features. OGBL-CITATION2 is a directed graph, representing the citation network between a subset of papers extracted from MAG. Each

Table 2: The test effectiveness comparison of LP decoders on four OGB datasets (DDI, COLLAB, PPA, and CITATION2) [16]. We report the results of the standard metrics averaged over 10 runs following the existing work [6, 16]. HadamardMLP is more effective than other decoders. Flashlight effectively retrieves the top scoring neighbors for HadamardMLP and keep its exact outputs.

Decoder	Dot Product	Bilinear	ConcatMLP	HadamardMLP	HadamardMLP w/ Flashlight
OGBL-DDI					
GCN [5]	13.8 ± 1.8	16.1 ± 1.2	12.9 ± 1.4	37.1 ± 5.1	37.1 ± 5.1
GraphSAGE [12]	36.5 ± 2.6	39.4 ± 1.7	34.2 ± 1.9	53.9 ± 4.7	53.9 ± 4.7
Node2Vec [27]	11.6 ± 1.9	13.8 ± 1.6	10.8 ± 1.7	23.3 ± 2.1	23.3 ± 2.1
OGBL-COLLAB					
GCN [5]	42.9 ± 0.7	43.2 ± 0.9	42.3 ± 1.0	44.8 ± 1.1	44.8 ± 1.1
GraphSAGE [12]	37.3 ± 0.9	41.5 ± 0.8	37.0 ± 0.7	48.1 ± 0.8	48.1 ± 0.8
Node2Vec [27]	27.7 ± 1.1	31.5 ± 1.0	27.2 ± 0.8	48.9 ± 0.5	48.9 ± 0.5
OGBL-PPA					
GCN [5]	5.1 ± 0.4	5.8 ± 0.5	6.2 ± 0.6	18.7 ± 1.3	18.7 ± 1.3
GraphSAGE [12]	3.2 ± 0.3	6.5 ± 0.7	5.8 ± 0.4	16.6 ± 2.4	16.6 ± 2.4
Node2Vec [27]	4.2 ± 0.5	7.8 ± 0.6	8.3 ± 0.4	22.3 ± 0.8	22.3 ± 0.8
OGBL-CITATION2					
GCN [5]	65.3 ± 0.4	69.0 ± 0.8	62.7 ± 0.3	84.7 ± 0.2	84.7 ± 0.2
GraphSAGE [12]	62.2 ± 0.7	65.4 ± 0.9	60.8 ± 0.6	80.4 ± 0.1	80.4 ± 0.1
Node2Vec [27]	52.7 ± 0.8	54.1 ± 0.6	51.4 ± 0.5	61.4 ± 0.1	61.4 ± 0.1

node is a paper with 128-dimensional word2vec features. OGBL-PPA is an undirected, unweighted graph. Nodes represent proteins from 58 different species, and edges indicate biologically meaningful associations between proteins. The statistics of these datasets is presented in Table. 1.

6.2 Hyper-parameter Settings

For all experiments in this section, we report the average and standard deviation over ten runs with different random seeds. The results are reported on the the best model selected using validation data. We set hyper-parameters of the used techniques and considered baseline methods, e.g., the batch size, the number of hidden units, the optimizer, and the learning rate as suggested by their authors. We use the recent MIPS method ScaNN [18] in the implementation of our Flashlight. For the hyper-parameters of our Flashlight, we have found in the experiments that the performance of Flashlight is robust to the change of hyper-parameters in a board range. Therefore, we simply set the number of iterations of our Flashlight as $T = 3$ and the number of retrieved neighbors constant as 200 per iteration by default. We run all experiments on a machine with 80 Intel(R) Xeon(R) E5-2698 v4 @ 2.20GHz CPUs, and a single NVIDIA V100 GPU with 16GB RAM.

6.3 Effectiveness of Link Prediction Decoders

We follow the standard benchmark settings of OGB datasets to evaluate the effectiveness of LP with different decoders. The benchmark setting of OGBL-DDI is to predict drug-drug interactions given information on already known drug-drug interactions. The performance is evaluated by Hits@20: each true drug interaction is ranked among a set of approximately 100,000 randomly-sampled negative drug interactions, and count the ratio of positive edges that are ranked at 20-place or above. The task of OGBL-COLLAB is to predict the future author collaboration relationships given the past collaborations. Evaluation metric is Hits50, where each true collaboration is ranked among a set of 100,000 randomly-sampled negative collaborations. The task of OGBL-PPA is to predict new association edges given the training edges. Evaluation metric is Hits@100, where each positive edge is ranked among 3,000,000 randomly-sampled negative edges. The task of OGBL-CITATION2 is predict missing citation given existing citations. The evaluation metric is Mean Reciprocal Rank (MRR), where the reciprocal rank of the true reference among 1,000 sampled negative candidates is calculated for each source nodes, and then the average is taken over all source nodes.

We implement different decoders as introduced in Sec. 2, including the Dot Product, Bilinear, ConcatMLP, and the HadamardMLP decoders, over the LP encoders, including GCN [5], GraphSAGE [12], and Node2Vec [27], to compare the effects of different decoders on the LP effectiveness. We present the results on the OGBL-DDI, OGBL-COLLAB, OGBL-PPA, and OGBL-CITATION2 datasets in Table. 2. We observe that the HadamardMLP decoder outperforms other decoders on all

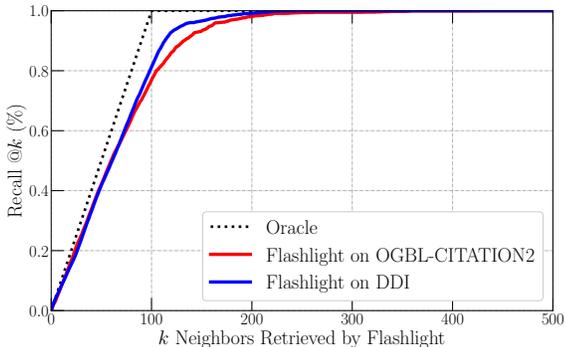
254 encoders and datasets. Our Flashlight algorithm can effectively retrieve the top scoring neighbors for
 255 the HadamardMLP decoder and keep the exact LP probabilities of HadamardMLPs’ output, which
 256 leads to the same results of the HadamardMLP decoder with and without Flashlight.

257 Note that the benchmark settings of these datasets sample a small portion of negative edges for the test
 258 evaluation, which is not challenging enough to evaluate the scalability of LP decoders on retrieving the
 259 top scoring neighbors from massive candidates in practice.

260 6.4 The Flashlight Algorithm Effectively Finds the Top Scoring Neighbors

261 To evaluate the effectiveness of our Flashlight on retrieving the top scoring neighbors for the
 262 HadamardMLP decoder, we propose a more challenging test setting for the OGB LP datasets.
 263 Given a source node, we takes its top 100 scoring neighbors of the HadamardMLP decoder as the
 264 ground-truth for retrievals. We set the task as retrieving k neighbors for a source node that can match
 265 the ground-truth neighbors as much as possible. We formally define the metric as Recall@ k , which is
 266 the portion of the ground-truth neighbors being in the top k neighbors retrieved by different methods.

267 We sample 1000 nodes as the source nodes from the OGBL-DDI and OGBL-
 268 CITATION2 datasets respectively for evaluation. We evaluate the effectiveness of our
 269 Flashlight algorithm by checking whether it can find the top scoring neighbors for every
 270 source node. We set the number of Flashlight iterations as 10 and the number of re-
 271 trieved neighbors per iteration as 50. We present the Recall@ k for k from 1 to 500
 272 averaged over all the source nodes in Fig. 3. The “oracle” curve represents the perfor-
 273 mance of a optimum searcher, of which the top k scoring neighbors are exactly the top
 274 retrieved top k neighbors are exactly the top k scoring neighbors of HadamardMLP.



280 When $k = 100$, the 100 neighbors retrieved
 281 by our Flashlight can cover more than 80%
 282 ground-truth neighbors. When $k \geq 200$,
 283 the recall reaches 100%. As a comparison,
 284 if we randomly sample the candidate neighbors for retrievals, the Recall@ k grows linearly with k
 285 and is less than 1×10^{-4} for $k = 100$ on the OGBL-CITATION2 dataset. The curves of Flashlight is
 286 close the optimum curve of the “oracle”. These results demonstrate the highly effectiveness of our
 287 Flashlight on finding the top scoring neighbors.

Figure 3: Recall@ k is the fraction of the 100 top scoring neighbors of HadamardMLP ranked in the top k neighbors retrieved by Flashlight. We report Recall@ k averaged over all the source nodes on OGBL-CITATION2 and OGBL-DDI.

288 Given the large OGBL-Citation2 dataset and smaller DDI dataset, our Flashlight exhibits similar
 289 Recall@ k performance given different numbers k of retrieved neighbors. This implies that our
 290 Flashlight can accurately find the top scoring neighbors for both small and large graphs.

293 6.5 Inference Efficiency of Link Prediction with Our Flashlight Algorithm

294 We use the throughputs to evaluate the inference speed of neighbor retrieval of different methods.
 295 The throughput is defined as how many source nodes that a method can serve to retrieve the top
 296 100 scoring neighbors per second. Except for the LP models that follow the encoder and decoder
 297 architectures, e.g., GraphSAGE [12], GCN [5], and PLNLP [6], there are some subgraph based LP
 298 models, e.g., SUREL [7] and SEAL [58]. The common issue of the subgraph based models is the poor
 299 efficiency: they have to crop a separate subgraph for every node pair to calculate the LP probability
 300 on the node pair. In this sense, the node embeddings cannot be shared on the LP calculation for
 301 different node pairs. This leads to the much lower inference speed of the subgraph based LP models
 302 than the encoder-decoder LP models. We compare the inference efficiency of different methods on
 303 the OGBL-CITATION2 dataset in Fig. 4, where we present the inference speed of different methods
 304 when achieving the 100% recall@100 for the top 100 scoring neighbors.

305 We observe that our Flashlight significantly accelerate the inference speed of LP models GraphSAGE
 306 [12], GCN [5], and PLNLP [6] with the HadamardMLP decoders by more than 100 times. This gap

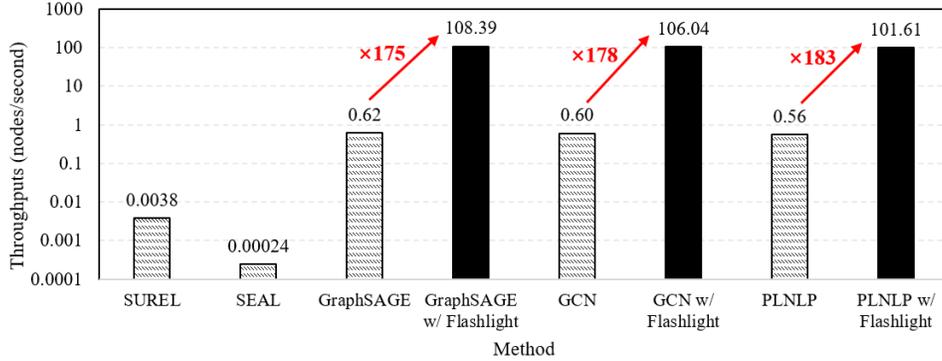


Figure 4: The inference speed of different LP methods on the OGBL-CITATION2 dataset. The y-axis (throughputs) is in the logarithmic scale.

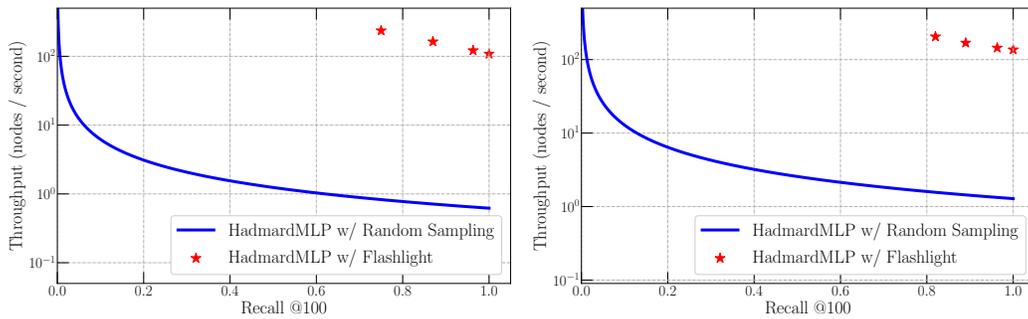


Figure 5: The tradeoff between the inference speed (y-axis) and the effectiveness of finding the top scoring neighbors (x-axis) on the OGBL-CITATION2 (left) and OGBL-PPA (right) datasets.

307 will be even larger for the datasets of larger scales, because the inference with our Flashlight holds the
 308 sublinear time complexity while the HadamardMLP decoders holds the linear complexity. Note that
 309 the y-axis is in logarithmic scale. The subgraph based methods SUREL [7] and SEAL [58] hold the
 310 inference speed of throughputs lower than 1×10^{-2} and 1×10^{-3} respectively, which is not applicable to
 311 the practical services that require the low latency of milliseconds.

312 Taking a further step, we comprehensively evaluate the tradeoff between the inference speed and the
 313 effectiveness of finding the top scoring neighbors. Taking GraphSAGE as the encoder, we present
 314 the tradeoff curves between the throughputs and the Recall@100 on the OGBL-CITATION2 and
 315 OGBL-PPI datasets in Fig. 5. In comparison with our Flashlight, we take the HadamardMLP decoder
 316 with the Random Sampling as the baseline for comparison. For example, on the OGBL-CITATION2
 317 dataset, when achieving the Recall@100 as more than 80%, the HadamardMLP with our Flashlight
 318 can serve more than 200 source nodes per second, while the HadamardMLP with the random sampling
 319 can only serve less than 1 node per second. Overall, our Flashlight achieves much better inference
 320 speed and effectiveness tradeoff than the HadamardMLP with random sampling.

321 7 Conclusion

322 Our theoretical and empirical analysis suggests that the HadamardMLP decoders are a better default
 323 choice than the Dot Product in terms of LP effectiveness. Because there does not exist a well-
 324 developed sublinear complexity top scoring neighbor searching algorithm for HadamardMLP, the
 325 HadamardMLP decoders are not scalable and cannot support the fast inference on large graphs. To
 326 resolve this issue, we propose the Flashlight algorithm to accelerate the inference of LP models with
 327 HadamardMLP decoders. Flashlight progressively operates the well-studied MIPS techniques for a
 328 few iterations. We adaptively adjust the query embeddings at every iteration to find more high scoring
 329 neighbors. Empirical results show that our Flashlight accelerates the inference of LP models by more
 330 than 100 times on the large OGBL-CITATION2 graph. Overall, our work paves the way for the use
 331 of strong LP decoders in practical settings by greatly accelerating their inference.

References

- 332
- 333 [1] Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications*, 390(6):1150–1170, 2011. 1
- 334
- 335 [2] Víctor Martínez, Fernando Berzal, and Juan-Carlos Cubero. A survey of link prediction in
336 complex networks. *ACM computing surveys (CSUR)*, 49(4):1–33, 2016. 1
- 337 [3] Lada A Adamic and Eytan Adar. Friends and neighbors on the web. *Social networks*, 25(3):
338 211–230, 2003. 1, 6
- 339 [4] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recom-
340 mender systems. *Computer*, 42(8):30–37, 2009. 1
- 341 [5] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional
342 networks. *arXiv preprint arXiv:1609.02907*, 2016. 1, 2, 3, 6, 7, 8
- 343 [6] Zhitao Wang, Yong Zhou, Litao Hong, Yuanhang Zou, and Hanjing Su. Pairwise learning for
344 neural link prediction. *arXiv preprint arXiv:2112.02936*, 2021. 1, 2, 3, 6, 7, 8, 13
- 345 [7] Haoteng Yin, Muhan Zhang, Yanbang Wang, Jianguo Wang, and Pan Li. Algorithm and
346 system co-design for efficient subgraph-based graph representation learning. *arXiv preprint*
347 *arXiv:2202.13538*, 2022. 1, 8, 9
- 348 [8] Chuxiong Sun and Guoshi Wu. Adaptive graph diffusion networks with hop-wise attention.
349 *arXiv preprint arXiv:2012.15024*, 2020. 1, 2, 3, 6, 13
- 350 [9] Tao Zhou, Linyuan Lü, and Yi-Cheng Zhang. Predicting missing links via local information.
351 *The European Physical Journal B*, 71(4):623–630, 2009. 1, 6
- 352 [10] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. Neural collaborative filtering vs.
353 matrix factorization revisited. In *Fourteenth ACM conference on recommender systems*, pages
354 240–248, 2020. 1, 2, 3, 4, 6, 13, 14
- 355 [11] Rui Liu, Tianyi Wu, and Barzan Mozafari. A bandit approach to maximum inner product search.
356 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4376–4383,
357 2019. 1, 6
- 358 [12] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large
359 graphs. *Advances in neural information processing systems*, 30, 2017. 2, 3, 6, 7, 8
- 360 [13] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint*
361 *arXiv:1611.07308*, 2016. 2, 6
- 362 [14] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure
363 Leskovec. Graph convolutional neural networks for web-scale recommender systems. In
364 *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery &*
365 *data mining*, pages 974–983, 2018. 2
- 366 [15] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of*
367 *control, signals and systems*, 2(4):303–314, 1989. 2
- 368 [16] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-
369 lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*,
370 2021. 3, 7
- 371 [17] Xinyan Dai, Xiao Yan, Kelvin KW Ng, Jiu Liu, and James Cheng. Norm-explicit quantization:
372 Improving vector quantization for maximum inner product search. In *Proceedings of the AAAI*
373 *Conference on Artificial Intelligence*, volume 34, pages 51–58, 2020. 4, 6
- 374 [18] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv
375 Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International*
376 *Conference on Machine Learning*, pages 3887–3896. PMLR, 2020. 4, 6, 7
- 377 [19] Shulong Tan, Zhixin Zhou, Zhaozhuo Xu, and Ping Li. Fast item ranking under neural network
378 based measures. In *Proceedings of the 13th International Conference on Web Search and Data*
379 *Mining*, pages 591–599, 2020. 4
- 380 [20] Rihan Chen, Bin Liu, Han Zhu, Yaoyuan Wang, Qi Li, Buting Ma, Qingbo Hua, Jun Jiang,
381 Yunlong Xu, Hongbo Deng, et al. Approximate nearest neighbor search under neural similarity
382 metric for large-scale recommendation. *arXiv preprint arXiv:2202.10226*, 2022. 4

- 383 [21] Gobinda G Chowdhury. *Introduction to modern information retrieval*. Facet publishing, 2010.
384 6
- 385 [22] David Liben-Nowell and Jon Kleinberg. The link prediction problem for social networks. In
386 *Proceedings of the twelfth international conference on Information and knowledge management*,
387 pages 556–559, 2003.
- 388 [23] Glen Jeh and Jennifer Widom. Simrank: a measure of structural-context similarity. In *Proceed-*
389 *ings of the eighth ACM SIGKDD international conference on Knowledge discovery and data*
390 *mining*, pages 538–543, 2002. 6
- 391 [24] Aditya Krishna Menon and Charles Elkan. Link prediction via matrix factorization. In *Joint*
392 *european conference on machine learning and knowledge discovery in databases*, pages 437–
393 452. Springer, 2011. 6
- 394 [25] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social repre-
395 sentations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge*
396 *discovery and data mining*, pages 701–710, 2014.
- 397 [26] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-
398 scale information network embedding. In *Proceedings of the 24th international conference on*
399 *world wide web*, pages 1067–1077, 2015.
- 400 [27] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In
401 *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and*
402 *data mining*, pages 855–864, 2016. 7
- 403 [28] Zhitao Wang, Chengyao Chen, and Wenjie Li. Predictive network representation learning for
404 link prediction. In *Proceedings of the 40th international ACM SIGIR conference on research*
405 *and development in information retrieval*, pages 969–972, 2017. 6
- 406 [29] Muhan Zhang and Yixin Chen. Weisfeiler-lehman neural machine for link prediction. In
407 *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and*
408 *data mining*, pages 575–583, 2017. 6
- 409 [30] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *Advances in*
410 *neural information processing systems*, 31, 2018. 6
- 411 [31] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning
412 architecture for graph classification. In *Proceedings of the AAAI conference on artificial*
413 *intelligence*, volume 32, 2018. 6
- 414 [32] Edith Cohen and David D Lewis. Approximating matrix multiplication for pattern recognition
415 tasks. *Journal of Algorithms*, 30(2):211–252, 1999. 6
- 416 [33] Hsiang-Fu Yu, Cho-Jui Hsieh, Qi Lei, and Inderjit S Dhillon. A greedy approach for budgeted
417 maximum inner product search. *Advances in neural information processing systems*, 30, 2017.
418 6
- 419 [34] Qiang Huang, Guihong Ma, Jianlin Feng, Qiong Fang, and Anthony KH Tung. Accurate
420 and fast asymmetric locality-sensitive hashing scheme for maximum inner product search. In
421 *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery &*
422 *Data Mining*, pages 1561–1570, 2018. 6
- 423 [35] Behnam Neyshabur and Nathan Srebro. On symmetric and asymmetric lshs for inner product
424 search. In *International Conference on Machine Learning*, pages 1926–1934. PMLR, 2015.
- 425 [36] Anshumali Shrivastava and Ping Li. Asymmetric lsh (alsh) for sublinear time maximum inner
426 product search (mips). *Advances in neural information processing systems*, 27, 2014.
- 427 [37] Xiao Yan, Jinfeng Li, Xinyan Dai, Hongzhi Chen, and James Cheng. Norm-ranging lsh for
428 maximum inner product search. *Advances in Neural Information Processing Systems*, 31, 2018.
429 6
- 430 [38] Jie Liu, Xiao Yan, Xinyan Dai, Zhirong Li, James Cheng, and Ming-Chang Yang. Understanding
431 and improving proximity graph based maximum inner product search. In *Proceedings of the*
432 *AAAI Conference on Artificial Intelligence*, volume 34, pages 139–146, 2020. 6
- 433 [39] Stanislav Morozov and Artem Babenko. Non-metric similarity graphs for maximum inner
434 product search. *Advances in Neural Information Processing Systems*, 31, 2018.

- 435 [40] Zhixin Zhou, Shulong Tan, Zhaozhuo Xu, and Ping Li. Möbius transformation for fast inner
436 product search on graph. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- 437 [41] Kazuo Aoyama, Kazumi Saito, Hiroshi Sawada, and Naonori Ueda. Fast approximate similarity
438 search based on degree-reduced neighborhood graphs. In *Proceedings of the 17th ACM SIGKDD
439 international conference on Knowledge discovery and data mining*, pages 1055–1063, 2011. 6
- 440 [42] Akhil Arora, Sakshi Sinha, Piyush Kumar, and Arnab Bhattacharya. Hd-index: Pushing the
441 scalability-accuracy boundary for approximate knn search in high-dimensional spaces. *arXiv
442 preprint arXiv:1804.06829*, 2018.
- 443 [43] Cong Fu, Chao Xiang, Changxu Wang, and Deng Cai. Fast approximate nearest neighbor search
444 with the navigating spreading-out graph. *arXiv preprint arXiv:1707.00143*, 2017.
- 445 [44] Yu A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search
446 using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and
447 machine intelligence*, 42(4):824–836, 2018.
- 448 [45] Philipp M Riegger. Literature survey on nearest neighbor search and search in graphs. 2010.
- 449 [46] Wenhui Zhou, Chunfeng Yuan, Rong Gu, and Yihua Huang. Large scale nearest neighbors
450 search based on neighborhood graph. In *2013 International Conference on Advanced Cloud
451 and Big Data*, pages 181–186. IEEE, 2013. 6
- 452 [47] Myron Flickner, Harpreet Sawhney, Wayne Niblack, Jonathan Ashley, Qian Huang, Byron Dom,
453 Monika Gorkani, Jim Hafner, Denis Lee, Dragutin Petkovic, et al. Query by image and video
454 content: The qbic system. *computer*, 28(9):23–32, 1995. 6
- 455 [48] Chun Jiang Zhu, Tan Zhu, Haining Li, Jinbo Bi, and Minghu Song. Accelerating large-scale
456 molecular similarity search through exploiting high performance computing. In *2019 IEEE
457 International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 330–333. IEEE,
458 2019. 6
- 459 [49] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on
460 information theory*, 13(1):21–27, 1967. 6
- 461 [50] Atsutake Kosuge and Takashi Oshima. An object-pose estimation acceleration technique for
462 picking robot applications by using graph-reusing k-nn search. In *2019 First International
463 Conference on Graph Computing (GC)*, pages 68–74. IEEE, 2019. 6
- 464 [51] Qiang Huang, Jianlin Feng, Qiong Fang, Wilfred Ng, and Wei Wang. Query-aware locality-
465 sensitive hashing scheme for l_p norm. *The VLDB Journal*, 26(5):683–708, 2017. 6
- 466 [52] Masajiro Iwasaki. Pruned bi-directed k-nearest neighbor graph for proximity search. In
467 *International Conference on Similarity Search and Applications*, pages 20–33. Springer, 2016.
468 6
- 469 [53] Yuan Cao, Heng Qi, Wenrui Zhou, Jien Kato, Keqiu Li, Xiulong Liu, and Jie Gui. Binary
470 hashing for approximate nearest neighbor search on big data: A survey. *IEEE Access*, 6:
471 2039–2054, 2017. 6
- 472 [54] Scott Cost and Steven Salzberg. A weighted nearest neighbor algorithm for learning with
473 symbolic features. *Machine learning*, 10(1):57–78, 1993. 6
- 474 [55] Yitong Meng, Xinyan Dai, Xiao Yan, James Cheng, Weiwen Liu, Jun Guo, Benben Liao,
475 and Guangyong Chen. Pmd: An optimal transportation-based user distance for recommender
476 systems. In *European Conference on Information Retrieval*, pages 272–280. Springer, 2020. 6
- 477 [56] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative
478 filtering recommendation algorithms. In *Proceedings of the 10th international conference on
479 World Wide Web*, pages 285–295, 2001. 6
- 480 [57] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
481 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
482 *Advances in neural information processing systems*, 33:22118–22133, 2020. 6
- 483 [58] Muhan Zhang, Pan Li, Yinglong Xia, Kai Wang, and Long Jin. Labeling trick: A theory of using
484 graph neural networks for multi-node representation learning. *Advances in Neural Information
485 Processing Systems*, 34:9061–9073, 2021. 8, 9

- 486 [59] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via
 487 over-parameterization. In *International Conference on Machine Learning*, pages 242–252.
 488 PMLR, 2019. 13
- 489 [60] Alexandr Andoni, Rina Panigrahy, Gregory Valiant, and Li Zhang. Learning polynomials with
 490 neural networks. In *International conference on machine learning*, pages 1908–1916. PMLR,
 491 2014. 13

492 A Learning a Dot Product decoder with a HadamardMLP decoder is Easy

493 Before we have discussed the limitations of the Dot Product decoder. An interesting questions is
 494 whether the HadamardMLP decoder can replace the Dot Product decoder by approximating it. If the
 495 MLP decoder can learn a dot product easily, it is safe to use MLP decoder instead of the dot product
 496 ones in most cases. There are similar problems actively studied in machine learning. Existing work
 497 imply that the difficulty scales polynomial with dimensionality d and $1/\epsilon$ in theory [10, 59, 60]. This
 498 motivates us to investigate the question empirically.

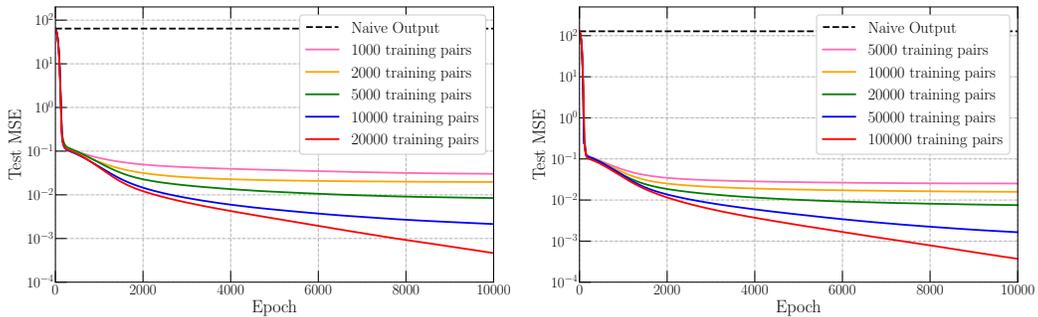


Figure 6: A MLP decoder can learn a Dot Product decoder well with enough training data. The left and right figures shows the MSE differences (y-axis) per epoch (x-axis) between the outputs of dot product and the MLP decoders given different training sizes with the input embedding dimensionality as $d = 64$ and $d = 128$ respectively. The naive output denotes the outputs of zeros.

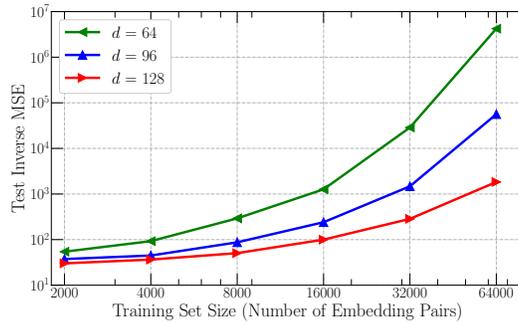


Figure 7: Test inverse MSE differences between the outputs of Dot Product and MLP decoders after convergence (y-axis) versus the training set size (x-axis).

499 We set up a synthetic learning task where given two embeddings $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^d$ and a label $\mathbf{x}_i \cdot \mathbf{x}_j$, we
 500 want to obtain a MLP function that approximates the $\mathbf{x}_i \cdot \mathbf{x}_j$ with the inputs $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^d$. For this
 501 experiment, we create the datasets including the embedding matrix as $\mathbf{E} \in \mathbb{R}^{10^6 \times d}$. We draw every
 502 row in \mathbf{E} from $\mathcal{N}(0, \mathbf{I})$ independently. Then, we uniformly sample (without replacement) 10^4 and S
 503 embedding pair combinations from \mathbf{E} to form the test and training sets (no overlap) respectively.

504 We train the MLP on the training and evaluate it on the test set. For the architecture of the MLP, we
 505 keep it simple: we follow the existing work [6, 8] to set the number of layers as 2 and the number of
 506 hidden units as same as the input embeddings: d . For the optimizer, we also follow the existing work
 507 [6, 8] to choose the Adam optimizer.

508 As for evaluation metrics, we compute the MSE (Mean Squared Error) differences between the
 509 predicted score of the MLP and the dot product decoders. We measure the MSE of a naive model that
 510 predicts always 0 (the average rating). Every experiment is repeated 5 times and we report the mean.

511 Fig. 6 shows the approximation errors on the MLP per epoch given different number of training
 512 pairs and dimensions. The figure suggests that an MLP can easily approximate the dot product with
 513 enough training data. Consistent with the theory, the number of samples needed scales polynomially
 514 with the increasing dimensions and reduced errors. Anecdotaly, we observe the number of needed
 515 training samples is about $\mathcal{O}(d^\alpha/\epsilon^\beta)$ for $\alpha \approx 2, \beta \ll 1$ (see Fig. 7). In all cases, the MSE errors of
 516 the MLP decoder are negligible compared with the naive output.

517 This experiment shows that an MLP can easily approximate the dot product with enough training
 518 data. We hope this can explain, at least partially, why the MLP decoder generally performs better
 519 than the dot product.

520 Our conclusion seems to be distinct to to the existing work [10], which claims that the ConcatMLP
 521 is hard to learn a Dot Product. Actually, our conclusion is not conflicted with that in [10]. This
 522 ConcatMLP decoder processes the concatenation of the paired embeddings instead of the Hadamard
 523 product of the paired embeddings as the HadamardMLP. The HadamardMLP holds the inductive bias
 524 similar to the Dot Product, which makes the former easily learns the latter. Actually, we show that
 525 a simple two-layer MLP with only two hidden units is equivalent to the Dot Product with specific
 526 weights. We assign the first layer weights for two hidden units as $\mathbf{1}$ and $-\mathbf{1}$ and the second layer
 527 weights as ones. Then, we have its output as:

$$s_{ij} = \phi^{\text{MLP}}(\mathbf{x}_i, \mathbf{x}_j) = \text{ReLU}(\mathbf{1} \bullet (\mathbf{x}_i \odot \mathbf{x}_j)) + \text{ReLU}(-\mathbf{1} \bullet (\mathbf{x}_i \odot \mathbf{x}_j)) = \mathbf{1} \bullet (\mathbf{x}_i \odot \mathbf{x}_j) = \mathbf{x}_i \bullet \mathbf{x}_j, \quad (14)$$

528 which is equivalent to the Dot Product decoder. From this result, we find that any MLP decoder with
 529 the careful initialization is equivalent to the Dot Product decoder and thus can learn the Dot Product
 530 easily.