Finding MNEMON: Reviving Memories of Node Embeddings

Yun Shen 1* Yufei Han 2 Zhikun Zhang 3 Min Chen 3 Ting Yu 4 Michael Backes 3 Yang Zhang 3 Gianluca Stringhini 5

NetApp ²INRIA ³CISPA Helmholtz Center for Information Security ⁴QCRI ⁵Boston University

Abstract

Previous security research efforts orbiting around graphs have been exclusively focusing on either (de-)anonymizing the graphs or understanding the security and privacy issues of graph neural networks. Little attention has been paid to understand the privacy risks of integrating the output from graph embedding models (e.g., node embeddings) with complex downstream machine learning pipelines. In this paper, we fill this gap and propose a novel model-agnostic graph recovery attack that exploits the implicit graph structural information preserved in the embeddings of graph nodes. We show that an adversary can recover edges with decent accuracy by only gaining access to the node embedding matrix of the original graph without interactions with the node embedding models. We demonstrate the effectiveness and applicability of our graph recovery attack through extensive experiments.

1 Introduction

Many complex systems can be represented as graphs, such as social networks, communication networks, function call graphs, biomedical graphs, and the World Wide Web [35, 42, 56]. Graph embedding algorithms [6, 19, 77] have been long researched to obtain effective graph representations to represent these networks concisely in low dimensional Euclidean vectors. Upon such transformation, these embedding vectors can make graph analytics tasks efficient and facilitate numerous solutions to real world problems, e.g., node classification [69], community detection [48], link prediction/recommendation [43], binary similarity detection [18, 82, 85], malware detection [16, 52], fraud detection [70], and bot detection [2].

It is well recognized that graphs contain sensitive and private information about the nodes (e.g., node attributes, the relationships among the nodes, etc.). Previous security research efforts orbiting around graphs have been focusing on either (de-)anonymizing the graphs [31, 45, 84] or understanding the security and privacy issues of graph neural networks [26, 62, 66, 71, 72, 80, 81]. Specifically, graph anonymization techniques [31, 45, 84] perturb the original

graph data to protect users' privacy while preserving as much data utility as possible. In contrast, graph de-anonymization techniques focus on unveiling sensitive private information from graphs. In recent years, inspired by the membership inference attack [63], we have witnessed several successful link re-identification attacks against graph neural networks that extract private links contained in the training data via these GNN models [26, 71, 78, 81]. Note that the node embeddings are not privacy preserving by design. Yet, they are pervasively used in many graph analytics tasks as aforementioned. To our surprise, understanding and quantifying the privacy risks of integrating them with the complex ML pipeline vin a model-agnostic setting remains unexplored, hence our focus in this paper.

As such, we fill this gap and quantify the privacy risks of integrating node embeddings with downstream data analytics/machine learning pipelines. Our attack's application scenarios (see Section 3.2) lie in the complex ML systems where raw graph data is part of the learning process but cannot be directly obtained by the attackers due to data segregation policy and/or privacy policy. Instead, the attackers only gain access to the transformed graph data (i.e., the node embeddings of the original graph). They cannot interact with the node embedding models since such pipelines usually operate in one direction. For instance, the data holder may have integrated with the malicious machine learning solution providers (i.e., MLaaS providers) from the AWS Marketplace [47,65], or the data holder is part of a vertical federated learning environment in an enterprise [71]. In both cases, the node embeddings are part of the learning process and can be obtained by either the malicious MLaaS providers [47,65] or the insiders [71] in the pipeline.

Concretely, our study addresses two research questions - can we recover the edges with decent accuracy from the node embedding matrix and can we recover a graph structure that is similar to the original graph with respect to the graph properties? - without knowledge of and the interactions with the node embedding models. Note that these two research questions were discussed in the link re-identification attacks [14, 26, 71]. They, however, follow the adversarial machine learning methodology and assume the interaction with the target model using shadow datasets and the supervision information from the feedback. Our attack does not

^{*}Work partially done while the author was with NortonLifeLock.

assume such capabilities (see Section 3), which is more practical in the real world.

Our Contributions. In this paper, we propose MNEMON - a joint graph metric learning and self-supervised learning based graph recovery attack - to tackle these two questions. MNEMON first leverages the background information (i.e., the origin of the node embedding matrix) to estimate the average node degree. It then uses graph metric learning with a multi-head attention mechanism to construct a data specific distance metric from a given node embedding matrix. Coupling with graph metric learning, MNEMON employs graph autoencoder framework to iteratively optimize a graph structure through self-supervised graph regularization (i.e., the learning objectives are generated from the data itself). Upon the termination of the process, the learned graph structure constitutes the recovered graph from the node embedding matrix.

We stress that our goal is not perfectly recovering a graph from its node embedding matrix. Rather, we focus on understanding and quantifying the privacy risks of integrating them with the complex ML pipeline.

A successful graph recovery attack can lead to severe consequences. For instance, in the context of social networks, MNEMON allows an adversary to gain direct knowledge of sensitive and private social relationships. Also, certain graph data is often expensive to obtain (e.g., protein interaction networks collected from lab studies). MNEMON can pose a direct threat to the intellectual property of the data holder as well. In summary, we make the following contributions.

- We propose a novel model-agnostic graph recovery attack that exploits the implicit graph structural information preserved in the node embedding vectors. We show that the attacker can unveil the private and sensitive graph structural information with decent accuracy from the node embeddings.
- We systematically define the threat model to characterize an adversary's background knowledge and realistic application scenarios. Extensive evaluation of four popular node embedding models using four benchmark graph datasets demonstrates the efficacy of our attacks.
- We discuss a preliminary mitigation mechanism to defend against the graph recovery attack. Our results demonstrate that MNEMON could be partially mitigated with some utility trade-off.

2 Preliminaries

2.1 Notations

We denote an undirected, attributed graph as $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{X})$, where $\mathbf{V} = \{v_i\}_{i=1}^n$ represents the nodes, $\mathbf{E} \subseteq \{(v,u)|v,u \in \mathbf{V}\}$ denotes the edges, and $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ denotes the node features, where \mathbf{x}_i represents the node feature of v_i . $|\mathbf{E}|$ denotes the graph size (i.e., the number of edges). The original and the recovered graphs are denoted as \mathbf{G}_O and \mathbf{G}_R respectively. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ represent the (weighted) adjacency matrix. As

Table 1: Summary of the notations. We use lowercase letters to denote scalars, bold lowercase letters to denote vectors and bold uppercase letters to denote matrices.

Notation	Description
G = (V, E, X)	graph (network)
G_O / G_R	original/recovered graph
n	number of nodes
$\mathbf{A} \in \mathbb{R}^{n imes n}$	(weighted) adjacency matrix
X	node features
v, u	node
d	dimension of node embeddings
$\mathbf{H} \in \mathbb{R}^{n \times d}$	node embedding matrix
\mathbf{h}_{v}	node embedding of node v
t	t-th iteration
k	(estimated) average node degree
f	node embedding model
φ	learnable embedding distance function
L	loss function

such, **G** can also be represented as $\mathbf{G} = (\mathbf{A}, \mathbf{X})$. The notations introduced here and used in the following sections are summarized in Table 1.

2.2 Node Embedding

Definition. In this paper, we focus on *node embedding*, which plays a central role in graph embedding techniques. As the name suggests, a node embedding model f maps nodes to d-dimensional vectors that capture their structural properties and node features (if available). Formally, a node embedding model is defined as $f: \mathbf{G} \to \mathbf{H}$, where $\mathbf{H} \in \mathbb{R}^{n \times d}$ represents node embedding matrix where d denotes the dimension of the embeddings ($d \ll n$) and $\mathbf{h}_v \in \mathbf{H}$ denotes the node embedding vector of node v. The node embeddings of connected nodes maintain "approximate closeness" to each other in the latent space (e.g., \mathbf{h}_v and \mathbf{h}_u should be close in the Euclidean space if v and u are connected in the graph).

Overview. There exists abundant previous work on node embedding models [6, 19, 77]. Broadly speaking, these techniques can be grouped into two categories - matrix factorization based approaches and deep learning based approaches.

- Matrix factorization based approaches. The essence of
 these approaches is treating node embedding as a dimensionality reduction problem and factorizing graph adjacency matrix or node proximity/similarity matrix to obtain
 node embedding [32]. The core idea of these approaches is
 that the graph property to be preserved can be interpreted
 as pairwise node similarities or node proximity in a low dimensional space by matrix factorization. In general, matrix
 factorization methods can be classified into two categories
 node proximity matrix factorization and graph Laplacian
 eigenmaps factorization.
- Deep learning (DL) based approaches. The pioneer DL-based approaches include DeepWalk [53], Node2Vec [20], and their variants. These approaches first generate a set of truncated random walk paths sampled from a graph, then apply deep learning techniques (e.g., SkipGram) to the sampled paths, consequently learning node embeddings. In recent years, we also witnessed the rise of Graph Neural

Networks (GNNs). These GNN models are a type of Neural Network which directly operates on the Graph structure via message passing between the nodes of graphs, and encoding the nodes into a low dimensional space (e.g., GCN [36], GraphSAGE [23]). They can take node features into consideration and do not need random walk paths.

We refer the audience to [6, 19, 77] for the overview of node embedding techniques and other graph tasks (e.g., graph-level embedding, graph-level classification, etc.).

3 Threat Model

3.1 Attack Setting

We frame our attack in a *model agnostic* setting. We assume that the adversary only has access to the *node embedding matrix* **H** together with the *background information* of the origin from which the embedding matrix was leaked (see Section 3.2 for detailed application scenario discussion). The attackers do not have any knowledge of the node embedding model, and they cannot tamper with its internals (e.g., model parameters, model architecture). We strictly require that the attackers cannot interact with the target model, and do not have the auxiliary data to train a shadow model using the feedback from the target model.

Remarks. It is important to note that our attack setting is different from the existing adversarial machine learning settings, whereas the interaction with the target model (i.e., querying the target model via publicly accessible API) and the availability of auxiliary data (e.g., nodes with features and labels, etc.) are indispensable. Our attack, however, assumes neither. That is, we strictly require that the attackers cannot interact with the target model, and do not have the auxiliary data to query a target model and use the query results to train a shadow model. In other words, this setting eliminates the supervise information from the target model and consequently renders the previous link re-identification attacks inapplicable, hence the novelty of our attack. We provide a detailed discussion in Section 7 to distinguish our attack from the existing ones.

3.2 Attack Scenarios

We consider our attack's application scenarios lie in those complex ML systems where graph data is part of the learning process but cannot be directly obtained by the attackers due to data segregation policy and/or privacy policy. As such, we discuss three real world scenarios below.

• The first attack scenario is the insider threat in a complex enterprise ML environment. In this scenario, a company enforces rigid data protection and segregation policies to guard the security of raw data. As a result, one department may have sensitive user private profile and relationship information (i.e., graph data with node features), and another department has the user purchase history. To train a joint model (e.g., a personalized recommender system) that leverages the data from different departments, the company needs to perform vertical federated learning [75]. In-

stead of supplying the graph data to the central model, the department that holds the graph data may generate node embeddings that preserve the utility (i.e., user closeness without disclosing the exact edges) and facilitate the learning task. The insider then obtains the node embeddings during this learning process and leaks them to the attackers. This attack scenario is in line with the setting recently discussed by Wu et al. [71].

- The second attack scenario is the malicious third-party provider that is already part of the data holder's data analytics or machine learning pipeline. For example, the data holder may have integrated with the malicious machine learning solution providers (i.e., MLaas providers) from the AWS Marketplace. In this case, the upstream data holder, without knowing the implications, passes the node embedding matrix to the rouge provider for downstream analytical tasks, such as data visualization, link prediction, node classification, profiling, etc. The attackers can then obtain the node embedding matrix from the data holder through the rouge provider. This attack scenario is in line with the malicious machine learning provider scenario discussed by Song et al. [65] and Malekzadeh et al. [47].
- The third attack scenario is security misconfiguration in the ML environment. For instance, researchers may leverage the free computing resources (e.g., GPUs) offered by Colab, and connect it to their private Github repository. Due to such misconfigurations, the notebooks containing the node embeddings are leaked (i.e., wrongly using "anyone on the Internet with this link can view" instead of "send to the specific users"). This attack scenario is in line with the real world misconfigured S3 buckets leakage discussed by Continella et al. [12].

Background Information Acquisition. Besides, given the first two attack scenarios, the attackers can easily obtain the background information of the origin of the embedding matrix (e.g., from which companies the matrices come from). With fair reconnaissance efforts (e.g., correlating the owner of Colab notebooks with Github handles), the attacker may also infer the origin of the embedding matrix in the third scenario. In summary, these three attack scenarios are tangible and match our attack setting.

3.3 Attack Goals

The primary goal of the attackers is uncovering the edges with decent accuracy from the node embedding matrix. Attaining this goal would enable the attacker to expose private and sensitive relationships among the nodes rather than the "approximate closeness" offered by the node embeddings (see Section 2). Nevertheless, due to the strict attack setting, it is impractical for the attackers to faultlessly retrieve all the edges from the node embedding matrix. As a result, the secondary goal of the attackers is recovering a graph structure \mathbf{A}_R that is similar to the original graph \mathbf{A}_O with respect to the graph properties. Achieving this goal would enable the attackers to gain additional knowledge of the original graph as

a whole and perform graph mining tasks, which in turn violates the intellectual property of the data holder or can facilitate advanced attacks, such as re-identifying individuals [31], structural data de-anonymization [30], etc. For example, recovering a graph with similar triangle counts and joint degree distribution to the original graph would enable the attacker to gain insights into the underlying user engagement in a social network. This information itself is sensitive and proprietary.

Non-goals. Recall our attack setting in Section 3.1 that the attackers only have the node embedding matrix and the background information of the origin of the embedding matrix, and cannot interact with the target model with auxiliary data. We thereby cannot infer node features (i.e., attribute inference attack) since we do not have any auxiliary data (i.e., we do not know the format of the original node features). Similarly, we cannot steal the target model (i.e., model extraction attack) nor can we understand the privacy leakage from the target model itself as we do not interact with it. Finally, our attack focuses on the node-level embeddings. We thus do not attack the graph-level embeddings [6, 19, 77].

4 MNEMON: Graph Recovery Attack

4.1 Attack Overview

At a high level, MNEMON contains three main components.

- The first component (see Section 4.2) leverages the background information (i.e., the origin of the node embedding matrix) to estimate the average node degree. The goal is to estimate a rough average node degree k and the graph size (i.e., $|\mathbf{E}| = \frac{k \times n}{2}$).
- The second component (see Section 4.3) leverages graph metric learning (GML) to learn a data-specific distance function since it is often difficult to choose a standard metric that fits all the datasets. The goal is to learn multi-head attention weights and tailor the distance function on a per node embedding matrix basis.
- The third component (see Section 4.4) learns a graph structure through Graph AutoEncoder (GAE) framework using self supervised graph regularization. The goal is to optimize the graph structure and reduce the false positive edges incurred by the learned metric from the second component.

We iteratively optimize the second and third components as they are inter-connected.

Specifically, GML learns a distance function to measure the closeness of two nodes and builds the input graph for GAE (T = t in Figure 1). GAE then learns to reconstruct this input graph. If GAE finds certain parts of the input graph are hard to reconstruct, which is reflected by the self supervised graph learning loss, it may be due to the input graph built by GML partially capturing the graph structure. We then merge the graph structures by combining both the input graph and output graph of GAE, which enables us to retain the most confident edges (the transition from T = t to T = t + 1 in




Figure 1: Overview of MNEMON. At timestamp T=0, MNEMON estimates the average node degree k and initializes the seed graph using Gumbel-Top-k trick. At timestamp T=t, it iteratively learns a data-specific graph distance metric using GML and optimizes a graph structure in a self-supervised way using GAE.

Figure 1). The combined graph is then used to guide GML to update its metric learning process in the next iteration. We outline MNEMON's workflow in Figure 1. In the following sections, we discuss the technical details of our attack.

4.2 Estimate the Average Node Degree

The only clue that the attackers have is the background information about the origin of the node embedding matrix. For instance, in our first attack scenario where the node embedding matrix is leaked by an insider, it is straightforward for the attackers to obtain such background information. The attacker's immediate task is thereby estimating the average node degree k. The rationale is straightforward. The attackers already know the number of nodes from the embedding matrix (i.e., n). Yet, due to the combinatorial nature of the graph, there exist $n(n-1)^2$ possible edges. As such, if the attacker can estimate the average node degree k, they can trivially obtain the estimated size of the original graph (i.e., the number of edges) which is equivalent to $(k \times n)/2$. In this way, the estimated graph size enables them to effectively learn the graph structure as we will discuss in Section 4.3 and Section 4.4.

Abundant previous work [13, 24, 25, 38, 41] has already exemplified that the graphs of similar origins may share similar graph properties (e.g., node degree, graph density, small world phenomenon, local clustering coefficient, etc.). Our core idea is that the attackers can estimate the average node degree from the graphs of similar origins and transfer the estimated node degree from these graphs to facilitate the attack. This alleviates the attackers from stealing a sample training data from the data holder, which, in turn, makes our attack realistic. For instance, if the attackers know that the node embeddings come from a Facebook network, they can leverage graph sampling methods to sample Facebook networks publicly available in the Network Data Repository [59] and estimate the average node degree of the network that they target (see Section 5 for how we use graph sampling to sample real world data). These graph sampling methods have been proven accurate in estimating the average node degree [60]. In this paper, we use the state-of-the-art spikyball sampling [41] implemented in the latest Little Ball of Fur python library [60] to estimate the average node degree. Additional details can be found in Section 5.2.

Notes. The graph sampling process does not interact with the original node embedding models. It also does not need the supervision information from the target models as required by previous research [14, 26, 71]. We also stress that MNEMON does not estimate or require the precise average node degree. MNEMON can accommodate the inevitable estimation error. We provide a detailed study in Section 5.5 to illustrate this capability. For instance, we show that our attack can still achieve good performance even when the estimated average node degree is twice the real average node degree (see Section 5.5) thanks to graph metric learning (see Section 4.3) and self-supervised graph structure learning (see Section 4.4).

4.3 Graph Metric Learning

Upon estimating the average node degree, the common approach to recover a graph from the node embedding matrix is using kNN algorithm. kNN builds a graph in which two nodes v and u are connected by an edge if the distance between the embedding vectors h_v and h_u is among the kth smallest distances. The drawback of kNN algorithm is that it requires a manually predefined distance function for neighbor selection. However, it is often difficult to choose a standard metric that fits all the datasets and tasks of interest. Take a barbell graph for example, which consists of two dense cliques connected by a long chain. Reflected in the latent space, the node embedding vectors from two dense cliques are close to each other (i.e., dense regions), while those from the long chain are relatively farther to each other (i.e., sparse regions). A standard distance function, such as Eucliean or cosine distance, used by kNN may not recover the edges from the long chain as the distances among them are inevitably large. Yet, they are equally connected from a graph perspective. As such, we propose to leverage graph metric learning in the node embedding space to learn a dataspecific distance function and automatically adjust for both the dense and sparse regions in the node embedding matrix.

Graph Initialization with Gumbel-Top-*k* **Trick.** We follow the approach discussed by Kazi [34] to initialize a seed graph. We first generate a fully connected graph with edge normalized distance score using Equation 1.

$$p_{vu} = e^{-\tau \delta(h_v, h_u)}, \forall v, u \in \mathbf{V}$$
 (1)

Here δ is a distance function and τ is a temperature parameter controlling the smoothness of the distance scores between node embedding vectors. Instead of using the Euclidean distance function adopted by Kazi [34], we opt for cosine distance as the distance function δ , i.e., $\delta(h_v,h_u)=1-cos(h_v,h_u)$.

Note that most of the node embeddings are normalized to facilitate downstream tasks. In this case, Euclidean distance is proportional to cosine distance, given a well normalized value range of the node embeddings. When node embeddings are not normalized, our framework can also be adjusted to Euclidean distance. Let $\mathbf{P} = \{p_{vu}\}$ denote the edge probability matrix. We then leverage Gumbel-Top-k trick [39] to sample from \mathbf{P} , which generalizes Gumbel-Max trick [22] to draw an ordered sample of size k without replacement from a

categorical distribution by taking the indices of the k largest perturbed log-probabilities.

That is, we perturb each p_{vu} by adding a Gumbel random variate $\vartheta_{vu} \sim Gumbel(0,1)$. We then select the indices of the k largest perturbed log-probabilities without replacement. This process makes the sampling a stochastic relaxation of kNN [34]. This sampled adjacency matrix (denoted as \mathbf{A}^0) constitutes our seed graph structure. Note that δ is used for sampling purposes only and is not part of our learning targets. This corresponds to T=0 in Figure 1.

Learnable Distance Function (ϕ). Due to the stochastic nature of Gumbel-Top-k trick, we inevitably obtain an initial noisy graph structure from the above graph initialization process. That is, an edge (v,u) in \mathbf{A}^0 may not exist in the original graph \mathbf{G}_O , i.e., a false positive edge. To reduce such false positives, we propose a learnable distance function ϕ to learn a better graph structure. The core idea is that, instead of using a predefined distance function, we leverage metric learning [74] to learn a distance metric for the input space of data (i.e., the node embedding matrix \mathbf{H}) from the adjacency matrix \mathbf{A} that preserves the node relationships (i.e., \mathbf{A} is used to supervise the distance learning). In this paper, we adopt a weighted cosine distance (defined in Equation 2) [8, 83] as our learnable distance function ϕ .

$$\phi(\mathbf{h}_{v}, \mathbf{h}_{u}) = 1 - \cos(\mathbf{w} \circ \mathbf{h}_{v}, \mathbf{w} \circ \mathbf{h}_{u}) \tag{2}$$

Here **w** is a learnable weight vector that is the same dimension as \mathbf{h}_{ν} and \mathbf{h}_{u} , and \circ denotes the Hadamard product. Following the procedure discussed in [10,68], we further extend Equation 2 to a multi-head version as in Equation 3 to increase the expressiveness and stablize the learning process.

$$\phi(\mathbf{h}_{v}, \mathbf{h}_{u}) = 1 - \frac{1}{m} \sum_{i=1}^{m} cos(\mathbf{w}_{i} \circ \mathbf{h}_{v}, \mathbf{w}_{i} \circ \mathbf{h}_{u})$$
(3)

Here m refers to the number of attention heads. In this way, we can learn the distance function from multiple perspectives. Note that all node embeddings share the same metric parameters $\mathbf{W} = \{\mathbf{w}_i\}_{i=1}^m$.

Graph Sparsification. We plug ϕ into Equation 1 (i.e., replacing δ) and use the aforementioned Gumbel-Top-k-based sampling trick to extract an adjacency matrix. Our graph sparsification method is different from the ϵ -neighborhood approach used by [10] which cannot easily control the graph size (i.e., ϵ is fixed and may lead to different graph sizes as the learned weighted adjacency matrix also evolves during the learning process).

4.4 Self-Supervised Graph Structure Learning

Having introduced how we apply the graph metric learning technique to tune a data-specific distance function in the previous section, we move on to discuss how we optimize a graph structure and learn the graph distance metric jointly via self-supervised learning. The core idea is that we refine the initial noisy graph structure through self-supervised graph regularization. To this end, we propose to use Graph AutoEncoder [37] (GAE) with an adaptive graph structure

Algorithm 1: Graph recovery attack MNEMON

Input: Node embedding matrix \mathbf{H}^0 , background information B, maximum iteration T, hyperparameters τ , α , β , η , m

Output: Learned graph structure A^T (i.e., G_R)

- 1 $k \leftarrow EstimateAvgDegree(B)$
- 2 A⁰ ← Apply Gumbel-Top-k trick on the fully connected probabilistic graph P with τ (Equation 1) to generate the initial seed graph

```
3 for t \leftarrow 0 to T-1 do

4 | \mathbf{A}^t \leftarrow GML(\mathbf{A}^t, \mathbf{H}^t, m)

5 | \mathbf{A}^{t+1}, \mathbf{H}^{t+1} \leftarrow GAE(\mathbf{A}^t, \mathbf{H}^t)

6 | \mathcal{L} \leftarrow \mathcal{L}_{lap}(\mathbf{A}^{t+1}, \mathbf{H}^0) + \mathcal{L}_{spa}(\mathbf{A}^{t+1}, \alpha, \beta)

7 | + \mathcal{L}_{rec}(\mathbf{A}^{t+1}, \mathbf{A}^t)

8 | Backpropagate \mathcal{L}

9 | \mathbf{A}^{t+1} \leftarrow Combine(\mathbf{A}^0, \mathbf{A}^{t+1}, \eta)

10 | \mathbf{A}^{t+1} \leftarrow Binarize(\mathbf{A}^{t+1})
```

combination mechanism to iteratively refine the graph structure learned from the node embedding matrix.

Graph Autoencoder (GAE). Given the adjacency matrix roughly estimated using the multi-headed distance metric in Eq.3 and the node embedding vectors as the input, GAE learns to refine the adjacency matrix as the output. The initial input of our GAE is $\mathbf{G}^0 = (\mathbf{A}^0, \mathbf{H}^0)$. Here \mathbf{H}^0 represents the node embedding matrix obtained by the attackers. \mathbf{A}^0 represents the initialized seed graph. In this way, we treat \mathbf{H}^0 as the node features \mathbf{X} of \mathbf{G}^0 . Note that we add a superscript for ease of description of the following iterative learning process.

• *Encoder*: The encoder is a *Z*-layer graph convolutional network (GCN) [36]. At the *t*-th iteration, its input is a graph $\mathbf{G}^t = (\mathbf{A}^t, \mathbf{H}^t)$. The encoder (see Equation 4) learns a latent representation $\mathbf{H}^{t+1} \in \mathbb{R}^{n \times d}$ where each row represents a node *v*'s latent representation after encoding.

$$\mathbf{H}^{t+1} = GCN(\mathbf{A}^t, \mathbf{H}^t) \tag{4}$$

• *Decoder.* We use an inner-product decoder in this paper [36]. The adjacency matrix can be reconstructed using Equation 5, where $\sigma(x) = 1/(1+e^{-x})$ and the output \mathbf{A}^{t+1} is a weighted adjacency matrix.

$$\mathbf{A}^{t+1} = \mathbf{\sigma}(\mathbf{H}^{t+1}\mathbf{H}^{t+1^T}) \tag{5}$$

Note that GAE is a generic framework.

We follow the design by Kipf et al. [37] and use GCN as the encoder and inner-product as the decoder. This design allows us to use linear GCN [61] to accelerate the computation and compare to our baseline [14] in Section 5. The adversary can plug in other GNN models into GAE framework. The audience can use different architectures as encoders and decoders.

Self-Supervised Graph Regularization. MNEMON cannot interact with the target model (i.e., the node embedding model). We therefore rely on several graph regularization objectives to guide the above GAE-based learning process in a self-supervised way.

• Graph Laplacian regularization (\$\mathcal{L}_{lap}\$) [4]. A graph Laplacian regularization assumes that the learned weighted adjacency matrix is smooth with respect to a set of node features. In our case, the weighted adjacency matrix is \$\mathbf{A}^{t+1}\$ and the set of node features is the node embedding matrix \$\mathbf{H}^0\$. Note that our goal is to optimize the graph structure (i.e., \$\mathbf{A}^{t+1}\$). As we can see in Equation 6, we stress that we always force that the learned weighted adjacency matrix \$\mathbf{A}^{t+1}\$ is smooth with respect to the initial node embedding matrix \$\mathbf{H}^0\$. As such, graph Laplacian regularization can be interpreted that two connected nodes in the learned graph structure should be close enough in the latent node embedding space defined by \$\mathbf{H}^0\$.

$$\mathcal{L}_{lap}(\mathbf{A}^{t+1}, \mathbf{H}^{0}) = \frac{1}{2n^{2}} \sum_{v,u} \mathbf{A}_{vu}^{t+1} ||h_{v}^{0} - h_{u}^{0}|| = \frac{1}{2} \mathbf{tr}(\mathbf{H}^{0^{T}} \mathbf{L}^{t+1} \mathbf{H}^{0})$$
(6)

where **tr** denotes trace of matrix, $\mathbf{L}^{t+1} = \mathbf{D}^{t+1} - \mathbf{A}^{t+1}$ and $\mathbf{D}^{t+1} = \sum_{v} \mathbf{A}_{vu}^{t+1}$.

• Graph sparsity regularization (\mathcal{L}_{spa}) [33]. In the real world, the graphs are normally sparse. We use graph sparsity regularization proposed by Kalofolias et al. [33] to learn graphs that meet such expectations. As we can see in Equation 7, graph sparsity regularization encourages that each node connects to at least another node in the first term, and penalizes large degrees in the second term naturally arising from the first term. Graph sparsity regularization can be interpreted as using α to force the graph degrees to be positive and β to control the graph sparsity.

$$\mathcal{L}_{spa}(\mathbf{A}^{t+1}, \alpha, \beta) = -\alpha \mathbf{1}^{T} log(\mathbf{A}^{t+1} \mathbf{1}) + \frac{\beta}{2} \|\mathbf{A}^{t+1}\|$$
 (7)

where $\alpha>0$ and $\beta\geq0$ are two controlling hyperparameters.

• Graph reconstruction loss (\mathcal{L}_{rec}) [37]. Graph reconstruction loss forces GAE to learn a latent representation \mathbf{H}^{t+1} to faithfully rebuild the input adjacency matrix \mathbf{A}^t .

In this paper, we adopt the link prediction as the way to interpret the reconstruction loss [44] and minimize the binary cross entropy loss between negative (i.e., non-existing edges) and positive samples (i.e., existing edges). The loss function can be found in Equation 8.

$$\mathcal{L}_{rec} = \frac{1}{2n^2} \| \mathbf{A}^t \circ log(A^{t+1}) + (\mathbf{1} - \mathbf{A}^t) \circ log(\mathbf{1} - \mathbf{A}^{t+1}) \|_F^2$$
(8)

where \circ is elementwise product and **1** is an all-ones matrix.

To summarize, \mathcal{L}_{rec} forces GAE to learn simultaneously the updated latent representation \mathbf{H}^{t+1} and a graph adjacency matrix \mathbf{A}^{t+1} decoded from \mathbf{H}^{t+1} to faithfully rebuild the input adjacency matrix \mathbf{A}^t . \mathcal{L}_{lap} and \mathcal{L}_{spa} makes the learned graph smooth and sparse. Note that all these three supervisory signals are from the data itself.

Learning Objective. Equation 9 summarizes the objective function of the self-supervised graph structure learning.

$$\mathcal{L} = \mathcal{L}_{lap} + \mathcal{L}_{spa} + \mathcal{L}_{rec}$$

$$\mathbf{W}^*, \mathbf{A}^* = \underset{\mathbf{W}, \mathbf{A}}{\arg \min} \mathcal{L}(\mathbf{W}, \mathbf{A}, \mathbf{H}^0)$$
(9)

By minimizing Equation 9, we can jointly refine the graph structure and the graph metric function ϕ . The learning process is executed with two alternating steps. First, we refine the distance metric ϕ by updating the multi-heads parameters $\mathbf{w}_{i=1...m}$, given the current estimation of the graph structure (Section 4.3). Second, we estimate the graph structure using the current distance metric ϕ (Section 4.4). The two steps are complementary to each other and boost the overall accuracy of graph structure recovery.

It is worth noting that all three loss functions are empirically comparable in magnitude in our evaluation. The adversary can weigh the losses to accommodate their specific attack targets in Equation 9.

Adaptive Graph Structure Combination. The learned weighted graph structure \mathbf{A}^{t+1} is then combined with the input graph structure \mathbf{A}^0 using Equation 10. This structure combination step can be interpreted as a denoising function to reduce false positive edges incurred by the initial adjacency matrix \mathbf{A}^0 . That is, GAE learns to reconstruct a graph structure \mathbf{A}^{t+1} given its structure \mathbf{A}^t and node feature \mathbf{H}^t . The edges reconstructed with high confidence are likely to appear the original graph. we use $\mathbf{\eta}$ in Equation 10 to control the update rate of \mathbf{A}^0 using \mathbf{A}^{t+1} , and iteratively filter out the false positive edges from the initial graph structure.

$$\mathbf{A}^{t+1} = (1 - \eta)\mathbf{A}^0 + \eta \mathbf{A}^{t+1}$$
 (10)

Note that \mathbf{A}^{t+1} remains a weighted adjacency matrix after combination. At the end of each iteration, however, we need to obtain the learned graph structure in a binary form to guide graph metric learning in the next iteration. To this end, we first apply an entrywise clipping function, clip(x) = min(max(0,x),1), to \mathbf{A}^{t+1} . We then use the same Bernoulli binarization strategy outlined in [9] to obtain the binary adjacency matrix \mathbf{A}^{t+1} .

Specifically, we treat each element of the weighted adjacency matrix \mathbf{A}^{t+1} as the parameter of a Bernoulli distribution and sample independently to produce the final binary adjacency matrix.

Summary. We summarize the whole learning process (i.e., Section 4.2, Section 4.3 and Section 4.4) in Algorithm 1. Additional details (e.g., complexity analysis) can be found in Appendix A.

Table 2: Summary of datasets.

Dataset	Category	$ \mathbf{V} $	E	$ \mathbf{X} $	$\lceil E / V \rceil$	Density
Cora	Citation	2,708	5,429	1,433	4	0.0014
Citeseer	Citation	4,230	5,358	602	3	0.0006
Actor	Co-Occurrence	7,600	33,544	931	9	0.0011
Facebook	Social	4,039	88,234	1,283	43	0.011

5 Evaluation

5.1 Experimental Setup

Datasets. We use 4 public benchmark datasets to evaluate the performance of our graph reconstruction attack, including Cora [76], Citeseer [76], Actor [51], and Facebook [49]. Cora and Citeseer are citation networks with nodes representing publications and edges indicating citations among these publications. Actor is the actor-only induced subgraph of the film-director-actor-writer network used in [51]. Each node corresponds to an actor, and the edge between two nodes denotes co-occurrence on the same Wikipedia page. Facebook is a social network where nodes represent Facebook users and edges are friendships. We use these datasets to verify the efficacy of our attack given graphs with different characteristics (e.g., origin, graph size, density, node feature size, etc.). For example, Facebook is a social network, it has a well known small world phenomenon and tight community structures among the nodes while the other networks are relatively sparse. Statistics of these datasets are summarized in Table 2.

Node Embedding Models (f). We use four popular node embedding models - network embedding as sparse matrix factorization (NetSMF) [54], Deepwalk (abbreviated as DW) [53], Node2Vec (abbreviated as N2V) [20] and graph convolutional network (GCN) [36] - to generate node embeddings for our evaluation. These four node embedding models are representative of the existing node embedding model families. Network embedding as sparse matrix factorization (NetSMF) [54] improves NetMF [55] and represents the state-of-the-art matrix factorization based approach to generate node embeddings. Deepwalk and Node2Vec are two well known shallow neural network-based (i.e., a neural network with one hidden layer) node embedding techniques. Graph convolutional network (GCN) is a widely used deep neural network based approach for graph representation learning. Note that NetMF, Deepwalk, and Node2Vec generate node embedding using graph structural information only, while GCN considers both node feature and graph structure. As such, these models also cover different real world use cases whereas node embeddings can be generated with different inputs. For reproducibility purposes, we outline their details below.

- **NetSMF.** We use the Pytorch implementation by the original authors [7]. The window size of approximate matrix is 10. The number of negative nodes in sampling is 1. We run the path sampling algorithm for 100 iterations.
- **Deepwalk.** We use the DGL implementation of Deepwalk. The learning rate is set to 0.1. The number of negative

nodes in sampling is 5. The random walk length is fixed at 80, and we run 10 random walks per node.

- Node2Vec. We also use the DGL implementation of Node2Vec. The number of negative nodes in sampling is 5. The random walk length is fixed at 50, and we run 100 random walks per node. p and q are set to 0.25 and 4 respectively by default.
- Graph Convolutional Network (GCN). We use the Pytorch Geometric implementation of GCN. Our GCN model consists of 2 layers as suggested by the original authors. For the first hidden layer, we set the hidden unit size to twice the size of input vectors. For the second layer, we set the hidden unit size to the embedding size. We use ReLU as the activation function between layers. Node embeddings are generated using link prediction as the objective function. We train the GCN model for 400 epochs.

For all node embedding models, we set their output embedding size (i.e., d) to 64, 128, and 256 for our evaluation. These sizes are commonly used in the real world practices balancing between the expressiveness of the node embeddings and the computational complexity of the downstream tasks. Besides, we use the largest connected components from all four datasets to accommodate these node embedding models in our evaluation.

Competitors. We implement three baseline methods detailed below for comparison study.

- **Direct Recovery.** This baseline computes the pairwise similarity matrix from the embeddings of the original graph and reconstructs the graph by choosing the top $k \times n/2$ pairs (i.e., edges) of the largest pairwise similarity scores. It is a straightforward attack strategy that can be leveraged by the adversaries since the embeddings of similar nodes should be close in the latent spaces (see Section 2). Note that our implementation of direct recovery is identical to the decoder used by Duddu et al. [14] to reconstruct graphs.
- kNN Graph. We employ the widely used kNN algorithm (see Section 2) as the second baseline. kNN builds a graph in which two nodes v and u are connected by an edge if the distance between h_v and h_u is among the k-th smallest distances. We use cosine similarity as the distance function.
- Invert Embedding [9]. We adapt the optimization algorithm (Algorithm 2&3 in Chanpuriya et al. [9]) as our third baseline to recover a graph from the node embeddings. Since the attackers cannot obtain the real eigenvalues from the PPMI matrix in a model agnostic setting, we thereby use a random diagonal eigenvalue matrix together with the node embedding matrix to generate the low-rank approximation matrix. We set the other hyperparameters as outlined in Chanpuriya et al. [9]. Additional discussion about invert embedding can be found in Section 7.

The graph size (i.e., the number of edges) of all baselines are set to $k \times n/2$. We detail how we estimate k in Section 5.2

and how k influences the graph recovery performance in Section 5.5.

Hyperparamter Configurations. We set the number of attention heads m to 16. The temperature τ , graph sparsity hyperparameters α and β , and the update rate η are set to 1, 0.3, 0.1 and 0.5 respectively. We set the maximum iteration T to 400. We use a linear graph autoencoder (i.e., Z is set to 1) proposed by Salha et al. [61], which is an effective alternative to multilayer GCNs. These hyperparameter values offer consistent performance across different datasets and models in our evaluation.

Evaluation Metrics. Recall that the attackers have two main goals. Their primary goal is uncovering the edges with decent accuracy from the node embedding matrix, and their secondary goal is recovering a graph structure that is similar to the original graph with respect to the graph properties. Bearing them in mind, we use two categories of metrics to evaluate MNEMON's performance.

• Edge Metrics. We first use four edge related metrics - precision (P), recall (R), F1 score (F1), and joint degree distribution (JDD) - to measure how MNEMON attains the primary goal. Precision, recall, and F1 are commonly used, and we apply them to measure the overall capability of MNEMON recovering the exact edges. The joint degree distribution is a metric relating to the edge distribution and provides an additional measurement about 1-hop neighborhoods around a node. It examines each pair of connected nodes and notes their respective nodal degrees. It is defined as $P(k_1, k_2) = \mu(k_1, k_2) \times m(k_1, k_2)$, where $\mu(k_1, k_2) = 1$ if $k_1 = k_2$ otherwise 2, and $m(k_1, k_2)$ denotes the number of edges connecting nodes of degree k_1 and k_2 .

We use SecGraph [29] to calculate the Jaccard similarity among two JDDs. For all edge metrics, values close to 1 are the best.

• Global Metrics. We then employ three global metrics relative Frobenius error, relative triangle error, and relative average clustering coefficient error - to measure how MNEMON achieves its secondary goal. The relative error is defined as the absolute error (i.e., the difference between the measured value and ground truth value) divided by the ground truth value. It gives an indication of how good a measurement is relative to the ground truth value, or in other words, how much the observed value deviates from actual value. We use the relative Frobenius error, which measures the difference between the adjacency matrix A_{O} and \mathbf{A}_R , i.e., $\|\mathbf{A}_O - \mathbf{A}_R\|_F / \|\mathbf{A}_O\|_F$. Similarly, we count the absolute difference between the number of triangles (respectively average clustering coefficient) of G_R and that of G_O , then divided by the number of triangles (respectively average clustering coefficient) of G_O to calculate the relative triangle error (the relative average clustering coefficient error). For all global metrics, values close to 0 are the best.

Similar relative error metrics are also used in Chanpuriya et al. [9].




Figure 2: Distribution of estimated average node degrees. The mean and standard deviation of our estimated average node degree of the citation/co-occurrence graphs (Figure 2a) are 4.6 and 0.8. The respective values of social networks (Figure 2b) are 45.7 and 8.4.




Figure 3: F1 scores and relative triangle error scores of all basesline methods and MNEMON when given different node embedding sizes (i.e., 64, 128 and 256). We use Node2Vec to generate node embedding matrices.

In practice, the audience could potentially leverage Narayanan-Shmatikov's attack [50] (and other appropriate de-anonymization attacks) to measure, to what extent, the recovered graph can assist the graph de-anonymization task given different graphs and different levels of background knowledge.

Runtime Configuration. All the experiments in this paper are repeated 5 times. For each run, we follow the same experimental setup laid out before. We report the mean and standard deviation of each metric to evaluate the attack performance. In this way, we can delineate objective performance results without reporting opportunistically optimal results.

5.2 How to Estimate the Average Node Degree?

Recall that the only clue that the attackers have is the background information about the origin of the node embedding matrix. In this section, we exemplify how the attackers can estimate the average node degree k from the graphs of similar origins by leveraging state-of-the-art graph sampling methods. Note that the attackers can estimate the graph size (i.e., the number of edges) which equals to $k \times n/2$. We use the state-of-the-art spikyball sampling [58] to estimate the average node degree for our evaluation. It generalizes several exploration-based sampling schemes (e.g., Snowball sampling, Forest Fire sampling, graph-expander sampling etc.), and can be applied to any large graphs due to its flexibility [58].

Specifically, for citation/co-occurrence graphs (e.g., Cora, Citeseer, and Actor), we use the publicly available citation graphs - Pubmed and DBLP - to estimate the average node

degree. For each graph, we use spikyball sampling to sample 30% of the whole graph then estimate the average node degree from the sampled graph. This process is repeated 300 times. We calculate the mean average node degree as our final estimation of citation/co-occurrence graphs. For social network graphs (e.g., Facebook), we randomly select six graphs (e.g., socfb-BU10, socfb-Carnegie49, socfb-JMU79, socfb-Lehigh96, socfb-Maine59 and socfb-UCSC68) from the publicly available FB100 dataset [57] plus one Twitter graph [49] from SNAP. We also sample 30% of each graph to estimate the average node degree and repeat this process 300 times per graph. This strategy enables the adversary to sample enough graphs to cover a wide spectrum of graph properties.

The estimation results are shown in Figure 2. What can be seen in Figure 2 is that the estimated average node degrees may not exactly match the real values but are roughly within the same order of magnitude. For instance, the mean and standard deviation of our estimated average node degree of the citation/co-occurrence graphs (Figure 2a) are 4.6 and 0.8, while the mean and standard deviation of our estimated average node degree of social networks are 45.7 and 8.4 respectively. Comparing to the real values in Table 2, the estimated values are not precise. For instance, the above graph sampling process overestimates the average node degree for the Cora and Citeseer datasets, while underestimating the average node degree of the Actor dataset. However, they can offer the attackers a reasonable starting point to estimate the graph sizes. We use these estimated values (i.e., 5 for citation/cooccurrence graphs and 46 for social network graphs) in the rest of our evaluation. We provide a detailed study on how MNEMON can attain good performance even when the estimated average node degree is almost twice the ground truth value in Section 5.5.

Takeaways. When only having the background information about the origin of the node embedding matrix, our sampling process represents a feasible way that the attackers take to estimate the average node degree. The estimated average node degrees are in the vicinity of the ground truth values but not exactly matching them.

5.3 Is MNEMON Better than the Baselines?

In this section, we aim at studying whether MNEMON is effective to recover a graph from the node embedding matrix, or whether the existing baseline methods would be enough for the task at hand. To address this research question, we compare MNEMON to the baseline methods discussed in Section 5.1. All methods use the estimated average node degrees outlined in Section 5.2. The node embedding size is fixed to 256. Due to space limitations, we only report the attack results on the Cora dataset. The results of the other three datasets can be found in Appendix B.

Performance. The performance comparison results are shown in Table 3. Overall, direct graph recovery and invert embedding graph recovery cannot recover graphs from the node embedding matrices given all four node embedding models. For instance, the F1 scores of these two methods are no greater than 0.027, indicating that they cannot attain the

	able 3: Comparison of all baseline methods and MNEMON. We use the Cora dataset and the node er	nbedding size is 256.
--	--	-----------------------

Graph	f		Edge 1	Metric			Global Metr	ic
Reocvery Method	,	Precision	Recall	F1	JDD	Frobenius Error	Triangle Error	Clustering Coef. Error
	DW	0.001±0.000	0.002±0.000	0.001±0.000	0.000±0.000	1.647±0.000	6.867±0.000	0.753±0.000
Direct	N2V	0.003 ± 0.000	0.006 ± 0.000	0.004 ± 0.000	0.000 ± 0.000	1.645 ± 0.000	7.559 ± 0.000	0.759 ± 0.000
Recovery	NetSMF	0.013 ± 0.000	0.022 ± 0.000	0.016 ± 0.000	0.311 ± 0.000	1.647 ± 0.000	0.621 ± 0.000	0.228 ± 0.000
	GCN	0.001 ± 0.000	0.002 ± 0.000	0.002 ± 0.000	0.000 ± 0.000	1.647 ± 0.000	6.391 ± 0.000	0.653 ± 0.000
	DW	0.007±0.005	0.012±0.010	0.009±0.007	0.667±0.092	1.660±0.010	0.866±0.661	0.198±0.008
Invert Embedding	N2V	0.021 ± 0.008	0.037 ± 0.014	0.027 ± 0.010	0.665 ± 0.030	1.645 ± 0.008	1.869 ± 0.160	0.148 ± 0.005
	NetSMF	0.003 ± 0.001	0.005 ± 0.003	0.004 ± 0.002	0.462 ± 0.064	1.676 ± 0.007	0.842 ± 0.789	0.221 ± 0.011
	GCN	0.015 ± 0.003	$0.026{\pm}0.006$	0.019 ± 0.004	0.675 ± 0.004	$1.657 {\pm} 0.005$	$0.255{\pm}0.158$	$0.188 {\pm} 0.005$
	DW	0.401±0.000	0.492±0.000	0.442±0.000	0.340±0.000	1.114±0.000	3.029±0.000	0.286±0.000
kNN	N2V	0.397 ± 0.000	0.487 ± 0.000	0.438 ± 0.000	0.338 ± 0.000	1.119 ± 0.000	2.185 ± 0.000	0.276 ± 0.000
Graph	NetSMF	0.469 ± 0.000	0.575 ± 0.000	0.517 ± 0.000	0.334 ± 0.000	1.037 ± 0.000	2.379 ± 0.000	0.325 ± 0.000
	GCN	0.378 ± 0.000	$0.463 {\pm} 0.000$	$0.416 {\pm} 0.000$	$0.333 {\pm} 0.000$	1.140 ± 0.000	2.172 ± 0.000	$0.286{\pm}0.000$
	DW	0.492±0.004	0.578±0.003	0.531±0.003	0.840±0.011	1.010±0.005	1.118±0.036	0.215±0.006
MNEMON	N2V	0.506 ± 0.001	0.554 ± 0.003	0.529 ± 0.001	0.724 ± 0.007	0.993 ± 0.001	0.973 ± 0.027	0.228 ± 0.004
MINEMON	NetSMF	0.579 ± 0.002	0.640 ± 0.003	0.608 ± 0.003	0.732 ± 0.006	0.908 ± 0.003	1.263 ± 0.019	0.288 ± 0.004
	GCN	0.462 ± 0.002	0.506 ± 0.001	0.483 ± 0.001	0.753 ± 0.006	1.040 ± 0.003	$0.864 {\pm} 0.032$	0.230 ± 0.005

attacker's primary goal. At the same time, the global metrics of these two baselines are equally underwhelming. Our results show that such optimization based approach is less effective in a model agnostic setting. kNN algorithm represents a de facto approach to recover the edges from node embeddings. Our results show that kNN graph can partially recover the edges from the node embedding matrix. For example, it can recover the edges from the node embeddings generated by Node2Vec with a 0.438 F1 score. As we can see in Table 3, MNEMON outperforms all baseline methods. Take the node embeddings generated by Node2Vec for example, MNEMON achieves 0.529 F1 score, which is 0.091 higher than that of kNN graph recovery. In other words, MNEMON's F1 score relatively improves that of kNN graph recovery by 0.208 (i.e., 0.091/0.438=0.208). If we take the edges recovered by kNN graph as the upper bound of the existing privacy risk assessment, MNEMON empirically improves this upper bound by 0.208 per our evaluation results. Given the combinatorial nature of graph edges (i.e., n(n-1)/2 possibilities) and our strict attack setting (i.e., no interaction with the node embedding models), such 0.208 relative improvement by MNEMON is substantial. Practically speaking, if we position MNEMON in the privacy risk assessment framework, it would lead to 0.208 increase of the estimated privacy loss than the de facto risk assessment using kNN algorithm. We also provide a visual explanation to exemplify MNEMON's capability in recovery graphs from the node embedding matrices in Appendix B.

Impact of Node Embedding Size. We use two metrics - F1 score and relative triangle error - to understand the impact of node embedding size on both baselines and MNEMON. We use Node2Vec to generate node embedding matrices. The results are shown in Figure 3. It is straightforward to see that MNEMON consistently performs better than the baselines given different node embedding sizes.

Takeaways. The proposed learnable distance function and adaptive graph structure combination can reduce a reasonable amount of false edges. They, in turn, enable MNEMON to recover better graph structure from the node embedding matrix given different node embedding models and embedding sizes. Besides, *k*NN graph remains a viable approach

to recover edges from the node embedding matrix. However, due to its non-learning-based nature, *k*NN graph is outperformed by MNEMON.

5.4 How Effective is MNEMON?

In this section, we evaluate MNEMON on all four datasets to understand its overall performance. The results are summarized in Table 4. Due to space limitations, we only show the results when the node embedding size is fixed to 256. The performance results of 64- and 128-dimensional node embeddings follow similar patterns and can be found in Appendix C.

Edge Metrics. Recall that the primary goal of the attackers is uncovering the edges with decent accuracy from the node embedding matrix. We thereby use edge metrics outlined in Section 5.1 to measure MNEMON's performance. Besides, kNN graph remains a viable approach to recover edges from the node embedding matrix as we explicate in Section 5.3. We also show the relative improvement scores in Table 4 to demonstrate to what extent MNEMON can relatively improve from kNN graph. We add a positive sign (+)next to the relative improvement score to highlight the improvement. As we can see from Table 4, MNEMON can enable the adversary to recover edges from all node embedding matrices generated by all four node embedding models with good precision and recall. Take the Cora dataset and the node embedding matrix generated by NetMF for example. MNEMON achieves 0.579 precision, 0.640 recall and 0.608 F1 scores. These scores relatively improve 0.235, 0.113, and 0.176 from those of kNN graph recovery. Similarly, take the Actor dataset and the node embedding matrix generated by Deepwalk for example, MNEMON achieves 0.687 precision, 0.435 recall, and 0.533 F1 scores. These scores relatively improve 0.222, 0.088, and 0.138 from those of kNN graph recovery. The other datasets given all node embedding models follow similar patterns. At the same time, MNEMON overwhelmingly outperforms kNN graph given joint degree distribution similarity metric. Given the above two examples, JDD similarity scores of MNEMON respectively improve 1.191 and 0.587 from those of kNN graph. Our results demonstrate that MNEMON can recover a graph

Table 4: The performance results of MNEMON using all four datasets. We fix the node embedding size to 256. We show the relative improvement scores in edge metrics to demonstrate to what extent MNEMON can relatively improve from kNN graph. We add a positive sign (+) next to the relative improvement score to highlight the improvement. We also show the relative error reduction scores in global metrics to demonstrate to what extent MNEMON can relatively reduce errors incurred by kNN graph. We add a negative sign (-) next to the relative error reduction score to highlight the difference.

Dataset	f		Edge I	Metrics		Global Metrics			
Dataset	J	Precision	Recall	F1	JDD	Frobenius Error	Triangle Error	Clustering Coef. Error	
	DW	0.492±0.004 (+0.226)	0.578±0.003 (+0.174)	0.531±0.003 (+0.201)	0.840±0.011 (+1.471)	1.010±0.005 (-0.104)	1.118±0.036 (-1.911)	0.215±0.006 (-0.071)	
Cora	N2V	$0.506\pm0.001~(+0.276)$	0.554±0.003 (+0.137)	$0.529 \pm 0.001 \ (+0.208)$	0.724 ± 0.007 (+1.143)	0.993±0.001 (-0.126)	0.973±0.027 (-1.212)	0.228 ± 0.004 (-0.048)	
Cora	NetSMF	$0.579\pm0.002~(+0.235)$	0.640±0.003 (+0.113)	0.608 ± 0.003 (+0.176)	0.732±0.006 (+1.191)	0.908±0.003 (-0.129)	1.263±0.019 (-1.116)	0.288±0.004 (-0.037)	
	GCN	$0.462{\pm}0.002~(+0.223)$	$0.506 {\pm} 0.001 ~(+0.092)$	$0.483 \pm 0.001 \ (+0.162)$	$0.753 \pm 0.006 \ (+1.260)$	1.040±0.003 (-0.100)	0.864±0.032 (-1.308)	0.230±0.005 (-0.056)	
	DW	0.403±0.002 (+0.193)	0.555±0.005 (+0.149)	0.467±0.003 (+0.174)	0.617±0.011 (+1.635)	1.125±0.003 (-0.085)	1.877±0.075 (-2.651)	0.341±0.009 (-0.080)	
	N2V	0.445±0.001 (+0.271)	0.575±0.002 (+0.149)	0.502±0.001 (+0.217)	0.506±0.007 (+1.137)	1.069±0.002 (-0.127)	1.734±0.039 (-1.665)	0.357±0.005 (-0.059)	
Citeseer	NetSMF	0.530±0.002 (+0.229)	0.672±0.001 (+0.091)	0.592±0.001 (+0.168)	0.461±0.005 (+1.048)	0.961±0.002 (-0.133)	2.001±0.056 (-1.419)	0.432±0.003 (-0.056)	
	GCN	0.414±0.003 (+0.206)	0.529±0.002 (+0.080)	0.465±0.001 (+0.153)	0.527±0.011 (+1.344)	1.105±0.004 (-0.099)	1.467±0.057 (-1.734)	0.330±0.004 (-0.067)	
	DW	0.687±0.001 (+0.222)	0.435±0.002 (+0.088)	0.533±0.002 (+0.138)	0.417±0.001 (+0.587)	0.874±0.001 (-0.081)	0.203±0.009 (-0.105)	0.229±0.002 (-0.017)	
	N2V	0.465±0.001 (+0.356)	0.313±0.000 (+0.282)	0.374±0.000 (+0.312)	0.477 ± 0.001 (+0.293)	1.023±0.001 (-0.083)	0.179±0.007 (-0.387)	0.176±0.001 (-0.035)	
Actor	NetSMF	0.562±0.001 (+0.240)	0.366±0.001 (+0.136)	0.443±0.001 (+0.179)	0.475 ± 0.003 (+0.406)	0.959 ± 0.001 (-0.074)	0.147±0.007 (-0.387) 0.147±0.013 (-0.758)	0.285 ± 0.002 (-0.025)	
	GCN	0.373±0.001 (+0.226)	0.263±0.000 (+0.211)	0.308±0.001 (+0.218)	0.505 ± 0.003 (+0.446)	1.086±0.001 (-0.045)	0.280±0.008 (-0.349)	0.153±0.002 (-0.049)	
	DW	0.441±0.001 (+0.028)	0.471±0.001 (+0.066)	0.456±0.001 (+0.046)	0.519±0.006 (+1.745)	1.061±0.001 (-0.009)	0.494±0.002 (+0.213)	0.077±0.001 (+0.006)	
Facebook	N2V	0.468±0.000 (+0.018)	0.487±0.001 (+0.026)	0.477±0.001 (+0.022)	0.444±0.002 (+1.581)	1.033±0.001 (-0.007)	0.545±0.001 (+0.499)	0.090±0.001 (+0.050)	
	NetSMF GCN	0.454±0.001 (+0.022) 0.342±0.001 (+0.061)	0.502±0.002 (+0.098) 0.364±0.001 (+0.100)	0.476±0.001 (+0.059) 0.352±0.001 (+0.078)	0.457±0.002 (+0.570) 0.371±0.004 (+1.026)	1.050±0.001 (-0.006) 1.157±0.001 (-0.012)	0.424±0.007 (+0.418) 0.452±0.002 (+0.380)	0.081±0.001 (+0.041) 0.056±0.001 (-0.031)	
00			0.700		0.700		0.700		
00			0.600		0.600		0.600		
			Φ		Φ		Ψ		
00			0.500 OS		0.500		0.500 S		
00		-+- DW	₫ 0.400	-+- DW	₫ 0.400	-+- DW	₩ 0.400	-+- D	
00		-•·· N2V -■·· NetSMF	0.300	-•·· N2V -•·· NetSMF	0.300	-e·· N2V -■·· NetSM		N	
		····· GCN		····· GCN	Ţ	····· GCN		G	
64	128 d	256	0.200 64	128 256 d	0.200 64	128 2	0.200 64	128 d	
((a) Cora	1	(b) Cit	eseer	(c)	Actor		(d) Facebook	
00			는 0.500		ō 0.500		는 0.500		
00			0.400		Ē 0.400		ш , 400		
00 *			0.400		0.400		0.400		
00			0.300		0.300		0.300		
00			0.200		0.200	·	0.200		
		-+- DW	0.200	-+- DW	9	-+- DW	9	-+-	
00		N2V NetSMF	0.100 1.00 1.00 1.00 1.00 1.00 1.00 1.0	-0- N2V -8- NetSMF	0.100	N2V NetSM	8 0.100 H	-0-	
64	128	256	윤 64	128 256	e 64	128 25	6 64	128	
UH	120 d	230	04	d 230	U 4	d 23	0 04	d	
	(e) Cora		(f) Cite			Actor		Facebook	

Figure 4: F1 scores and relative average clustering coefficient error scores of MNEMON given all four datasets. We fix the node embedding size to 256.

in which each pair of connected nodes share similar 1-hop neighborhood as they are in the original graph.

Recall that the secondary goal of the attackers is recovering a graph structure that is similar to the original graph with respect to the graph properties. We use the global metrics outlined in Section 5.1 to understand MNEMON's performance. Similar to the above edge metrics, we also compare MNEMON to kNN graph. We show the relative error reduction scores in Table 4 to demonstrate to what extent MNEMON can relatively reduce errors incurred by kNN graph. We add a negative sign (-) next to the relative error reduction score to highlight the difference between MNEMON and kNN graph. As we can see from Table 4, MNEMON can incur relatively low error scores given all three global metrics. Take the Actor dataset and the node embedding matrix generated by GCN for example. MNEMON's relative triangle error is 0.280. This indicates that the graph recovered by MNEMON contains a similar number of triangles to that of the original graph. At the same time, this score reduces the relative error made by kNN graph for 0.349. Note that the estimated average node degree (i.e., 5) is larger than the ground truth values of both Cora and Citeseer. This leads to higher relative triangle errors. However, combined with the edge metrics, we can assert that such error is due to a combination of reorientation of the specific edges between the true and the recovery networks, and extra edges incurred by the overestimation of k. Besides, we compare MNEMON with the invert embedding using the overlapping Citeseer dataset and its 256-dimensional NetSMF node embedding matrix. MNEMON can achieve 0.908 relative Frobenius error score which is close to that of the invert embedding (see Figure 4 in [9]). Note that the invert embedding in [9] is under the white box setting while MNEMON is under the black box setting. In summary, our performance results demonstrate that MNEMON can also recover a graph that is structurally similar to the original graph with respect to the global graph properties.

Note that the clustering coefficient of a node c_v is defined as $c_v = \frac{2*\mathcal{N}(v)}{deg(v)(deg(v)-1)}$ where $\mathcal{N}(v)$ represents the number of edges between the neighbors of v and deg(v) represents the degree of v. The average clustering coefficient of the whole




Figure 5: F1 scores and relative average clustering coefficient error scores of MNEMON given all four datasets and various average node degrees. We fix the node embedding size to 256. The vertical bar indicates the actual average node degree.

graph $c_{\mathbf{G}}$ is defined as $c_{\mathbf{G}} = \frac{1}{n} \sum_{\nu=1}^{n} c_{\nu}$. A smaller k may lead to the increasing possibility that the number of triangles recovered from a graph drops closer to 0. We therefore use different relative errors to objectively evaluate the performance from multiple perspectives.

Impact of Node Embedding Size. We use two metrics -F1 score and relative average clustering coefficient error - to understand the impact of node embedding size MNEMON across all four datasets. The results are shown in Figure 4. As we can see in the figure, MNEMON can offer stable graph recovery performance given different embedding sizes and all embedding models. We only observe a marginal F1 score decrease given the Deepwalk node embedding model in Cora and Citeseer datasets. Overall, our results imply that reducing the embedding size (a common defense mechanism) may not work for MNEMON. More details can be found in Section 6.

Stability. We run MNEMON 5 times on a given node embedding matrix. Such runtime configuration enables us to measure how widely all those metric values are dispersed from the average value (i.e., standard deviation). At the same, it eliminates the chance of reporting opportunistically good results. A low standard deviation indicates low volatility. As we can observe in Table 4, the standard deviation values are low in all cases. The results show that MNEMON can recover graphs from node embedding matrices with statistically stable performance.

Ablation Study. We also carry out an ablation study to understand the impact of GML on the performance of MNEMON. To this end, we customize MNEMON and remove GML from the optimization. Specifically, we initialize a graph using Gumbel-Top-*k* trick and run GAE once. We use edge metrics in this study and summarize the results in Table 5. We observe that MNEMON performs better given all edge metrics. The results exemplify that jointly optimizing GAE and GML enables us to learn more information from the node embedding matrix and further reduce noise from the recovered graph.

Takeaways. We can observe that MNEMON achieves good performance on all datasets. Such results demonstrate

Table 5: Ablation study on the impact of GML to MNEMON's performance. We use Citeseer dataset and fix the node embedding size to 256.

Method	Precision	Recall	F1
kNN	0.338	0.483	0.398
MNEMON w/o GML	0.391	0.511	0.443
MNEMON	0.404	0.557	0.468

that jointly optimizing the learnable distance function and adaptive graph structure combination is effective for recovering graphs from node embeddings.

5.5 How does *k* Affect the Attack Performance?

Recall that the attackers use the graph sampling algorithms to estimate the average node degree from the graphs of similar origins and transfer the estimated node degree from these graphs to facilitate the attack (see Section 4.2). We show that the adversary cannot obtain the precise average node degree in Section 5.2. However, the estimated average node degree k directly affects the graph size of the recovered graph G_R . More importantly, our attack uses this estimated k to seed the initial graph using the Gumbel-Top-k trick and iterative graph structure optimization during the learning process. It is therefore essential to study the impact of the estimated average node degree k on the graph recovery performance. To this end, we run MNEMON 5 times on every k that falls within at least one standard deviation of the estimated average node degree in Section 5.2. For instance, the mean and standard deviation of our estimated average node degree of the Facebook dataset are 45.7 and 8.4 respectively. In this case, we run MNEMON 5 times for every value between 38 and 52. We use two metrics - F1 score and relative average clustering coefficient error - to understand the impact of k across all four datasets. Due to space limitations, we only show the attack results when the node embedding size is fixed to 256. The performance results using 64- and 128dimensional node embeddings follow similar patterns and can be found in Appendix D.




Figure 6: trade-off between node embedding utility (average link prediction precision) and MNEMON's graph recovery performance (F1 score). We use the Cora dataset and GCN as the node embedding model. The node embedding size is fixed to 256.

Performance. The results are shown in Figure 5. We show a vertical line in each figure to mark the ground truth value of the average node degree for each dataset (see Table 2). In general, we can see that MNEMON can accommodate the inevitable estimation error. Take the Citeseer dataset and the node embedding matrix generated by NetSMF for example. The ground truth average node degree is 3, and the estimated average node degree is 5. The F1 scores achieved by MNEMON are respectively 0.592 and 0.605. This means that, even though our estimated k is almost twice the ground truth value, MNEMON can still deal with such estimation error and attain good results (i.e., the F1 score difference is only 0.013). Similar to the F1 score metric, we can see that MNEMON can also achieve low relative average clustering coefficient error. Take the Citeseer dataset and the node embedding matrix generated by NetSMF for instance, the relative average clustering coefficient errors of MNEMON are 0.432 and 0.332. The difference is approximately 0.100.

Observations. We observe that the node embedding matrices generated by GCN are relatively harder to recover than those by the other three models. Two factors make the graph recovery task difficult. First, GCN considers both node features and graph structure to generate node embeddings while the other three models only use graph structures. Second, we use ReLU as the activation function between layers. This non-linear, element-wise function outputs the input directly if it is positive, otherwise, it outputs zero. It is computationally efficient but leads to sparse representation, making the graph recovery harder.

Takeaways. Our evaluation results show that MNEMON can accommodate the inevitable estimation error of k. The root cause of the moderate decrease (increase) of F1 scores (relative average clustering coefficient error) in Figure 5 is due to the increasing graph size. However, given the real world application scenarios outlined in Section 3.2, we show that the estimated average node degrees can be in the vicinity of the real values as shown in Section 5.2. Combining with the adaptive learning process outlined in Section 4.4, MNEMON remains practical to recover graphs in the wild as exemplified by the results in Figure 5.

6 Defense

In this section, we discuss node embedding perturbation as a tentative defense mechanism and empirically evaluate its effectiveness.

Embedding Perturbation. One possible defense of MNEMON is adding perturbations (i.e., noise) to the original node embeddings \mathbf{H}_O . As such, the data holder only passes on a noisy but usable version $\widetilde{\mathbf{H}}_O$ to the ML pipeline. Formally, $\widetilde{\mathbf{H}}_O = \mathbf{H}_O + \Delta(\mu, b)$ where Δ denotes the Laplace distribution, μ is a location parameter and b>0 is a scale parameter. However, adding noise inevitably distorts the information contained in the node embeddings and can lead to utility loss. We therefore focus on evaluating the trade-off between the utility and the defense in this section. Similar defense mechanisms were also discussed in the previous literature [26, 81].

Experimental Setup. We use the Cora dataset and the 256 dimensional node embedding matrix generated by GCN for our evaluation. We fix μ to 0 and choose 10 evenly distributed values between 0 and 1 for b (i.e., $b = \{0,0.1,...,0.9\}$). We use average link prediction precision as the utility metric and F1 score as the attack performance metric.

Results. The results are shown in Figure 6. As we can see in the figure, adding perturbations could work with noticeable utility loss. For instance, when b = 0.2, the average link prediction precision using the perturbated node embedding matrix drops from 0.881 to 0.761. In turn, we can see that the MNEMON's F1 score drops from 0.486 to 0.088. This result shows that the data holder might choose the noise level to defend against MNEMON while preserving some utility. However, it is a delicate process. For instance, if the data holder chooses b = 0.1, the average link prediction precision drops from 0.881 to 0.868. In this case, MNEMON's F1 score drops from 0.486 to 0.401. In short, our results show that the trade-off is inevitable if using added perturbations to defend against MNEMON. We plan to explore such research direction in the future.

Notes. The node embedding size affects the expressiveness of the node embeddings. As such, another prospective defense mechanism is to reduce the dimension of the node embedding. Its core idea is reducing the knowledge that the attackers can obtain and consequently lessening the capability of the graph recovery attack. However, we show that MNEMON achieves stable graph recovery performance given different embedding sizes and all embedding models in Figure 4. Our results indicate that reducing the embedding size may not work for MNEMON. We plan to expand such research direction in the future.

7 Related Work

Graph Theory Based Graph Restoration. Graph restoration algorithms in the graph theory realm restore a hidden graph by repeatedly querying an oracle for certain types of information about the graph structure [40]. Depending on the algorithm, different types of information can be revealed

Table 6: Difference between our attack and the close work.

Method	Supervision from Auxiliary data	•		Attack setting
Chanpuriya et al. [9]	✓	Х	N/A	whitebox
Link reidentification [14,26,71]	1	1	1	blackbox
Zhang et al. [81]	1	1	1	blackbox
MNEMON	×	Х	х	model-agnostic

by the oracle, including node betweenness [1], distance or shortest path between nodes [27], edge counting [5], edge detection [3], etc. The common goal among these research is identifying strategies that recover the graph with low worst-case query complexity. However, those approaches are not learning-based and require the existence of an oracle knowing the structural information of the original graph. They cannot be adapted to reconstruct graphs from the node embeddings.

Graph Completion. Graph completion [11, 21] aims at inferring the unobserved part of the network (i.e., missing edges and nodes) given the partially observed network. Link prediction algorithms (see [15, 46] for an overview) have been actively investigated and successfully applied to identify missing edges [67, 79]. Probabilistic and deep learning models [28, 64] have also been investigated to deduce the missing nodes. However, these algorithms require the graphs to be substantially observed and high-quality attribute information provided. Our attack assumes neither.

Deep Graph Structure Learning for Robust Representations. This line of research centers on Graph Structure Learning (GSL) that jointly learns an optimized graph structure and corresponding representations [86]. The goal of GSL is to generate node representations robust to noisy graph structures. Common assumptions of these methods include the availability of node features, incomplete graph structure, and node labels. Different approaches then leverage metric learning [10], probabilistic modeling [17], direct optimization [73], etc. to learn an adjacency matrix as well as the corresponding node representations. In contrast to GSL, our attack does not assume the availability of node features and node labels. Besides the goal difference, our attack is self-supervised while GSL approaches use node labels to supervise the learning process.

Close Work. To our best knowledge, there exist five pieces of close work to our attack [9, 14, 26, 71, 81]. The closest work is Chanpuriya et al. [9] presenting two optimization algorithms to recover a graph from its node embeddings generated by NetMF [55]. Their algorithms assume the knowledge of the NetMF algorithm (i.e., the target model in our terminology), window size T, the low-ranking approximation of the finite-T positive pointwise mutual information (PPMI) matrix, and the exact degree of each node, hence a specific white-box attack against NetMF only. Another closely related work is link reidentification attack [14,26,71] from node-level information. In theory, those attacks can be used to reconstruct a graph upon querying the target model n^2 times. However, they train a shadow model using auxiliary data and their posterior scores obtained from the target model. Our attack assumes the attackers can not interact with the target model using auxiliary data, which renders this link

stealing attack infeasible in our setting. In addition to link re-identification attacks using node-level information, Zhang et al. [81] also introduce a reconstruction attack to rebuild a graph from its graph-level embedding within the context of graph classification. This attack suffers from the same pitfalls of the link re-identification attacks, and cannot be used in our setting. In short, our attack is fundamentally different from the existing work by removing the assumptions of the availability of supervision information from auxiliary data, the shadow model, and the interaction with the target model. We summarize the differences between our attack and the closely related work in Table 6.

8 Conclusion

In this paper, we presented a model-agnostic attack that uses the node embedding matrices to recover graphs. Extensive experiments show that an adversary can recover graphs with decent accuracy by only gaining access to the node embeddings of the original graph. Our results highlight the need for the data holders to rethink the privacy implications when integrating node embeddings for downstream analysis, even when the third party has extremely limited knowledge of the data.

Acknowledgments

We wish to thank the anonymous reviewers for their feedback and our shepherd Gergely Acs for his help in improving our paper. This work is partially funded by the Helmholtz Association within the project "Trustworthy Federated Data Analytics" (TFDA) (funding number ZT-I-OO1 4) and by the National Science Foundation under grant CNS-2127232.

References

- [1] Mikkel Abrahamsen, Greg Bodwin, Eva Rotenberg, and Morten Stöckel. Graph reconstruction with a betweenness oracle. In *STACS*, 2016. 14
- [2] Seyed Ali Alhosseini, Raad Bin Tareaf, Pejman Najafi, and Christoph Meinel. Detect me if you can: Spam bot detection using inductive representation learning. In WWW, 2019. 1
- [3] Dana Angluin and Jiang Chen. Learning a hidden graph using o(log n) queries per edge. *Journal of Computer and System Sciences*, 74(4), 2008. 14
- [4] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NeurIPS*, 2001. 6
- [5] Mathilde Bouvel, Vladimir Grebinski, and Gregory Kucherov. Combinatorial search on graphs motivated by bioinformatics applications: A brief survey. In *WG*, 2005. 14
- [6] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. A comprehensive survey of graph embedding: Problems, techniques and applications. *IEEE*

- Transactions on Knowledge and Data Engineering, 30(9), 2018. 1, 2, 3, 4
- [7] Yukuo Cen, Zhenyu Hou, Yan Wang, Qibin Chen, Yizhen Luo, Xingcheng Yao, Aohan Zeng, Shiguang Guo, Peng Zhang, Guohao Dai, Yu Wang, Chang Zhou, Hongxia Yang, and Jie Tang. Cogdl: Toolkit for deep learning on graphs. arXiv preprint arXiv:2103.00959, 2021. 7
- [8] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. Sequential recommendation with graph neural networks. In SIGIR, 2021. 5
- [9] Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, and Charalampos E Tsourakakis. Deepwalking backwards: From embeddings back to graphs. In *NeurIPS*, 2021. 7, 8, 11, 14
- [10] Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. In *NeurIPS*, 2020. 5, 14
- [11] Aaron Clauset, Cristopher Moore, and Mark EJ Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 2008. 14
- [12] Andrea Continella, Mario Polino, Marcello Pogliani, and Stefano Zanero. There's a hole in that bucket! a large-scale analysis of misconfigured s3 buckets. In *ACSAC*, 2018. 3
- [13] Anirban Dasgupta, Ravi Kumar, and Tamas Sarlos. On estimating the average degree. In *WWW*, 2014. 4
- [14] Vasisht Duddu, Antoine Boutet, and Virat Shejwalkar. Quantifying privacy leakage in graph embedding. *arXiv* preprint arXiv:1912.10979, 2020. 1, 5, 6, 8, 14
- [15] Daniel M Dunlavy, Tamara G Kolda, and Evrim Acar. Temporal link prediction using matrix and tensor factorizations. *TKDD*, 5(2), 2011. 14
- [16] Ming Fan, Xiapu Luo, Jun Liu, Meng Wang, Chunyin Nong, Qinghua Zheng, and Ting Liu. Graph embedding based familial analysis of android malware using unsupervised learning. In *ICSE*, 2019. 1
- [17] Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. Learning discrete structures for graph neural networks. In *ICML*, 2019. 14
- [18] Jian Gao, Xin Yang, Ying Fu, Yu Jiang, and Jiaguang Sun. Vulseeker: A semantic learning based vulnerability seeker for cross-platform binary. In *ASE*, 2018.
- [19] Palash Goyal and Emilio Ferrara. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 151, 2018. 1, 2, 3, 4
- [20] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, 2016. 2, 7
- [21] Roger Guimerà and Marta Sales-Pardo. Missing and spurious interactions and the reconstruction of complex networks. *PNAS*, 2009. 14

- [22] Emil Julius Gumbel. Statistical Theory of Extreme Values and Some Practical Applications: A Series of Lectures. US Government Printing Office, 1954. 5
- [23] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NeurIPS*, 2017. 3
- [24] Stephen J Hardiman and Liran Katzir. Estimating clustering coefficients and size of social networks via random walk. In *WWW*, 2013. 4
- [25] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. Accurate estimation of the degree distribution of private networks. In *ICDM*, 2009. 4
- [26] Xinlei He, Jinyuan Jia, Michael Backes, Neil Zhenqiang Gong, and Yang Zhang. Stealing Links from Graph Neural Networks. In USENIX Security, 2021. 1, 5, 13, 14
- [27] Danny Hermelin, Avivit Levy, Oren Weimann, and Raphael Yuster. Distance oracles for vertex-labeled graphs. In *ICALP*, 2011. 14
- [28] Darko Hric, Tiago P Peixoto, and Santo Fortunato. Network structure, metadata, and the prediction of missing nodes and annotations. *Physical Review X*, 2016. 14
- [29] Shouling Ji, Weiqing Li, Prateek Mittal, Xin Hu, and Raheem Beyah. Secgraph: A uniform and open-source evaluation system for graph data anonymization and deanonymization. In *USENIX Security*, 2015. 8
- [30] Shouling Ji, Weiqing Li, Mudhakar Srivatsa, and Raheem Beyah. Structural data de-anonymization: Quantification, practice, and implications. In *ACM CCS*, 2014. 4
- [31] Shouling Ji, Prateek Mittal, and Raheem Beyah. Graph data anonymization, de-anonymization attacks, and de-anonymizability quantification: A survey. *IEEE Communications Surveys & Tutorials*, 2016. 1, 4
- [32] Ian T Jolliffe and Jorge Cadima. Principal component analysis: A review and recent developments. *Philos. Trans. A Math. Phys. Eng Sci.*, 2016. 2
- [33] Vassilis Kalofolias. How to learn a graph from smooth signals. In *AISTATS*, 2016. 6
- [34] Anees Kazi, Luca Cosmo, Nassir Navab, and Michael Bronstein. Differentiable graph module (dgm) for graph convolutional networks. *arXiv preprint arXiv*:2002.04999, 2020. 5
- [35] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular Graph Convolutions: Moving Beyond Fingerprints. *Journal of Computer-Aided Molecular Design*, 2016. 1
- [36] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2016. 3, 6, 7

- [37] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016. 5, 6
- [38] Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *ACM STOC*, 2000. 4
- [39] Wouter Kool, Herke Van Hoof, and Max Welling. Stochastic beams and where to find them: The gumbeltop-k trick for sampling sequences without replacement. In *ICML*, 2019. 5
- [40] Josef Lauri and Raffaele Scapellato. *Topics in Graph Automorphisms and Reconstruction*. Cambridge University Press, 2016. 13
- [41] Jure Leskovec and Christos Faloutsos. Sampling from large graphs. In *KDD*, 2006. 4
- [42] Xiaoxiao Li, João Saúde, Prashant Reddy, and Manuela Veloso. Classifying and Understanding Financial Data Using Graph Neural Network. In *The AAAI Workshop on Knowledge Discovery from Unstructured Data in Financial Services (KDF)*. AAAI, 2020. 1
- [43] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *JAIST*, 58(7):1019–1031, 2007. 1
- [44] Jenny Liu, Aviral Kumar, Jimmy Ba, Jamie Kiros, and Kevin Swersky. Graph normalizing flows. *NeurIPS*, 32, 2019. 6
- [45] Kun Liu and Evimaria Terzi. Towards identity anonymization on graphs. In *SIGMOD*, 2008. 1
- [46] Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A*, 390(6), 2011. 14
- [47] Mohammad Malekzadeh, Anastasia Borovykh, and Deniz Gündüz. Honest-but-curious nets: Sensitive attributes of private inputs can be secretly coded into the entropy of classifiers' outputs. In *ACM CCS*, 2021. 1, 3
- [48] Fragkiskos D Malliaros and Michalis Vazirgiannis. Clustering and community detection in directed networks: A survey. *Physics reports*, 533(4), 2013. 1
- [49] Julian J McAuley and Jure Leskovec. Learning to discover social circles in ego networks. In *NeurIPS*, 2012. 7, 9
- [50] Arvind Narayanan and Vitaly Shmatikov. Deanonymizing social networks. In *IEEE S&P*, 2009. 9
- [51] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020. 7
- [52] Abdurrahman Pektaş and Tankut Acarman. Deep learning for effective android malware detection using api call graph embeddings. *Soft Computing*, 24(2), 2020. 1
- [53] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In KDD, 2014. 2, 7

- [54] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Chi Wang, Kuansan Wang, and Jie Tang. Netsmf: Large-scale network embedding as sparse matrix factorization. In *WWW*, 2019. 7
- [55] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *WSDM*, 2018. 7, 14
- [56] Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. DeepInf: Social Influence Prediction with Deep Learning. In *KDD*, 2018.
- [57] Veronica Red, Eric D Kelsic, Peter J Mucha, and Mason A Porter. Comparing community structure to characteristics in online collegiate social networks. *SIAM review*, 53(3), 2011. 9
- [58] Benjamin Ricaud, Nicolas Aspert, and Volodymyr Miz. Spikyball sampling: Exploring large networks via an inhomogeneous filtered diffusion. *Algorithms*, 2020. 9
- [59] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. 4
- [60] Benedek Rozemberczki, Oliver Kiss, and Rik Sarkar. Little ball of fur: A python library for graph sampling. In CIKM, 2020. 4
- [61] Guillaume Salha, Romain Hennequin, and Michalis Vazirgiannis. Simple and effective graph autoencoders with one-hop linear models. In *ECML-PKDD*, 2020. 6,
- [62] Yun Shen, Xinlei He, Yufei Han, and Yang Zhang. Model stealing attacks against inductive graph neural networks. In *IEEE S&P*, 2022. 1
- [63] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *IEEE S&P*, 2017. 1
- [64] Sigal Sina, Avi Rosenfeld, and Sarit Kraus. Solving the missing node problem using structure and attribute information. In *ASONAM*, 2013. 14
- [65] Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. Machine learning models that remember too much. In *ACM CCS*, 2017. 1, 3
- [66] Lichao Sun, Yingtong Dou, Carl Yang, Ji Wang, Philip S Yu, Lifang He, and Bo Li. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528*, 2018. 1
- [67] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, 2016. 14
- [68] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 5

- [69] Binghui Wang, Jinyuan Jia, and Neil Zhenqiang Gong. Graph-based security and privacy analytics via collective classification with joint weight learning and propagation. In NDSS, 2019.
- [70] Jianyu Wang, Rui Wen, Chunming Wu, Yu Huang, and Jian Xion. Fdgars: Fraudster detection via graph convolutional networks in online app review system. In WWW, 2019. 1
- [71] Fan Wu, Yunhui Long, Ce Zhang, and Bo Li. Linkteller: Recovering private edges from graph neural networks via influence analysis. In *IEEE S&P*, 2022. 1, 3, 5, 14, 17
- [72] Zhaohan Xi, Ren Pang, Shouling Ji, and Ting Wang. Graph backdoor. In *USENIX Security*, 2021. 1
- [73] Liang Yang, Zesheng Kang, Xiaochun Cao, Di Jin, Bo Yang, and Yuanfang Guo. Topology optimization based graph convolutional network. In *IJCAI*, 2019. 14
- [74] Liu Yang and Rong Jin. Distance metric learning: A comprehensive survey. *Michigan State Universiy*, 2(2), 2006. 5
- [75] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2019. 3
- [76] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *ICML*, 2016. 7
- [77] Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. Network representation learning: A survey. *IEEE transactions on Big Data*, 2018. 1, 2, 3, 4
- [78] Minxing Zhang, Zhaochun Ren, Zihan Wang, Pengjie Ren, Zhunmin Chen, Pengfei Hu, and Yang Zhang. Membership inference attacks against recommender systems. In ACM CCS, 2021. 1
- [79] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *NeurIPS*, 2018. 14
- [80] Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Zhenqiang Gong. Backdoor attacks to graph neural nnetworks. In *SACMAT*, 2021. 1
- [81] Zhikun Zhang, Min Chen, Michael Backes, Yun Shen, and Yang Zhang. Inference attacks against graph neural networks. In *USENIX Security*, 2022. 1, 13, 14, 17
- [82] Gang Zhao and Jeff Huang. Deepsim: Deep learning code functional similarity. In *ESEC/FSE*, 2018. 1
- [83] Jianan Zhao, Xiao Wang, Chuan Shi, Binbin Hu, Guojie Song, and Yanfang Ye. Heterogeneous graph structure learning for graph neural networks. In *AAAI*, 2021.
- [84] Bin Zhou, Jian Pei, and WoShun Luk. A brief survey on anonymization techniques for privacy preserving publishing of social network data. *ACM SIGKDD Explorations Newsletter*, 2008. 1

- [85] Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks. In *NeurIPS*, 2019.
- [86] Yanqiao Zhu, Weizhi Xu, Jinghao Zhang, Qiang Liu, Shu Wu, and Liang Wang. Deep graph structure learning for robust representations: A survey. *arXiv preprint arXiv:2103.03036*, 2021. 14

A MNEMON **Algorithm Details**

Walk Through. We summarize the whole learning process (i.e., Section 4.3 and Section 4.4) in Algorithm 1. Line 1 uses spikyball graph sampling algorithm to sample from graphs of similar origins to estimate a rough node degree of a given node embedding matrix. Line 2 applies Gumbel-Top-k trick on the fully connected probabilistic graph **P** to generate the initial seed graph structure \mathbf{A}^0 . Line 4 - 10 learn a distance function ϕ and a corresponding graph structure \mathbf{A}^t in each iteration (line 5-8). Upon termination of the learning process after a maximum iteration T, we obtain the recovered graph structure \mathbf{A}^T (i.e., \mathbf{G}_R).

Complexity Analysis. Our graph initialization generates a fully connected probabilistic graph using Equation 1. The time and space complexity of graph initialization are both $O(n^2)$. Our graph metric learning learns a global distance metric, leading to a $O(n^2)$ time complexity. The encoder of the GAE has a $O(n^2)$ time and space complexity, while the inner-product decoder of the GAE has a $O(n^2)$ time complexity. Overall, the time and space complexity of MNEMON are both $O(cn^2)$. Besides, Zhang et al. [81] has a time complexity of $O(n^4)$ while LinkTeller [71] has a space complexity of $O(n^3)$, which inevitably limit their scalability in the real world.

B Is MNEMON better than baselines?

Comparison Study on Other Datasets. We outline the comparison study results from the Citeseer, Actor and Facebook datasets in Table 7, Table 8 and Table 9 respectively. The node embedding size is fixed to 256.

Visual Explanation. We use bitmap images to further exemplify MNEMON's capability in recovery graphs from the node embedding matrices. Each adjacency matrix \mathbf{A} is represented as an image where the pixel at a coordinate (i,j) is blue if $\mathbf{A}_{ij}=1$ and white otherwise. In this way, bitmap images can give an overall impression of the graph topology, offering a straightforward qualitative visual assessment between the original graphs and the recovered graphs. We visualize the original graph and the recovered graphs by both baseline methods and MNEMON from the node embeddings generated by Node2Vec in Figure 7. As we can see in Figure 7, the graph structures recovered by direct recovery and invert embedding do not resemble the original graph.




Figure 7: Bitmap visualization of the recovered graphs by all baselines and MNEMON. MNEMON, visually, removes fair amount of false positive edges.

Table 7: Comparison of all baseline methods and MNEMON. We use the Citeseer dataset and the node embedding size is fixed to 256.

Graph	f		Edge 1	Metric	Global Metric			
Reocvery Method		Precision	Recall	F1	JDD	Frobenius Error	Triangle Error	Clustering Coef. Error
Direct Recovery	DW N2V NetSMF GCN	0.001±0.000 0.003±0.000 0.013±0.000 0.001±0.000	0.002±0.000 0.006±0.000 0.022±0.000 0.002±0.000	0.001±0.000 0.004±0.000 0.016±0.000 0.002±0.000	0.000±0.000 0.000±0.000 0.311±0.000 0.000±0.000	1.647±0.000 1.645±0.000 1.647±0.000 1.647±0.000	6.867±0.000 7.559±0.000 0.621±0.000 6.391±0.000	0.753 ± 0.000 0.759 ± 0.000 0.228 ± 0.000 0.653 ± 0.000
Invert Embedding	DW N2V NetSMF GCN	0.007±0.005 0.021±0.008 0.003±0.001 0.015±0.003	0.012±0.010 0.037±0.014 0.005±0.003 0.026±0.006	0.009±0.007 0.027±0.010 0.004±0.002 0.019±0.004	0.667±0.092 0.665±0.030 0.462±0.064 0.675±0.004	1.660±0.010 1.645±0.008 1.676±0.007 1.657±0.005	0.866±0.661 1.869±0.160 0.842±0.789 0.255±0.158	0.198 ± 0.008 0.148 ± 0.005 0.221 ± 0.011 0.188 ± 0.005
kNN Graph	DW N2V NetSMF GCN	0.338±0.000 0.350±0.000 0.431±0.000 0.343±0.000	0.483±0.000 0.500±0.000 0.616±0.000 0.490±0.000	0.398±0.000 0.412±0.000 0.507±0.000 0.403±0.000	0.234±0.000 0.237±0.000 0.225±0.000 0.225±0.000	1.210±0.000 1.196±0.000 1.094±0.000 1.204±0.000	4.528±0.000 3.399±0.000 3.420±0.000 3.201±0.000	0.421 ± 0.000 0.416 ± 0.000 0.488 ± 0.000 0.397 ± 0.000
MNEMON	DW N2V NetSMF GCN	0.403±0.002 0.445±0.001 0.530±0.002 0.414±0.003	0.555±0.005 0.575±0.002 0.672±0.001 0.529±0.002	0.467±0.003 0.502±0.001 0.592±0.001 0.465±0.001	0.617±0.011 0.506±0.007 0.461±0.005 0.527±0.011	1.125±0.003 1.069±0.002 0.961±0.002 1.105±0.004	1.877±0.075 1.734±0.039 2.001±0.056 1.467±0.057	0.341±0.009 0.357±0.005 0.432±0.003 0.330±0.004

At the same time, kNN graph can recover some graph topology. However, because kNN is not learning-based and uses a predefined distance function, we can see in Figure 7 that the recovered graph by kNN is noisy. Thanks to the learnable distance function and adaptive graph structure combination (see Section 4), MNEMON can reduce a reasonable amount of false edges and recover a better graph topology, consequently leading to better performance as shown in Table 3.

C How Effective is MNEMON?

Additional Experimental Results. We list additional experimental results using all four datasets in Table 10 and Ta-

ble 11.

D How does k Affect the Attack Performance

Additional Experimental Results. We list additional experimental results using all four datasets in Figure 8 and Figure 9.

Table 8: Comparison of all baseline methods and MNEMON. We use the Actor dataset and the node embedding size is fixed to 256.

Graph	f		Edge	Metric			Global Metr	ic
Reocvery Method	,	Precision	Recall	F1	JDD	Frobenius Error	Triangle Error	Clustering Coef. Error
Direct Recovery	DW N2V NetSMF GCN	0.001±0.000 0.003±0.000 0.013±0.000 0.001±0.000	0.002±0.000 0.006±0.000 0.022±0.000 0.002±0.000	0.001±0.000 0.004±0.000 0.016±0.000 0.002±0.000	0.000±0.000 0.000±0.000 0.311±0.000 0.000±0.000	1.647±0.000 1.645±0.000 1.647±0.000 1.647±0.000	6.867±0.000 7.559±0.000 0.621±0.000 6.391±0.000	0.753 ± 0.000 0.759 ± 0.000 0.228 ± 0.000 0.653 ± 0.000
Invert Embedding	DW N2V NetSMF GCN	0.007±0.005 0.021±0.008 0.003±0.001 0.015±0.003	0.012±0.010 0.037±0.014 0.005±0.003 0.026±0.006	0.009±0.007 0.027±0.010 0.004±0.002 0.019±0.004	0.667±0.092 0.665±0.030 0.462±0.064 0.675±0.004	1.660±0.010 1.645±0.008 1.676±0.007 1.657±0.005	0.866±0.661 1.869±0.160 0.842±0.789 0.255±0.158	0.198 ± 0.008 0.148 ± 0.005 0.221 ± 0.011 0.188 ± 0.005
kNN Graph	DW N2V NetSMF GCN	0.562±0.000 0.343±0.000 0.453±0.000 0.304±0.000	0.400±0.000 0.244±0.000 0.322±0.000 0.217±0.000	0.468±0.000 0.285±0.000 0.376±0.000 0.253±0.000	0.263±0.000 0.366±0.000 0.325±0.000 0.349±0.000	0.955±0.000 1.106±0.000 1.033±0.000 1.131±0.000	0.308±0.000 0.566±0.000 0.905±0.000 0.629±0.000	0.246 ± 0.000 0.211 ± 0.000 0.310 ± 0.000 0.202 ± 0.000
MNEMON	DW N2V NetSMF GCN	0.687±0.001 0.465±0.001 0.562±0.002 0.373±0.001	0.435±0.002 0.313±0.000 0.366±0.001 0.263±0.000	0.533±0.002 0.374±0.000 0.443±0.001 0.308±0.001	0.417±0.001 0.473±0.003 0.457±0.003 0.505±0.003	0.874±0.001 1.023±0.001 0.959±0.001 1.086±0.001	0.203±0.009 0.179±0.007 0.147±0.013 0.280±0.008	0.229 ± 0.002 0.176 ± 0.001 0.285 ± 0.002 0.153 ± 0.002

Table 9: Comparison of all baseline methods and MNEMON. We use the Facebook dataset and the node embedding size is fixed to 256.

Graph	f		Edge 1	Metric			Global Metr	ic
Reocvery Method	,	Precision	Recall	F1	JDD	Frobenius Error	Triangle Error	Clustering Coef. Error
Direct Recovery	DW N2V NetSMF GCN	0.001±0.000 0.003±0.000 0.013±0.000 0.001±0.000	0.002±0.000 0.006±0.000 0.022±0.000 0.002±0.000	0.001±0.000 0.004±0.000 0.016±0.000 0.002±0.000	0.000±0.000 0.000±0.000 0.311±0.000 0.000±0.000	1.647±0.000 1.645±0.000 1.647±0.000 1.647±0.000	6.867±0.000 7.559±0.000 0.621±0.000 6.391±0.000	0.753±0.000 0.759±0.000 0.228±0.000 0.653±0.000
Invert Embedding	DW N2V NetSMF GCN	0.007±0.005 0.021±0.008 0.003±0.001 0.015±0.003	0.012±0.010 0.037±0.014 0.005±0.003 0.026±0.006	0.009±0.007 0.027±0.010 0.004±0.002 0.019±0.004	0.667±0.092 0.665±0.030 0.462±0.064 0.675±0.004	1.660±0.010 1.645±0.008 1.676±0.007 1.657±0.005	0.866±0.661 1.869±0.160 0.842±0.789 0.255±0.158	0.198±0.008 0.148±0.005 0.221±0.011 0.188±0.005
kNN Graph	DW N2V NetSMF GCN	0.429±0.000 0.460±0.000 0.444±0.000 0.322±0.000	0.442±0.000 0.474±0.000 0.457±0.000 0.331±0.000	0.436±0.000 0.467±0.000 0.450±0.000 0.327±0.000	0.189±0.000 0.172±0.000 0.291±0.000 0.183±0.000	1.070±0.000 1.040±0.000 1.056±0.000 1.169±0.000	0.281±0.000 0.046±0.000 0.006±0.000 0.072±0.000	0.071±0.000 0.040±0.000 0.040±0.000 0.087±0.000
MNEMON	DW N2V NetSMF GCN	0.441±0.001 0.468±0.000 0.454±0.001 0.342±0.001	0.471±0.001 0.487±0.001 0.502±0.002 0.364±0.001	0.456±0.001 0.477±0.001 0.476±0.001 0.352±0.001	0.519±0.006 0.444±0.002 0.457±0.002 0.371±0.004	1.061±0.001 1.033±0.001 1.050±0.001 1.157±0.001	0.494±0.002 0.545±0.001 0.424±0.007 0.452±0.002	0.077 ± 0.001 0.090 ± 0.001 0.081 ± 0.001 0.056 ± 0.001

Table 10: The performance results of MNEMON using all four datasets. We fix the node embedding size to 128. We show the relative improvement scores in edge metrics to demonstrate to what extent MNEMON can relatively improve from kNN graph. We add a positive sign (+) next to the relative improvement score to highlight the improvement. We also show the relative error reduction scores in global metrics to demonstrate to what extent MNEMON can relatively reduce errors incurred by kNN graph. We add a negative sign (-) next to the relative error reduction score to highlight the difference.

Dataset	f		Edge !	Metrics		Glo	bal Metrics (Relative Er	ror)
Dataset	,	Precision	Recall	F1	Deg. Dist.	Frobenius Error	Triangle Error	Clus. Coef. Error
Cora	DW	0.570±0.001 (+0.224)	0.633±0.005 (+0.109)	0.600±0.003 (+0.170)	0.786±0.012 (+1.311)	0.919±0.002 (-0.122)	0.987±0.039 (-1.376)	0.232±0.004 (-0.055)
	N2V	0.504±0.002 (+0.276)	0.554±0.003 (+0.145)	0.528±0.002 (+0.214)	0.726±0.010 (+1.205)	0.995±0.002 (-0.126)	0.997±0.035 (-1.140)	0.223±0.004 (-0.050)
	NetSMF	0.516±0.003 (+0.225)	0.579±0.003 (+0.121)	0.545±0.003 (+0.175)	0.716±0.006 (+1.188)	0.982±0.004 (-0.111)	1.453±0.033 (-1.139)	0.297±0.004 (-0.047)
	GCN	0.456±0.002 (+0.231)	0.499±0.003 (+0.101)	0.476±0.001 (+0.170)	0.763±0.007 (+1.305)	1.048±0.003 (-0.101)	0.806±0.039 (-1.235)	0.213±0.004 (-0.049)
Citeseer	DW	0.488±0.001 (+0.204)	0.639±0.003 (+0.104)	0.553±0.001 (+0.159)	0.576±0.006 (+1.482)	1.016±0.000 (-0.112)	1.772±0.058 (-2.132)	0.362±0.006 (-0.071)
	N2V	0.453±0.001 (+0.270)	0.583±0.003 (+0.141)	0.510±0.002 (+0.214)	0.500±0.004 (+1.155)	1.058±0.001 (-0.129)	1.715±0.021 (-1.581)	0.353±0.003 (-0.057)
	NetSMF	0.458±0.001 (+0.197)	0.599±0.003 (+0.095)	0.519±0.001 (+0.154)	0.455±0.006 (+1.014)	1.053±0.002 (-0.102)	2.319±0.064 (-1.266)	0.429±0.003 (-0.055)
	GCN	0.404±0.003 (+0.201)	0.519±0.001 (+0.079)	0.454±0.002 (+0.146)	0.518±0.005 (+1.270)	1.117±0.004 (-0.095)	1.404±0.025 (-1.641)	0.316±0.005 (-0.061)
Actor	DW	0.678±0.002 (+0.225)	0.424±0.001 (+0.078)	0.521±0.002 (+0.133)	0.400±0.002 (+0.564)	0.882±0.001 (-0.079)	0.219±0.009 (-0.032)	0.223±0.002 (-0.016)
	N2V	0.461±0.001 (+0.348)	0.307±0.001 (+0.264)	0.369±0.001 (+0.298)	0.473±0.003 (+0.296)	1.026±0.001 (-0.081)	0.209±0.006 (-0.330)	0.168±0.002 (-0.037)
	NetSMF	0.498±0.002 (+0.221)	0.330±0.002 (+0.135)	0.397±0.002 (+0.172)	0.472±0.003 (+0.409)	1.001±0.001 (-0.062)	0.148±0.013 (-0.828)	0.270±0.002 (-0.033)
	GCN	0.374±0.002 (+0.220)	0.263±0.001 (+0.206)	0.309±0.001 (+0.212)	0.512±0.002 (+0.481)	1.085±0.001 (-0.044)	0.303±0.007 (-0.263)	0.150±0.001 (-0.043)
Facebook	DW	0.448±0.001 (+0.025)	0.472±0.002 (+0.049)	0.460±0.001 (+0.038)	0.509±0.004 (+1.649)	1.053±0.001 (-0.010)	0.500±0.003 (+0.337)	0.069±0.001 (-0.003)
	N2V	0.450±0.001 (+0.018)	0.467±0.001 (+0.025)	0.458±0.001 (+0.023)	0.421±0.003 (+1.407)	1.050±0.001 (-0.008)	0.535±0.002 (+0.506)	0.074±0.001 (+0.031)
	NetSMF	0.435±0.001 (+0.026)	0.474±0.002 (+0.087)	0.454±0.001 (+0.055)	0.415±0.004 (+0.579)	1.069±0.000 (-0.007)	0.417±0.004 (+0.407)	0.047±0.001 (-0.002)
	GCN	0.341±0.001 (+0.062)	0.364±0.002 (+0.100)	0.352±0.001 (+0.080)	0.370±0.004 (+0.851)	1.157±0.001 (-0.013)	0.450±0.004 (+0.377)	0.059±0.000 (-0.026)

Table 11: The performance results of MNEMON using all four datasets. We fix the node embedding size to 64. We show the relative improvement scores in edge metrics to demonstrate to what extent MNEMON can relatively improve from kNN graph. We add a positive sign (+) next to the relative improvement score to highlight the improvement. We also show the relative error reduction scores in global metrics to demonstrate to what extent MNEMON can relatively reduce errors incurred by kNN graph. We add a negative sign (-) next to the relative error reduction score to highlight the difference.

Dataset	f		Edge !	Metrics	Glo	bal Metrics (Relative Er	ror)	
Dutaset	,	Precision	Recall	F1	Deg. Dist.	Frobenius Error	Triangle Error	Clus. Coef. Error
Cora	DW	0.596±0.003 (+0.231)	0.636±0.003 (+0.073)	0.615±0.001 (+0.155)	0.729±0.015 (+1.195)	0.892±0.002 (-0.128)	0.918±0.044 (-1.137)	0.248±0.003 (-0.046)
	N2V	0.467±0.002 (+0.270)	0.524±0.002 (+0.162)	0.494±0.002 (+0.220)	0.730±0.007 (+1.231)	1.036±0.002 (-0.115)	1.046±0.017 (-1.156)	0.218±0.004 (-0.048)
	NetSMF	0.441±0.003 (+0.166)	0.532±0.002 (+0.149)	0.482±0.001 (+0.159)	0.727±0.009 (+1.237)	1.069±0.004 (-0.071)	2.193±0.169 (-0.512)	0.295±0.004 (-0.045)
	GCN	0.426±0.003 (+0.217)	0.470±0.002 (+0.097)	0.447±0.003 (+0.161)	0.777±0.007 (+1.306)	1.079±0.004 (-0.091)	0.684±0.032 (-1.235)	0.174±0.004 (-0.054)
Citeseer	DW	0.541±0.002 (+0.221)	0.681±0.003 (+0.074)	0.603±0.002 (+0.155)	0.516±0.002 (+1.295)	0.947±0.003 (-0.131)	1.669±0.036 (-1.618)	0.388±0.004 (-0.065)
	N2V	0.423±0.004 (+0.275)	0.557±0.004 (+0.172)	0.481±0.004 (+0.231)	0.479±0.004 (+1.028)	1.096±0.005 (-0.121)	1.763±0.038 (-1.469)	0.335±0.004 (-0.055)
	NetSMF	0.362±0.004 (+0.049)	0.567±0.003 (+0.153)	0.442±0.004 (+0.091)	0.462±0.004 (+1.000)	1.197±0.007 (-0.005)	7.019±0.920 (+3.300)	0.438±0.004 (-0.039)
	GCN	0.389±0.003 (+0.196)	0.501±0.002 (+0.077)	0.438±0.002 (+0.143)	0.524±0.010 (+1.297)	1.134±0.004 (-0.091)	1.435±0.030 (-1.586)	0.314±0.002 (-0.054)
Actor	DW	0.660±0.002 (+0.224)	0.414±0.001 (+0.080)	0.508±0.001 (+0.135)	0.401±0.002 (+0.520)	0.894±0.001 (-0.078)	0.212±0.003 (-0.019)	0.224±0.003 (-0.013)
	N2V	0.447±0.001 (+0.330)	0.298±0.001 (+0.246)	0.357±0.001 (+0.279)	0.468±0.001 (+0.283)	1.035±0.001 (-0.076)	0.231±0.010 (-0.252)	0.159±0.002 (-0.034)
	NetSMF	0.402±0.002 (+0.186)	0.274±0.001 (+0.135)	0.326±0.001 (+0.155)	0.480±0.005 (+0.370)	1.065±0.001 (-0.044)	0.238±0.089 (-0.566)	0.228±0.002 (-0.031)
	GCN	0.338±0.001 (+0.211)	0.239±0.001 (+0.209)	0.280±0.001 (+0.209)	0.523±0.003 (+0.460)	1.109±0.001 (-0.038)	0.325±0.008 (-0.192)	0.135±0.002 (-0.045)
Facebook	DW	0.445±0.001 (+0.025)	0.468±0.001 (+0.047)	0.456±0.001 (+0.037)	0.473±0.004 (+1.688)	1.056±0.001 (-0.010)	0.498±0.002 (+0.435)	0.062±0.001 (-0.007)
	N2V	0.418±0.001 (+0.022)	0.436±0.001 (+0.033)	0.427±0.001 (+0.028)	0.385±0.003 (+1.253)	1.082±0.001 (-0.007)	0.495±0.001 (+0.422)	0.044±0.001 (-0.016)
	NetSMF	0.427±0.001 (+0.082)	0.522±0.002 (+0.283)	0.470±0.001 (+0.172)	0.394±0.002 (+0.958)	1.085±0.001 (-0.018)	0.038±0.022 (+0.014)	0.052±0.000 (-0.029)
	GCN	0.342±0.001 (+0.062)	0.371±0.002 (+0.122)	0.356±0.001 (+0.093)	0.372±0.003 (+0.876)	1.159±0.001 (-0.010)	0.432±0.003 (+0.350)	0.059±0.001 (-0.025)




Figure 8: F1 scores and relative average clustering coefficient error scores of MNEMON given all four datasets. We fix the node embedding size to 128. The estimated average node degree of Cora, Citeseer and Actor datasets is 5. The estimated average node degree of Facebook is 46.




Figure 9: F1 scores and relative average clustering coefficient error scores of MNEMON given all four datasets. We fix the node embedding size to 64. The estimated average node degree of Cora, Citeseer and Actor datasets is 5. The estimated average node degree of Facebook is 46.