Deep Distance Sensitivity Oracles

Davin Jeong[1] [*], **Chau Pham**[2], **Arnav Bhakta**[3], **Sarel Cohen**[4], **Maximilian Katzmann**[5], **Tobias Friedrich**[6], **Sang Chin**[2, 7]

[1] Phillips Academy, Andover [2] Boston University [3] Yale University [4] The Academic College of Tel Aviv-Yafo [5] Karlsruhe Institute of Technology [6] Hasso Plattner Institute [7] Dartmouth College

[*] Corresponding author. Email: djeong23@andover.edu

Abstract

One of the most fundamental graph problems is finding a shortest path from a source to a target node. While in its basic forms the problem has been studied extensively and efficient algorithms are known, it becomes significantly harder as soon as parts of the graph are susceptible to failure. Although one can recompute a shortest replacement path after every outage, this is rather inefficient both in time and/or storage.

One way to overcome this problem is to shift computational burden from the queries into a pre-processing step, where a data structure is computed that allows for fast querying of replacement paths, typically referred to as a *Distance Sensitivity Oracle (DSO)*.

Such a DSO can be constructed, by utilizing the fact that the shortest path in a graph with failures consists of a concatenation of shortest paths in the original graph. That is, one can quickly compute a replacement path by determining a suitable *pivot* and concatenating the shortest paths from source to pivot and from pivot to target.

Unfortunately, finding such a pivot is far from trivial. In this paper, we utilize node2vec, graph attention networks, and multi-layer perceptrons to find pivots such that the lengths of the resulting paths are close to that of the shortest replacement paths. More precisely, we evaluate our technique on a collection of real-world networks and observe that on average the computed replacement paths are longer by merely a few percentages compared to the optimal solution.

While DSOs have been extensively studied in the theoretical computer science community, to the best of our knowledge this is the first work to construct DSOs using deep learning techniques. In particular, in this work we first prove the existence of a combinatorial structure and then build on top of it a deep learning algorithm that utilizes the observed combinatorial structure. As mentioned above - we utilize the combinatorial structure of replacement paths as a concatenation of shortest paths and use deep learning to find the pivot nodes for stitching shortest paths into replacement paths.

Introduction

The shortest path problem is frequently encountered in the real-world. In road networks, users want to know how long it will take to get from one place to another (Huang et al. 2021). In biological networks, consisting of genes and their products, the shortest paths are used to find clusters and identify core pathways (Ren, Ay, and Kahveci 2018). In social networks, the number of connections between users can be used for friend recommendation (Tian et al. 2012). In the web search, relevant web pages can be ranked by their distances from queried terms (Ukkonen et al. 2008).

For graphs in the real world, often consisting of millions of nodes, special data structures called *Distance Oracles* (DO) are used to store information about distances of an input graph $G = (V, E)$ with n vertices and m edges. Without storing the entire graph, they can quickly retrieve important distance information to answer the shortest path queries. These shift the computational burden to the preprocessing step, so that queries can be answered quickly.

However, in addition to being large in size, real-world networks are also frequently susceptible to failures. For example, in road networks, a construction, a traffic accident, or an event might temporarily block nodes. In social networks, users might temporarily deactivate their accounts, resulting in a null node. And on the internet, web servers may be temporarily down due to mechanical failures or malicious attacks (Billand et al. 2016). In these instances, we desire a method that can continue answering shortest path queries without stalling or having to recompute shortest paths on the entire graph again.

Distance Sensitivity Oracles (DSO) are a type of DO that can respond to queries of the form (s, t, f), requesting the shortest path between nodes s and t when a vertex f fails and is thus unavailable. Desirable DSOs should provide reasonable trade-offs among space consumption, query time, and *MRE* (i.e., quality of the estimated distance).

Contributions

We have tested our method on a variety of real-world networks and achieved state-of-the-art performance on all of them, our accuracy even outperforms models for shortest paths without node failures. Our contributions mainly lie in three aspects:

- We prove that one can use deep learning to find pivot nodes in distance sensitivity oracles. In fact, to the best of our knowledge, we are the first to use deep learning to build a distance sensitivity oracle.

- In our theoretical analysis we first present a simple proof for the existence of an underlying combinatorial structure for replacement paths, and later we utilize it to implement the first distance sensitivity oracle using deep learning.

- We empirically evaluate our method and compare it with related works to demonstrate near-exact accuracy across a diverse range of real-world networks.

Related Work

Distance-Sensitivity Oracles Suppose we are given an undirected, unweighted graph $G = (V, E)$, and we want to build a data structure that, given any two vertices $u, v \in V$ and a failing vertex $f \in V$, outputs the length of the shortest path from u to v that does not go through f. Such a data structure is called a distance sensitivity oracle (or DSO for short). If the output is only approximate, then we say that the DSO has stretch α if the length of the path reported is at most α times the length of the shortest path from u to v that does not go through f.

The problem of constructing DSOs is motivated by the fact that real-life networks often suffer from failures. Suppose we have a network with n nodes and m links, and we want to send a package from a node u to another node v. Normally, it suffices to compute the shortest path from u to v. However, if some node f in this network fails, then our path cannot use f, and our task becomes to find the shortest path from u to v that does not go through f. Usually, there is only a very small number of failures. In this paper, we consider the simplest case, in which there is only one failed node.

The problem of constructing a DSO is well-studied in the theoretical computer science community: Demetrescu *et al.* (Demetrescu et al. 2008) showed that given a graph $G = (V, E)$, there is a DSO which occupies $O(n^2 \log n)$ space, and can answer a query in constant time. The preprocessing of this DSO, that is the time it takes to construct this DSO, is $O(mn^2 + n^3 \log n)$. Several theoretical results attempted to improve the preprocessing time required by the DSO. Bernstein and Karger (Bernstein and Karger 2008, 2009) improved this time bound to $\widetilde{O}(mn)$ [1]. Note that the All-Pairs Shortest Paths (APSP) problem, which only asks the distances between each pair of vertices u, v, is conjectured to require $mn^{1-o(1)}$ time to solve (Lincoln, Williams, and Williams 2018). Since we can solve the APSP problem by using a DSO, the preprocessing time $\widetilde{O}(mn)$ is theoretically asymptotically optimal in this sense, up to a polylogarithmic factor (note that, in practice, such polylogarithmic factors may be very large). Several additional results improved upon the theoretical preprocessing time by using fast matrix multiplication (Weimann and Yuster 2013; Grandoni and Williams 2020; Chechik and Cohen 2020; Ren 2022; Gu and Ren 2021).

Considering the size of the oracle, Duan and Zhang (Duan and Zhang 2017) improved the space complexity of (Demetrescu et al. 2008) to $O(n^2)$, which is from a theoretical perspective asymptotically optimal for dense graphs (*i.e.*, $m = \Theta(n^2)$). To do so, Duan and Zhang store multiple data-structures, which is reasonable for a theoretical work, however from a practical perspective the hidden constant is large.

_____

[1] For a non-negative function $f = f(n)$, we use $\widetilde{O}(f)$ to denote $O(f \cdot \mathsf{polylog}(n))$.

Therefore, it may also be interesting to consider DSOs with smaller space, at the cost of an approximate answer.

Here are several DSOs that provide tradeoffs between the size of DSO and the approximation guarantee (also called the stretch).

- The DSO described in (Baswana and Khanna 2013), for every parameter $\epsilon > 0$ and integer $k \geq 1$ has stretch $(2k - 1)(1 + \epsilon)$ and size $O(k^5 n^{1+1/k} \log^3 n/\epsilon^4)$.
- The DSO described in (Chechik et al. 2012), for every integer parameter $k \geq 1$ has stretch $(16k - 4)$ and size $O(kn^{1+1/k} \log n)$.

Note that even though the size of the above two DSOs for $k \geq 2$ is asymptotically smaller than $O(n^2)$, the stretch guarantee is at least 3 in (Baswana and Khanna 2013) and at least 28 in (Chechik et al. 2012), which is far from the optimum and may not be practical in many applications.

In this work we construct the first DSO that is built using deep learning. Our DSO computes near-optimal paths, with an approximation ratio that is close to 1, as shown in the experiments described in Section . Our method uses deep learning to find pivots (as described in Section), utilizing a combinatorial structural property that we prove in Section .

Shortest Paths using Deep Learning Real-world networks, such as those representing social media interactions, often contain more than hundreds of thousands of nodes (Backstrom et al. 2006). The high-dimensionality and sparsity of these networks make traditional graph analytics methods impractical (Cai, Zheng, and Chang 2018). Thus, graph embeddings have been used to represent graphs in a low-dimensional space. More specifically, nodes and edges may be represented as vectors that still preserve information and properties of the original graph.

Approaches towards generating graph embeddings include matrix factorization, deep learning with and without random walks, edge reconstruction, graph kernels, and generative models (Cai, Zheng, and Chang 2018). By either optimizing embeddings for their specific task or using general embedding techniques like node2vec, these graph embeddings may then be combined with existing techniques to tackle tasks such as node classification, node clustering, and link detection (Wang et al. 2017) (Zhang et al. 2021) (Crichton et al. 2018) (Kipf and Welling 2016).

Using these graph embeddings, previous works towards answering shortest path queries using deep learning have employed a two-stage solution: 1) representation learning and 2) distance prediction.

Among the first to apply graph embeddings to the shortest paths problem was Orion (Wang et al. 2002). Inspired by the successes of virtual coordinate systems, a landmark labelling approach was employed, where positions of all nodes were chosen based on their relative distances to a fixed number of landmarks. Using the Simplex Downhill algorithm, representations were found in a Euclidean coordinate space, allowing constant time distance calculations and producing mean relative error (MRE) between 15% - 20% (Costa et al. 2004; Nelder and Mead 1965; Wang et al. 2002).

Other existing coordinate systems have also been used. Building off of network routing schemes in hyperbolic

spaces, Rigel used a hyperbolic graph coordinate system to reduce the MRE to 9% and found that the hyperbolic space performed empirically better across distortion metrics than Euclidean and spherical coordinate systems (Cvetkovski and Crovella 2009; Zhao et al. 2011). In road networks, geographical coordinates have been utilized with a multi-layer perceptron to predict distances between locations with 9% MRE (Jindal et al. 2017). In addition to these coordinate systems, general graph embedding techniques have also been employed to handle shortest path queries. Node2vec and Poincare embeddings were used with a feed forward neural network to approximate shortest paths distances on large social networks and reported an MRE between 3% to 7% (Rizi, Schloetterer, and Granitzer 2018). More recently, graph embeddings have also been learned alongside the distance predictors, to produce representations more specific to the shortest path task. Vdist2vec directly learned vertex embeddings by passing the gradient from the distance predictor back to a $N \times k$ matrix, achieving an MRE between 1% to 7% (Qi et al. 2020). Huang et al. computed shortest path distances on road networks using a hierarchical embedding model and achieved an MRE of 0.7% (Huang et al. 2021).

This common two-stage solution, of finding representations and predicting their distances, directly computes the distances between nodes. However, in the real-world, the paths themselves are also often needed. Using the above methods to query for the shortest path length, routing algorithms can be constructed by performing a breadth-first search, where at each step only neighbors whose distance to the target is close to the current nodes' distance to the target are visited (Zhao et al. 2011). Another framework finds the shortest path in the routing process, decomposing a large shortest path instance and solving it using reinforcement learning (Yin, Rao, and Zhang 2021). Though the accuracy of these routing algorithms does depend on their distance predictors, they remain limited in their efficiency as iterative algorithms require multiple predictions.

Thus, current works have found effective methods towards estimating the distance between two nodes using graph embeddings and deep learning. The main avenues of development are representations, which accurately capture both the graphs' structures and the shortest path task, and predictors, which can be used to efficiently query both distances and the corresponding paths. Taking these lessons to the replacement path task, we propose a *pivot-based method*, which combines theoretical computer science and deep learning. To the best of our knowledge, our method is the first to handle node failures, completing the DSO data structure. Further, we efficiently calculate not only the length but also the actual shortest path in an end-to-end framework by exploiting combinatorial structures, which avoids potential inaccuracies that arise in iterative path-finding processes.

Theoretical Analysis

In this section we consider a combinatorial structural property of replacement paths: such a path is a concatenation of a few original shortest paths. As we previously described, later on our deep learning algorithm builds on this lemma.

Our research is motivated by the following lemma from Afek *et al.*.

Lemma 1. *(Afek et al. 2002) After k edge failures in an unweighted graph, each new shortest path is the concatenation of at most $k + 1$ original shortest paths.*

In other words, the replacement path can be defined using so called pivot nodes that specify at which nodes in the graph the shortest paths may be stitched together. In this work we are interested in the failure of a single node, which is equivalent to the failure of its incident edges. The number of concatenations (and with that the number of corresponding pivots) required to obtain the replacement path then depends on the degree of the failed node. While in real-world networks the average degree is often rather small, finding suitable pivots remains a hard task. To overcome this problem, we consider an approximate setting, where we allow for a slack in the quality of the obtained paths (they may be longer than a shortest replacement path) but where only one pivot node is used. From the theoretical perspective Lemma 1 proves that for a single edge failure it is sufficient to find a single pivot. To give more insights about the theoretical background of the case where a single pivot is enough, we present a simplification of the proof of Lemma 1 below.

Lemma 2. *After an edge failure in an unweighted undirected graph, each new shortest path is the concatenation of at most two original shortest paths.*

Proof. Let $P(s, t)$ be a shortest path from s to t, let f be the failing edge along $P(s, t)$, and let $P(s, t, f) = \langle v_0, \ldots, v_k \rangle$ be a replacement path, that is a shortest path from s to t in $G - \{f\}$. If $P(s, t, f)$ is also a shortest path in the original graph G then the lemma follows trivially. For the rest of the proof assume that $P(s, t, f)$ is not a shortest path in G. We refer to Figure 1 for a visualization of the argumentation that follows. Let $0 \le i \le k$ be the maximum index such that $P(s, v_i) := \langle v_0, \ldots, v_i \rangle$ is a shortest path in G. Since $P(s, t, f)$ is not a shortest path in G it must hold that $i < k$ and hence $\langle v_0, \ldots, v_{i+1} \rangle$ is not a shortest path in G. Let $P(s, v_{i+1})$ be a shortest path from s to v_{i+1} in G, it follows that $|P(s, v_{i+1})| < |\langle v_0, \ldots, v_{i+1} \rangle| = i + 1$. Let P_1 be the concatenation of the path $P(s, v_{i+1})$ with the edge (v_{i+1}, v_i), then it holds that

$$|P_1| < |\langle v_0, \ldots, v_{i+1} \rangle| + 1 = i + 2 \qquad (1)$$

We claim that $\langle v_i, \ldots, v_k \rangle$ must be a shortest path in G, and then $P(s, t, f)$ is a concatenation of two shortest paths in G (the first being $\langle v_0, \ldots, v_i \rangle$ and the second being $\langle v_i, \ldots, v_k \rangle$). For the rest of the proof, we assume by contradiction that $\langle v_i, \ldots, v_k \rangle$ is not a shortest path in G. Let $P_2 := P(v_i, t)$ be a shortest path from v_i to t in G, then it follows that

$$|P_2| = |P(v_i, t)| < |\langle v_i, \ldots, v_k \rangle| = k - i \qquad (2)$$

Note that both P_1 and P_2 contain the edge f:

- We claim that the path P_1 contains f as P_1 contains the shortest path $P(s, v_{i+1})$. Note that by definition of $i < k$ as the maximum index for which the subpath of
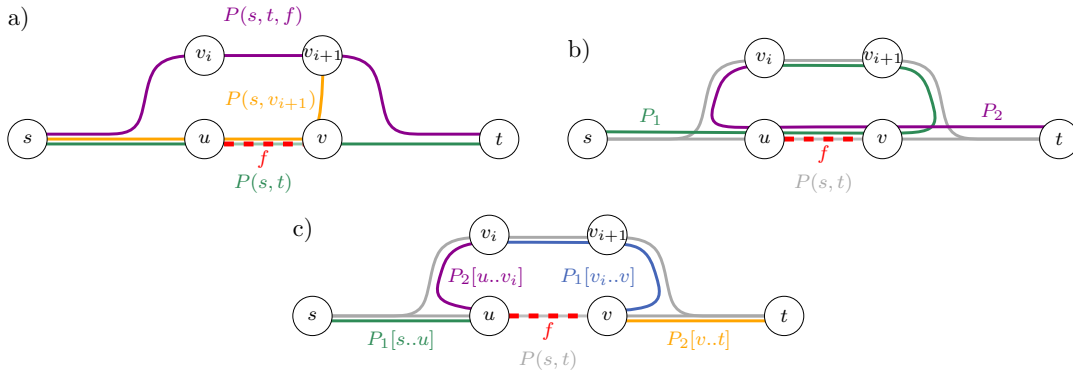
Figure 1: Visualization of the proof of Lemma 2. **a)** A shortest path $P(s,t)$ from s to t (green) and a replacement path $P(s,t,f)$ (purple) that is a shortest path from s to t that avoids the failing edge f (red). Vertex v_i is the last along $P(s,t,f)$ such that the subpath of $P(s,t,f)$ from s to v_i is a shortest path in G, but the subpath of $P(s,t,f)$ from s to v_{i+1} is not a shortest path in G. Thus, the shortest path $P(s,v_{i+1})$ (orange) from s to v_{i+1} must pass through f. **b)** Path P_1 (green) is a concatenation of the shortest path $P(s,v_{i+1})$ from s to v_{i+1} and the edge (v_{i+1},v_i), and P_2 (purple) is a shortest path from v_i to t. Path $P(s,v_{i+1})$ passes through f and thus P_1 does as well. We assume by contradiction that P_2 also passes through f. **c)** Paths $P_1[s..u]$ (green) is the subpath of P_1 from s to u, $P_2[u..v_i]$ (purple) is the subpath of P_2 from u to v_i, $P_1[v_i..v]$ (blue) is the subpath of P_1 from v_i to v and $P_2[v..t]$ (orange) is the subpath of P_2 from v to t. If we concatenate these paths, we get an alternative path from s to t that avoids f and is shorter than the shortest path $P(s,t,f)$ that avoids f, which is a contradiction.

$P(s,t,f)$ from s to v_i is also a shortest path, it follows that the subpath of $P(s,t,f)$ from s to v_{i+1} is not a shortest path in G, and therefore the shortest path $P(s,v_{i+1})$ from s to v_{i+1} must contain the failing edge f. As P_1 contains $P(s,v_{i+1})$ then P_1 contains the failing edge f.

- The path P_2 contains f as this is what we assumed by contradiction above.

Let u,v be the endpoints of f such that u appears before v along $P(s,t)$. It follows that u is closer to s than v, or in other words, that $d(s,u) < d(s,v)$. Similarly, it follows that v is closer to t than u, or in other words, that $d(v,t) < d(u,t)$. Therefore, it must be that u appears before v along the shortest path P_1 (as u is closer to s than v), and it must be that u appears before v along the shortest path P_2 (as v is closer to t than u). Let $P_1[s..u]$ be the subpath of P_1 from s to u, $P_1[v_i..v]$ be the subpath of P_1 from v to v_i, $P_2[u..v_i]$ be the subpath of P_2 from v_i to u, $P_1[v..t]$ be the subpath of P_2 from v to t. Finally, let P be the concatenation of $P_1[s..u]$, $P_2[u..v_i]$, $P_1[v_i..v]$, $P_1[v..t]$, we get that P is a path from s to t that avoids f. Furthermore, P contains all the edges of P_1 and P_2 except of two copies of the edge f, and thus $|P| = |P_1| + |P_2| - 2 < k$, where the last inequality follows from Equations 1 and 2. We conclude that any shortest path of G from v_i to t does contain f, and therefore both $\langle v_0, \ldots, v_i \rangle$ and $\langle v_i, \ldots, v_k \rangle$ are shortest paths in the original graph G. □

Given (s,t,f), let $P(s,t,f)$ be a shortest path from s to t in $G - \{f\}$. According to Lemma 2 it follows that $P(s,t,f)$ is a concatenation of two original shortest paths, or in other words, there exists a pivot vertex v such that $P(s,t,f)$ is the concatenation of the two shortest paths $P(s,v)$ and $P(v,t)$, where $P(s,v)$ is a shortest path from s to v in G and $P(v,t)$ is a shortest path from v to t in G.

As mentioned above, in the remainder of this paper we show that it is possible to use deep learning to find such pivot nodes.

Our Approach

Pivot Selection

By the above argumentation, we can now reduce the problem of finding a replacement path to finding pivot candidates. Our solution framework has two stages. First, we use a graph convolutional network to find vertex representations that encode relevant graph information. Next, we select pivots by formulating the problem as a multi-label classification task and apply a multi-layer perceptron.

We match the performance of state-of-the-art shortest path works with less training cases (numerically and proportionally) and no special selection process, yielding a simple and effective pipeline, which is depicted in Figure 2. We train our model using BCELoss such that we learn both representations and MLP weights towards the pivot selection task. The final output of our network is then a vector representing the likelihood of each node being a pivot node. After sorting the node candidates in non-ascending likelihood, we check if the path they represent (the precomputed shortest path from s to candidate c appended with the shortest path from c to t) contains the node failure. The first candidate for which this is not the case is selected as our pivot node.

Training Data

We train the network on a dataset containing potential start, target, failing, and pivot nodes. The input of our network is a triplet of indices for the start, target, and failing nodes. The output is an n-dimensional vector, where the ith entry represents the likelihood of node i being a pivot node. Training data was then computed by randomly sampling
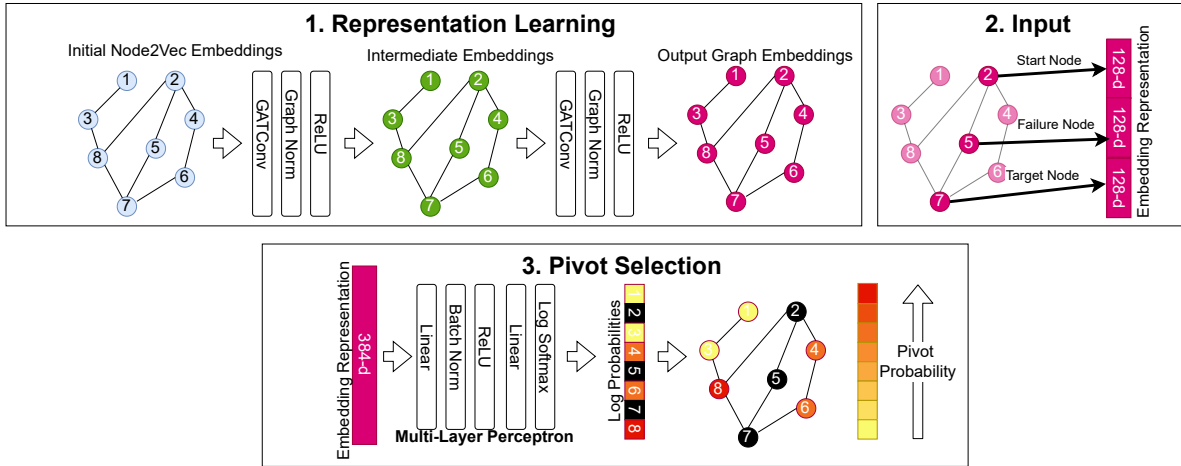
Figure 2: The overall neural network architecture. Our initial node2vec embeddings are processed through two GAT Layers to produce vector embeddings for each node. For each query or training sample, the vectors corresponding to the starting, target, and failure nodes are concatenated. These are processed through two Linear Layers to produce the log-likelihoods of each node being a pivot. During back-propagation, the gradient is passed through the MLP and to the relevant GCN parameters, so that the representations learned to encode relevant task-specific attributes.

$\min(100000, n^2)$ input triplets for a considered graph with n nodes. For each sample, we chose start, target, and failing node at random and ensured that the start and target nodes are connected but are not direct neighbors. Otherwise, the sample was rejected. We then simulated the node failure, computed replacement paths, and determined the vertices on them that fulfill the criteria for being pivot nodes, as described above. This list of pivots was encoded in a binary vector, representing our output. We split our dataset by 80/10/10 for our training, validation, and testing sets. Previous works use landmark labelling to improve the quality of their candidates (Huang et al. 2021). Thus, we have a comparatively simple selection process and a small sample size to learn from. However, we show in Section that we can still identify obtain paths with minimal MRE.

Representation Learning with Graph Convolutional Networks

In order to learn representations from, not only the shortest path data but also the network structure, we leverage graph convolutional networks (GCNs). Inspired by the successes of convolutional neural networks in computer vision, GCNs extend 2D convolutions by updating node information using their neighbors. These models utilize node attributes and the graph structure to update the current node by aggregating its neighbors' features with its own. By stacking many graph convolutional layers together, GCNs can obtain high-level feature representations that incorporate information from across the graph (Zhang et al. 2019). The specific layer we use is the Graph Attention (GAT) layer. This attention-based convolutional layer computes each nodes' embeddings by implicitly learning different importance to different nodes within a neighborhood (Veličković et al. 2017). These have

been used successfully for tasks like biological link prediction, Android malware detection, and text classification (Pham and Dang 2021; Hei et al. 2021; Linmei et al. 2019). To the best of our knowledge, we are the first to apply the GAT mechanism to shortest path problems.

We initialize our node embeddings using node2vec, a general embedding technique which uses deep learning and random walks to flexibly explore and encode neighborhoods (Grover and Leskovec 2016). As mentioned above, this technique has been applied successfully for recommender systems, community detection, and, most relevantly, the shortest path task (Chen et al. 2019; Hu et al. 2020; Rizi, Schloetterer, and Granitzer 2018). Our representation learning framework uses node2vec to generate the initial embeddings and two GAT layers to refine them to our task. We introduce non-linearities using a ReLU layer, and we speed up convergence using GraphNorm layers, a normalization technique that was recently found to be more effective than similar techniques for training graph neural networks (Cai et al. 2021). We also use dropout (0.1) to deal with overfitting (Srivastava et al. 2014).

Multi-Layer Perceptron

We frame the task of selecting our pivot embeddings as a multi-label classification task. As shown in Figure 2, after the GCN component of our model, we select the three embeddings corresponding the start, target, and failing node triplet. These are flattened into a 384-sized vector, which serves as the input for our feedforward neural network.

The multilayer perceptron consists of two linear layers and an output layer. We use ReLU layers as an activation function between the first two layers to introduce non-linearities (Nair and Hinton 2010). We also use BatchNorm

layers to enable faster convergence and introduce regularization (Ioffe and Szegedy 2015). The output layer is a log-softmax layer (commonly used for multi-label classification) that normalizes the vector values such that they represent the log probabilities of being a pivot node.

Experiments

Dataset

We test our method on a variety of networks from the Network Repository (Rossi and Ahmed 2015), which covers a diverse set of areas, such as road networks, biological networks, and communication networks, all representing potential input for real-world applications. For a list of all considered networks, we refer to Table 1.

Evaluation Metrics

In line with previous approaches to computing shortest paths using deep learning, we evaluate our method using the **Mean Relative Error (MRE)** metric. Let $\bar{d}_f(u, v)$ denote the distance between u and v in $G - \{f\}$. After defining the stretch $\alpha(G, s, t, f)$ as

$$\alpha(G, s, t, f) = \frac{\bar{d}_f(s, p) + \bar{d}_f(p, t)}{\bar{d}_f(s, t)},$$

which denotes the factor by how much the replacement is longer than the shortest path, the MRE is simply given by

$$\mathrm{MRE}(G, s, t, f, p) = \alpha(G, s, t, f) - 1. \qquad (3)$$

We note that, as is commonly done, we report the MRE in percentage points throughout the paper.

Reproductibility

To generate our initial embeddings on each real-world network, we train node2vec for 16 iterations using walk length = 10, num walks = 80, $p = 1$, $q = 1$, and window size = 10 (Grover and Leskovec 2016). Embeddings are all 128-dimensional, in line with node2vec and previous shortest paths works (Grover and Leskovec 2016) (Qi et al. 2020) (Rizi, Schloetterer, and Granitzer 2018).

We train our neural network for 32 epochs using the Adam optimizer (Kingma and Ba 2014). Our learning rate is 1e-3, our betas are 0.9 and 0.999, and our epsilon value is 1e-8.

We implemented everything in Python and run the experiments on a Quadro RTX 8000 and an Intel(R) Xeon(R) Silver 4214R CPU @ 2.40GHz.

Results

In order to evaluate whether our approach can be used to obtain suitable pivot vertices, and with them workable replacement paths, we wanted to answer the following questions using our experimental data:

1. How much longer than shortest paths are replacement paths based on pivots that are found using deep learning?

2. Is the resulting performance an achievement of our approach or merely an artifact of the structure of the input graph?

The motivation behind the first question is obvious. Computed replacement paths are only suitable if they are not much longer than the actual shortest paths in the graph. Consequently, it is important that the pivots determined using our approach represent replacement paths with a small stretch.

The second question is aimed at the validity of our approach in the sense that computed outputs are good because the method is and not because any output would have been. For example, consider a graph that is almost a clique (almost all vertices are pairwise connected). Then, all paths are short and after a failure (having almost no impact on the graph structure) most replacement paths are short as well. In this setting almost any node can serve as a suitable pivot, yielding a replacement path with a small stretch and one would expect that even randomly chosen pivots would yield good results. While the networks considered in our experiments are not as dense (see Table 1), other graph properties like a small diameter may make finding good pivots easier.

Regarding the first question, Table 1 lists the MRE values obtained on all considered networks. As can be clearly seen, most stretches are very close to the optimal value. Except for **bio-grid-yeast**, **ia-wiki-Talk**, and **tech-RL-caida** all values are below 1%. This means that the computed pivots represented replacement paths that are almost as short as the shortest paths in the graph, showing that DSOs enhanced via deep learning exhibit very good performance.

Regarding the second question, Table 1 also lists the MRE values we get when considering random pivots, which we obtained by replacing the output of our pipeline with random noise. As can be clearly seen, the MRE is much larger in this setting. For most networks the MRE is larger than 200%, which corresponds to a stretch of more than 3, meaning the found paths are more than 3 times longer than the shortest replacement paths. In order to compare our method with the random approach, Table 1 also lists the improvement factor, denoting the ratio of the MRE obtained with random pivots and the one obtained using our method. Except for **bio-grid-yeast**, **ia-wiki-Talk**, and **tech-RL-caida** this factor is always larger than 200, meaning on most networks our approach is over 200 times better than the random method, clearly indicating that the close to optimal performance is due to the quality of our approach and not an artifact of properties of the considered inputs.

We would like to mention another motivation for comparing our method with random pivot selection. In Li, Chen, and Koltun (2018), the authors used a combination of tree search and a GCN to find maximum independent sets in graphs. However, recent work by Böther et al. (2022) found that the GCN was not learning meaningful representations of the graphs, as results of similar quality were obtained by replacing the GCN outputs with random noise. Thus, we wanted to validate our results and measure the quality of our representation learning by comparing our method with the case that the GCN output is replaced by random vectors.

As identified above, the best performing shortest path using deep learning work achieves an MRE below 0.7% on road networks (Huang et al. 2021). We perform similarly a mean MRE of 1.54% a median MRE of 0.32% across

Network Name	Nodes	Density	Average Degree	Dataset Size	MRE (Random Pivots)	MRE (Our Method)	Improvement
chem-ENZYMES-g118	95	2.71E-02	2.547	9.03E+03	215.96%	0.27%	811.89
chem-ENZYMES-g296	125	1.82E-02	2.256	1.56E+04	257.69%	0.49%	530.22
infect-dublin	410	3.30E-02	13.488	1.00E+05	191.02%	0.18%	1,091.55
bio-celegans	453	1.98E-02	8.940	1.00E+05	176.28%	0.04%	4,299.51
bn-mouse-kasthuri-graph-v4	987	3.16E-03	3.112	1.00E+05	204.20%	0.04%	5,105.11
can_1072	1,072	1.18E-02	12.608	1.00E+05	211.58%	0.37%	573.40
scc_retweet	1,150	9.98E-02	114.713	1.00E+05	167.64%	0.05%	3,287.11
power-bcspwr09	1,723	2.78E-03	4.779	1.00E+05	237.46%	0.16%	1,493.46
inf-openflights	2,905	3.71E-03	10.771	1.00E+05	202.05%	0.72%	280.63
inf-power	4,941	5.40E-04	2.669	1.00E+05	228.67%	0.23%	1,013.10
ca-Erdos992	4,991	5.97E-04	2.977	1.00E+05	206.26%	0.33%	630.77
power-bcspwr10	5,300	9.66E-04	5.121	1.00E+05	232.83%	0.32%	739.13
bio-grid-yeast	6,008	8.70E-03	52.245	1.00E+05	173.80%	6.25%	27.80
soc-gplus	23,613	1.41E-04	3.319	1.00E+05	200.20%	0.31%	654.61
ia-email-EU	32,430	1.03E-04	3.355	1.00E+05	202.93%	0.94%	216.11
ia-wiki-Talk	92,117	8.50E-05	7.833	1.00E+05	203.29%	7.67%	26.50
tech-RL-caida	190,914	3.33E-05	6.365	1.00E+05	206.02%	7.91%	26.04

Table 1: Results of proposed deep learning DSO for node failures across several real-life networks.

a wide range of real-world networks. In addition to outperforming other state of the art shortest paths works, our approach increases the functionality in two key ways (Yin, Rao, and Zhang 2021) (Qi et al. 2020) (Rizi, Schloetterer, and Granitzer 2018). First, and most obviously, we accommodate node failures without recomputing shortest paths again. Second, using our representation and pivot selector learning, our method is able to achieve near-exact approximations of stretch by naively sampling the dataset. The previous top performing shortest path work chose candidates for its dataset using a complex hierarchical learning strategy (Huang et al. 2021). Thus our method is able to extract signals from noise more efficiently.

Conclusion & Outlook

We have shown that distance sensitivity oracles with close to optimal performance can be obtained by utilizing the power of deep learning. Our method builds on a combinatorial property that allows for finding replacement paths based on pivot vertices. On a variety of real-world networks in the presence of failures, we can reliably find suitable pivots where the lengths of the corresponding replacement paths are very close to those of optimal paths. Moreover, our experiments suggest that these results are not artifacts of the inherent structure of the inputs, but are instead based on the fact that the different building blocks of our pipeline successfully capture the relevant structural information about the input graph.

As a consequence, it would be interesting to apply this method to related tasks where similar structural information needs to be captured. One such example is the approximation of shortest paths without failures, as was considered before (Rizi, Schloetterer, and Granitzer 2018). Another example is *local routing*, where the goal is to find short paths in a graph without the use of a central data structure. One way to perform local routing is to compute an embedding of the graph in a metric space and to find paths by greedily routing to vertices that are closer to the target with respect to the metric. Prior work has shown that close to optimal greedy routing can be performed when embedding networks into hyperbolic space (Bläsius et al. 2020). However, the resulting embeddings were susceptible to numerical inaccuracies, and network failures decreased routing performance a lot. It would thus be interesting to see whether our approach can be extended to the greedy routing setting as well, in order to overcome the previously observed issues.

Apart from extending our method to further applications, the approach itself can be extended as follows. As mentioned above, we do not use a landmark-labeling strategy for collecting our training data set. Pursuing this option may improve the quality of the learned representations. We may also pursue different strategies towards the pivot selection task. Customizing the loss function to be multi-task, for example classifying pivots as we do now but also explicitly minimizing stretch, may improve the model's performance as well. Additionally, we may experiment with recently developed GCN layers, to see if they might better exploit the structures found in the challenging networks of Table 1. We may also extend our work to answer queries with edge failures.

Finally, our approach has currently not been tested on larger networks containing millions of nodes. By calculating the APSP information using a distance oracle and using an improved node2vec implementation, we plan to test our networks' scalability in the future.

References

Afek, Y.; Bremler-Barr, A.; Kaplan, H.; Cohen, E.; and Merritt, M. 2002. Restoration by path concatenation: fast recovery of MPLS paths. *Distributed Comput.*, 15(4): 273–283.

Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks: membership, growth, and evolution. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 44–54.

Baswana, S.; and Khanna, N. 2013. Approximate Shortest Paths Avoiding a Failed Vertex: Near Optimal Data Structures for Undirected Unweighted Graphs. *Algorithmica*, 66: 18–50.

Bernstein, A.; and Karger, D. R. 2008. Improved distance sensitivity oracles via random sampling. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, 34–43. SIAM.

Bernstein, A.; and Karger, D. R. 2009. A nearly optimal oracle for avoiding failed vertices and edges. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, 101–110. ACM.

Billand, P.; Bravard, C.; Iyengar, S. S.; Kumar, R.; and Sarangi, S. 2016. Network connectivity under node failure. *Economics Letters*, 149: 164–167.

Bläsius, T.; Friedrich, T.; Katzmann, M.; and Krohmer, A. 2020. Hyperbolic Embeddings for Near-Optimal Greedy Routing. *ACM J. Exp. Algorithmics*, 25.

Böther, M.; Kißig, O.; Taraz, M.; Cohen, S.; Seidel, K.; and Friedrich, T. 2022. What's Wrong with Deep Learning in Tree Search for Combinatorial Optimization. In *International Conference on Learning Representations*.

Cai, H.; Zheng, V. W.; and Chang, K. C.-C. 2018. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 30(9): 1616–1637.

Cai, T.; Luo, S.; Xu, K.; He, D.; Liu, T.-y.; and Wang, L. 2021. Graphnorm: A principled approach to accelerating graph neural network training. In *International Conference on Machine Learning*, 1204–1215. PMLR.

Chechik, S.; and Cohen, S. 2020. Distance sensitivity oracles with subcubic preprocessing time and fast query time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, 1375–1388. ACM.

Chechik, S.; Langberg, M.; Peleg, D.; and Roditty, L. 2012. ƒ-Sensitivity Distance Oracles and Routing Schemes. *Algorithmica*, 63: 861–882.

Chen, J.; Wu, Y.; Fan, L.; Lin, X.; Zheng, H.; Yu, S.; and Xuan, Q. 2019. N2vscdnnr: A local recommender system based on node2vec and rich information network. *IEEE Transactions on Computational Social Systems*, 6(3): 456–466.

Costa, M.; Castro, M.; Rowstron, R.; and Key, P. 2004. PIC: Practical Internet coordinates for distance estimation. In *24th International Conference on Distributed Computing Systems, 2004. Proceedings.*, 178–187. IEEE.

Crichton, G.; Guo, Y.; Pyysalo, S.; and Korhonen, A. 2018. Neural networks for link prediction in realistic biomedical graphs: a multi-dimensional evaluation of graph embedding-based approaches. *BMC bioinformatics*, 19(1): 1–11.

Cvetkovski, A.; and Crovella, M. 2009. Hyperbolic embedding and routing for dynamic graphs. In *IEEE INFOCOM 2009*, 1647–1655. IEEE.

Demetrescu, C.; Thorup, M.; Chowdhury, R. A.; and Ramachandran, V. 2008. Oracles for Distances Avoiding a Failed Node or Link. *SIAM J. Comput.*, 37(5): 1299–1318.

Duan, R.; and Zhang, T. 2017. Improved Distance Sensitivity Oracles via Tree Partitioning. In *Algorithms and Data Structures - 15th International Symposium, WADS 2017, St. John's, NL, Canada, July 31 - August 2, 2017, Proceedings*, volume 10389 of *Lecture Notes in Computer Science*, 349–360. Springer.

Grandoni, F.; and Williams, V. V. 2020. Faster Replacement Paths and Distance Sensitivity Oracles. *ACM Trans. Algorithms*, 16(1): 15:1–15:25.

Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864.

Gu, Y.; and Ren, H. 2021. Constructing a Distance Sensitivity Oracle in $O(n^2.5794M)$ Time. In Bansal, N.; Merelli, E.; and Worrell, J., eds., *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 76:1–76:20. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-195-5.

Hei, Y.; Yang, R.; Peng, H.; Wang, L.; Xu, X.; Liu, J.; Liu, H.; Xu, J.; and Sun, L. 2021. Hawk: Rapid android malware detection through heterogeneous graph attention networks. *IEEE Transactions on Neural Networks and Learning Systems*.

Hu, F.; Liu, J.; Li, L.; and Liang, J. 2020. Community detection in complex networks using Node2vec with spectral clustering. *Physica A: Statistical Mechanics and its Applications*, 545: 123633.

Huang, S.; Wang, Y.; Zhao, T.; and Li, G. 2021. A Learning-based Method for Computing Shortest Path Distances on Road Networks. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 360–371. IEEE.

Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456. PMLR.

Jindal, I.; Chen, X.; Nokleby, M.; Ye, J.; et al. 2017. A unified neural network approach for estimating travel time and distance for a taxi trip. *arXiv preprint arXiv:1710.04350*.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

Li, Z.; Chen, Q.; and Koltun, V. 2018. Combinatorial Optimization with Graph Convolutional Networks and Guided Tree Search. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 537–546.

Lincoln, A.; Williams, V. V.; and Williams, R. R. 2018. Tight Hardness for Shortest Cycles and Paths in Sparse Graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, 1236–1252. SIAM.

Linmei, H.; Yang, T.; Shi, C.; Ji, H.; and Li, X. 2019. Heterogeneous graph attention networks for semi-supervised short text classification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4821–4830.

Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Icml*.

Nelder, J. A.; and Mead, R. 1965. A simplex method for function minimization. *The computer journal*, 7(4): 308–313.

Pham, C.; and Dang, T. 2021. Link prediction for biomedical network. In *The 12th international conference on advances in information technology*, 1–5.

Qi, J.; Wang, W.; Zhang, R.; and Zhao, Z. 2020. A Learning Based Approach to Predict Shortest-Path Distances. In *EDBT*, 367–370.

Ren, H. 2022. Improved distance sensitivity oracles with subcubic preprocessing time. *J. Comput. Syst. Sci.*, 123: 159–170.

Ren, Y.; Ay, A.; and Kahveci, T. 2018. Shortest path counting in probabilistic biological networks. *BMC bioinformatics*, 19(1): 1–19.

Rizi, F. S.; Schloetterer, J.; and Granitzer, M. 2018. Shortest path distance approximation using deep learning techniques. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 1007–1014. IEEE.

Rossi, R. A.; and Ahmed, N. K. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958.

Tian, X.; Song, Y.; Wang, X.; and Gong, X. 2012. Shortest path based potential common friend recommendation in social networks. In *2012 Second International Conference on Cloud and Green Computing*, 541–548. IEEE.

Ukkonen, A.; Castillo, C.; Donato, D.; and Gionis, A. 2008. Searching the wikipedia with contextual information. In *Proceedings of the 17th ACM conference on Information and knowledge management*, 1351–1352.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Wang, H.-S.; Zhu, X.; Peh, L.-S.; and Malik, S. 2002. Orion: A power-performance simulator for interconnection networks. In *35th Annual IEEE/ACM International Symposium on Microarchitecture, 2002.(MICRO-35). Proceedings.*, 294–305. IEEE.

Wang, X.; Cui, P.; Wang, J.; Pei, J.; Zhu, W.; and Yang, S. 2017. Community preserving network embedding. In *Thirty-first AAAI conference on artificial intelligence*.

Weimann, O.; and Yuster, R. 2013. Replacement Paths and Distance Sensitivity Oracles via Fast Matrix Multiplication. *ACM Trans. Algorithms*, 9(2): 14:1–14:13.

Yin, J.; Rao, W.; and Zhang, C. 2021. Learning shortest paths on large dynamic graphs. In *2021 22nd IEEE International Conference on Mobile Data Management (MDM)*, 201–208. IEEE.

Zhang, S.; Tong, H.; Xu, J.; and Maciejewski, R. 2019. Graph convolutional networks: a comprehensive review. *Computational Social Networks*, 6(1): 1–23.

Zhang, X.; Mu, J.; Liu, H.; and Zhang, X. 2021. Graphnet: Graph Clustering with Deep Neural Networks. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3800–3804. IEEE.

Zhao, X.; Sala, A.; Zheng, H.; and Zhao, B. Y. 2011. Efficient shortest paths on massive social graphs. In *7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, 77–86. IEEE.