

AliGraph: A Comprehensive Graph Neural Network Platform

Rong Zhu, Kun Zhao, Hongxia Yang^{*}, Wei Lin, Chang Zhou, Baole Ai, Yong Li, Jingren Zhou

Alibaba Group

{red.zr, kun.zhao, yang.yhx, weilin.lw, ericzhou.zc, jiufeng.ly, jingren.zhou}@alibaba-inc.com

ABSTRACT

An increasing number of machine learning tasks require dealing with large graph datasets, which capture rich and complex relationship among potentially billions of elements. Graph Neural Network (GNN) becomes an effective way to address the graph learning problem by converting the graph data into a low dimensional space while keeping both the structural and property information to the maximum extent and constructing a neural network for training and referencing. However, it is challenging to provide an efficient graph storage and computation capabilities to facilitate GNN training and enable development of new GNN algorithms. In this paper, we present a comprehensive graph neural network system, namely *AliGraph*, which consists of distributed graph storage, optimized sampling operators and runtime to efficiently support not only existing popular GNNs but also a series of in-house developed ones for different scenarios. The system is currently deployed at Alibaba to support a variety of business scenarios, including product recommendation and personalized search at Alibaba's E-Commerce platform. By conducting extensive experiments on a real-world dataset with 492.90 million vertices, 6.82 billion edges and rich attributes, *AliGraph* performs an order of magnitude faster in terms of graph building (5 minutes vs hours reported from the state-of-the-art PowerGraph platform). At training, *AliGraph* runs 40%-50% faster with the novel caching strategy and demonstrates around 12 times speed up with the improved runtime. In addition, our in-house developed GNN models all showcase their statistically significant superiorities in terms of both effectiveness and efficiency (e.g., 4.12%–17.19% lift by F1 scores).

PVLDB Reference Format:

Zhu, R., Zhao, K., Yang, H., Lin, W., Zhou, C., Ai, B. and Zhou, J., AliGraph: A Comprehensive Graph Neural Network Platform. *PVLDB*, 12 (12): 2094 - 2105, 2019.

DOI: <http://doi.org/10.14778/3352063.3352127>

1. INTRODUCTION

^{*}Corresponding author.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 12, No. 12

ISSN 2150-8097.

DOI: <http://doi.org/10.14778/3352063.3352127>

As a sophisticated model, graph has been widely used to model and manage data in a wide variety of real-world applications. Typical examples include social networks [24], physical systems [1], biological networks [18], knowledge graphs [23]. Graph analytics, which explore underlying insights hidden in graph data, have drawn significant research attention in the last decade. They have been witnessed to play important roles in numerous areas, i.e., node classification [2], link prediction [38], graph clustering [4], recommendation [54], among many others.

As conventional graph analytic tasks often suffer from high computation and space costs [13, 24], a new paradigm, called *graph embedding* (GE), paves an efficient yet effective way to address such problems. Specifically, GE converts the graph data into a low-dimensional space such that the structural and content information in the graph can be preserved to the maximum extent. After that, the generated embeddings are fed as features into the downstream machine learning tasks. Furthermore, by incorporating with deep learning techniques, *graph neural networks* (GNN) are proposed by integrating GE with *convolutional neural network* (CNN) [32, 11, 27, 25]. In CNN, shared weights and multi-layered structure are applied to enhance its learning power [33]. And graphs are the most typical locally connected structures, with shared weights to reduce the computational cost and the multi-layer structure being the key to deal with hierarchical patterns while capturing features of various sizes. GNNs find such generalizations of CNNs to graphs. Thus, GNN not only embraces the flexibility of GE but also showcases its superiority in terms of both effectiveness and robustness with generalizations of CNNs.

Challenges. In the literature, considerable research efforts have been devoted in developing GE and GNN algorithms. These works mainly concentrate on simple graphs with no or little auxiliary information. However, the rising of big data and complex systems reveal new insights in graph data. As a consensus [13, 24, 5, 21], the vast majority of graph data related to real-world commercial scenarios exhibits four properties, namely *large-scale*, *heterogeneous*, *attributed* and *dynamic*. For example, nowadays e-commerce graphs often contain billions of vertices and edges with various types and rich attributes, and quickly evolve over time. These properties bring great challenges for embedding and representing graph data as follows:

- The core steps in GNN are particularly optimized for grid structures, such as images, but not feasible for graphs in irregular Euclidean space. Thus, existing GNN methods can not scale on real-world graphs with exceedingly large sizes. ***The first problem is how to improve the time and space efficiencies of GNN on large-scale graphs?***
- Different types of objects characterize the data from multiple perspectives. They provide richer information but increase the

difficulty to map the graph information into a *singleton* space. Thus, *the second problem is how to elegantly integrate the heterogeneous information to be an unified embedding result?*

- The attribute information can further enhance the power of the embedding results and make inductive GE possible [13, 24, 5, 21]. Without considering attribute information, the algorithms can only consider the transductive settings and ignore the need for predicting unseen instances. However, the topological structure information and unstructured attribute information are usually presented in two different spaces. Thus, *the third problem is how to unify them to define the information to be preserved?*
- As GNN suffers from low efficiency, recomputing the embedding results from scratch with respect to structural and contextual updates are expensive. Thus, *the fourth problem is how to design efficient incremental GNN methods on dynamic graphs?*

Contributions. To tackle with the above challenges, considerable research efforts have been devoted to design efficient yet effective GNN methods. In Table 1, we categorize a series of popular GE and GNN models according to the aspects that they focus on, as well as our in-house developed models shaded in yellow. As shown, most existing methods concentrate on one or two properties at the same time. However, real-world commercial data usually faces more challenges. To mitigate this situation, in this paper, we present a comprehensive and systemic solution to GNN. We design and implement a platform, called *AliGraph*, which provides both system and algorithms to tackle more practical problems that exhibit the four summarized arising problems, to well support a variety of GNN methods and applications. The main contributions are summarized as follows.

System. In the underlying components of *AliGraph*, we build a system to support GNN algorithms and applications. The system architecture is abstracted from general GNN methods, which consists of a *storage* layer, a *sampling* layer and an *operator* layer. Specifically, the storage layer applies three novel techniques, namely structural and attributed specific storage, graph partition and caching neighbors of some important vertices, to store large-scale raw data to fulfill the fast data access requirements of high-level operations and algorithms. The sampling layer optimizes the key sampling operation in GNN methods. We categorize sampling methods into three classes, namely TRAVERSE, NEIGHBORHOOD and NEGATIVE sampling, and propose lock-free methods to perform sampling operations in distributed environment. The operator layer provides optimized implementation of two common applied operators in GNN algorithms, namely AGGREGATE and COMBINE. We apply a cache strategy to store some intermediate results to accelerate the computation process. These components are co-designed and co-optimized to make the whole system effective and scalable.

Algorithms. The system provides a flexible interface to design GNN algorithms. We show that all existing GNN methods can be easily implemented upon our system. Besides, we also in-house developed several new GNNs for practical requirement and detail six here. As illustrated in Table 1, our in-house developed methods are shaded in yellow and each of them are more flexible and practical to deal with real-world problems.

Evaluation. Our *AliGraph* platform is practically deployed in the Alibaba company. The experimental results verify its effectiveness and efficiency from both the system and algorithm aspects. As shown in Figure 1, our in-house developed GNN models on the *AliGraph* platform improves the normalized evaluation metrics by 4.12%–17.19%. The data are collected from Alibaba’s e-commerce

Table 1: The property of different methods.

| Category | Method | Heterogeneous | | Attributed | Dynamic | Large-Scale |
|-------------------------|------------------|---------------|------|------------|---------|-------------|
| | | Node | Edge | | | |
| Classic Graph Embedding | DeepWalk | ✗ | ✗ | ✗ | ✗ | ✗ |
| | Node2Vec | ✗ | ✗ | ✗ | ✗ | ✗ |
| | LINE | ✗ | ✗ | ✗ | ✗ | ✗ |
| | NetMF | ✗ | ✗ | ✗ | ✗ | ✗ |
| | TADW | ✗ | ✗ | ✓ | ✗ | ✗ |
| | LANE | ✗ | ✗ | ✓ | ✗ | ✗ |
| | ASNE | ✗ | ✗ | ✓ | ✗ | ✗ |
| | DANE | ✗ | ✗ | ✓ | ✗ | ✗ |
| | ANRL | ✗ | ✗ | ✓ | ✗ | ✗ |
| | PTE | ✗ | ✓ | ✗ | ✗ | ✗ |
| | Methpath2Vec | ✗ | ✓ | ✗ | ✗ | ✗ |
| | HERec | ✗ | ✓ | ✗ | ✗ | ✗ |
| | HNE | ✗ | ✗ | ✗ | ✗ | ✗ |
| | PMNE | ✗ | ✓ | ✓ | ✗ | ✗ |
| | MVE | ✗ | ✓ | ✓ | ✗ | ✗ |
| MNE | ✗ | ✓ | ✓ | ✗ | ✗ | |
| Mvn2Vec | ✗ | ✓ | ✓ | ✗ | ✗ | |
| GNN | Structural2Vec | ✗ | ✗ | ✓ | ✗ | ✗ |
| | GCN | ✗ | ✗ | ✓ | ✗ | ✗ |
| | FastGCN | ✗ | ✗ | ✓ | ✗ | ✗ |
| | AS-GCN | ✗ | ✗ | ✓ | ✗ | ✗ |
| | GraphSAGE | ✗ | ✗ | ✓ | ✗ | ✗ |
| | HEP | ✓ | ✓ | ✓ | ✗ | ✗ |
| | AHEP | ✓ | ✓ | ✓ | ✗ | ✓ |
| | GATNE | ✓ | ✓ | ✓ | ✗ | ✓ |
| | Mixture GNN | ✓ | ✓ | ✓ | ✗ | ✗ |
| | Hierarchical GNN | ✓ | ✓ | ✓ | ✗ | ✗ |
| Bayesian GNN | ✗ | ✓ | ✓ | ✗ | ✗ | |
| Evolving GNN | ✗ | ✓ | ✓ | ✓ | ✗ | |

platform Taobao and we contribute this dataset to the community to nourish further development¹.




Figure 1: Comparison of normalized evaluation metric on the effectiveness of different methods. The red text indicates the lifting range of each in-house developed method w.r.t. its competitors.

2. PRELIMINARIES

In this section, we introduce the basic concepts and formalize the graph embedding problem. The symbols and notations frequently used throughout this paper are summarized in Table 2. We start with the acyclic, simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ where \mathcal{V} and \mathcal{E} represent the set of vertices and edges, respectively; and $\mathcal{W} : \mathcal{E} \rightarrow \mathbb{R}^+$ is a function assigning each edge $(u, v) \in \mathcal{E}$ a weight $\mathcal{W}(u, v)$ indicating the strength of relationships between vertex u and v . Let $n = |\mathcal{V}|$ and $m = |\mathcal{E}|$ denote the number of vertices and edges in \mathcal{G} , respectively. Notice that, the graph \mathcal{G} can be either directed or undirected. If \mathcal{G} is directed, (u, v) and (v, u) represent two different edges and may have different weights; otherwise, (u, v) and (v, u) are the same edge and we have $\mathcal{W}(u, v) = \mathcal{W}(v, u)$. For each vertex u , we use $Nb(u)$ to denote the set of its (in and out) neighbors.

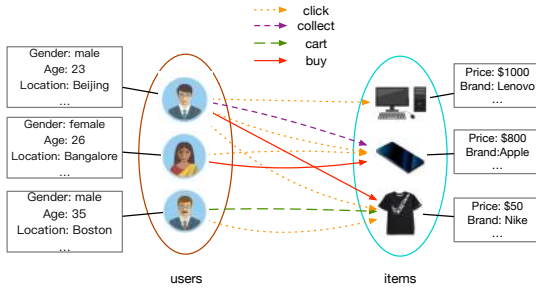
Attributed Heterogeneous Graph. To comprehensively characterize the real-world commercial data, the practical graphs often contain rich content information, e.g., multiple types of vertices, multiple types of edges, attributes and etc. Thus, we further define

¹<https://tianchi.aliyun.com/dataset/dataDetail?dataId=9716>.

Table 2: Summarization of symbols and notations.

| Symbols or Notations | Description |
|----------------------------------|--|
| \mathcal{G} | graph or attributed heterogeneous graph |
| $\mathcal{G}^{(t)}$ | the graph at the timestamp t |
| $\mathcal{V}(\mathcal{V}^{(t)})$ | vertex set (at timestamp t) |
| $\mathcal{E}(\mathcal{E}^{(t)})$ | edge set (at timestamp t) |
| n, m | number of vertices and edges |
| $\mathcal{W}(\mathcal{W}^{(t)})$ | edge weight assigning function (at timestamp t) |
| \mathcal{T}_V | vertex type mapping function |
| \mathcal{T}_E | edge type mapping function |
| \mathcal{A}_V | vertex attributes mapping function |
| \mathcal{A}_E | edge attributes mapping function |
| $\mathbf{x}_{v,i}$ | the i -th feature vector of vertex v |
| $\mathbf{w}_{e,i}$ | the i -th feature vector of edge e |
| d | the embedding dimension |
| $Nb(v)$ | neighbors set of vertex v |
| \mathbf{h}_v | the embedding vector of vertex v |
| $\mathbf{h}_{v,c}$ | the embedding vector of vertex v w.r.t. type c |
| $D_i^{(t)}(v)$ | number of t -hop in-neighbors of v |
| $D_o^{(t)}(v)$ | number of t -hop out-neighbors of v |
| $Imp^{(t)}(v)$ | importance of vertex v |

the *Attributed Heterogeneous Graph* (AHG). An AHG \mathcal{G} is a tuple $(\mathcal{V}, \mathcal{E}, \mathcal{W}, \mathcal{T}_V, \mathcal{T}_E, \mathcal{A}_V, \mathcal{A}_E)$ where \mathcal{V}, \mathcal{E} and \mathcal{W} have the same meaning as the simple graphs. $\mathcal{T}_V : \mathcal{V} \rightarrow \mathcal{F}_V$ and $\mathcal{T}_E : \mathcal{E} \rightarrow \mathcal{F}_E$ represent the vertex type and edge type mapping functions, where \mathcal{F}_V and \mathcal{F}_E are the set of vertex types and edge types, respectively. To ensure the heterogeneity, we request $|\mathcal{F}_V| \geq 2$ and/or $|\mathcal{F}_E| \geq 2$. \mathcal{A}_V and \mathcal{A}_E are two functions assigning each vertex $v \in \mathcal{V}$ and each edge $e \in \mathcal{E}$ some feature vectors representing its attributes. We denote the i -th feature vector of vertex v and edge e as $\mathbf{x}_{v,i}$ and $\mathbf{w}_{e,i}$, respectively. An example of AHG is shown in Figure 2, which contains two types of vertices, namely users and items, and four types of edges connecting them.


Figure 2: Illustrative example of AHG with multiple types of edges, nodes and rich attributes.

Dynamic Graph. Real-world graphs usually evolve with time. Given a time interval $[1, T]$, a *dynamic graph* is a series of graphs $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots, \mathcal{G}^{(T)}$. For each $1 \leq t \leq T$, $\mathcal{G}^{(t)}$ can be a simple graph or an AHG. For ease of notation, we add a superscript (t) to represent the corresponding state of the objects at timestamp t . For example, $\mathcal{V}^{(t)}$ and $\mathcal{E}^{(t)}$ represent the vertex set and edge set of graph $\mathcal{G}^{(t)}$, respectively.

Problem Definition. Given an input graph \mathcal{G} , which is a simple graph or an AHG, and a predefined number $d \in \mathbb{N}$ on the dimension of embedding where $d \ll |\mathcal{V}|$, the embedding problem is to convert the graph \mathcal{G} into the d -dimensional space such that the graph property is preserved as much as possible. GNN is a special kind of graph embedding method, which learns the embedding results by applying neural networks on graphs. Notice that, in this paper, we concentrate on the vertex-level embedding. That is, the embedding output is a d -dimensional vector $\mathbf{h}_v \in \mathbb{R}^d$ for each vertex $v \in \mathcal{V}$. In our future work as discussed in Section 7, we will also consider the embedding on edges, subgraphs or even the whole graph.

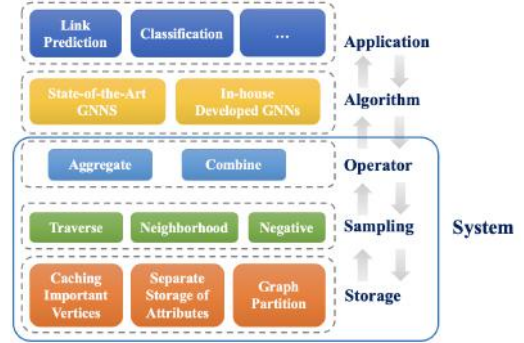
Algorithm 1: GNN Framework

Input: network \mathcal{G} , embedding dimension $d \in \mathbb{N}$, a vertex feature \mathbf{x}_v for each vertex $v \in \mathcal{V}$ and the maximum hops of neighbors $k_{max} \in \mathbb{N}$.
Output: embedding result \mathbf{h}_v of each vertex $v \in \mathcal{V}$

- 1 $\mathbf{h}_v^{(0)} \leftarrow \mathbf{x}_v$
- 2 **for** $k \leftarrow 1$ to k_{max} **do**
- 3 **for each** vertex $v \in \mathcal{V}$ **do**
- 4 $S_v \leftarrow \text{SAMPLE}(Nb(v))$
- 5 $\mathbf{h}'_v \leftarrow \text{AGGREGATE}(\mathbf{h}_u^{(k-1)}, \forall u \in S)$
- 6 $\mathbf{h}_v^{(k)} \leftarrow \text{COMBINE}(\mathbf{h}_v^{(k-1)}, \mathbf{h}'_v)$
- 7 normalize all embedding vectors $\mathbf{h}_v^{(k)}$ for all $v \in \mathcal{V}$
- 8 $\mathbf{h}_v \leftarrow \mathbf{h}_v^{(k_{max})}$ for all $v \in \mathcal{V}$ return \mathbf{h}_v as the embedding result for all $v \in \mathcal{V}$

3. SYSTEM

In our *AliGraph* platform, whose architecture is shown in Figure 3, we design and implement an underlying system (marked in blue square) to well support high-level GNN algorithms and applications. The details of this system will be described in this section. To start with, in Section 3.1, we abstract a general framework of GNNs to explain why our system is designed in this way. Sections 3.2 to 3.5 introduce the design and implementation details of each key component in the system.


Figure 3: Architecture of the *AliGraph* system.

3.1 Framework of GNN Algorithms

In this subsection, we abstract a general framework to GNN algorithms. A series of classic GNNs such as Structure2Vec [46], GCN [32], FastGCN [11], AS-GCN [27] and GraphSAGE [24] can be characterized by instantiating the operators in the framework. The input of the GNN framework includes a graph \mathcal{G} , the embedding dimension $d \in \mathbb{N}$, a vertex feature \mathbf{x}_v for each vertex $v \in \mathcal{V}$ and the maximum hops of neighbors $k_{max} \in \mathbb{N}$. The output of the GNN is an embedding vector $\mathbf{h}_v \in \mathbb{R}^d$ for each vertex $v \in \mathcal{V}$ and will be fed into the downstream machine learning tasks, such as classification, link prediction and etc.

The GNN framework is described in Algorithm 1. At the very beginning, the vertex embedding $\mathbf{h}_v^{(0)}$ of vertex v is initialized to be equal to the input attribute vector \mathbf{x}_v . Then, at each k , each vertex v aggregates the embeddings of its neighbors to update the embedding of itself. Specifically, we apply the SAMPLE function to fetch a subset S of vertices based on the neighbor set $Nb(v)$ of vertex v , aggregate the embeddings of all vertices $u \in S$ by the AGGREGATE function to obtain a vector \mathbf{h}'_v , and combine \mathbf{h}'_v with $\mathbf{h}_v^{(k-1)}$ to generate the embedding vector $\mathbf{h}_v^{(k)}$ by the COMBINE function. After processing all vertices, the embedding vectors are normalized. Finally, after k_{max} hops, $\mathbf{h}_v^{(k_{max})}$ is returned as the embedding result \mathbf{h}_v of vertex v .

System Architecture. Based on the GNN framework described above, we naturally construct the system architecture of the *AliGraph* platform, as shown in Figure 3. Notice that, the platform consists of five layers on the whole, where the three underlying layers form the system to support the *algorithm* layer and the *application* layer. Inside the system, the *storage* layer organizes and stores different kinds of raw data to fulfill the fast data access requirements of high-level operations and algorithms. Upon this, by Algorithm 1, we find that three main operators, namely SAMPLE, AGGREGATE and COMBINE, play important roles in various GNN algorithms. Among them, the SAMPLE operator lays foundation for AGGREGATE and COMBINE since it directly controls the scope of information to be processed by them. Therefore, we design the *sampling* layer to access the storage for fast and accurate generation of training samples. Above it, the *operator* layer specifically optimizes the AGGREGATE and COMBINE functions. On top of the system, the GNN algorithms can be constructed in the *algorithm* layer to serve real-world tasks in the *application* layer.

3.2 Storage

In this subsection, we discuss how to store and organize the raw data. Notice that, the space cost to store the real-world graphs is very large. Common e-commerce graphs can contain tens of billions of nodes and hundreds of billions of edges with storage cost over 10TB easily. The large graph size brings great challenges for efficient graph access, especially in a distributed environment of clusters. To well support the high-level operators and algorithms, we apply the following three strategies in the *storage* layer of *AliGraph*.

Graph Partition. Our *AliGraph* platform is build on a distributed environment, thus the whole graph is divided and separately stored in different workers. The goal of graph partition is to minimize the number of crossing edges whose endpoints are in different workers. To this end, literature work has proposed a series of algorithms. In our system, as recommended in [17], we implement four built-in graph partition algorithms: 1) METIS [30]; 2) Vertex cut and edge cut partitions [20]; 3) 2-D partition [3]; and 4) Streaming-style partition strategy [48]. These four algorithms are suitable to different circumstances. In short, the METIS method is specialized in processing sparse graphs; the vertex and edge cut method performs much better on dense graphs; 2-D partition is often used when the number of workers is fixed; and streaming-style partition method are often applied on graphs with frequently edge updates. Users can choose the best partition strategy based on their own needs, moreover, they can also implement other graph partition algorithms as plugins in the system.

In Algorithm 2, lines 1–4 present the interface of graph partition. For each edge $e \in \mathcal{E}$, the general function ASSIGN in line 4 computes which worker e will be in based on its endpoints.

Separate Storage of Attributes. Notice that, for AHGs, we need to store both the structural and attributes of the partitioned graphs in each worker. The structural information of graph can be simply stored by an adjacency table. That is, for each vertex v , we store its neighbor set $Nb(v)$. Whereas, for the attributes on both vertices and edges, it is inadvisable to store them together in the adjacency table. The reasons are two-fold: 1) Attributes often cost more spaces. For example, the space cost to store a vertex id is at most 8 bytes while the attributes on a vertex may range from 0.1KB to 1KB. 2) Attributes among different vertices or edges have largely overlaps. For example, many vertices may have the same tag “man” indicating its gender. Therefore, it is more reasonable to separately store attributes.

In our system, we do so by building two indices I_V and I_E to store the attributes on vertices and edges, respectively. Each entry in

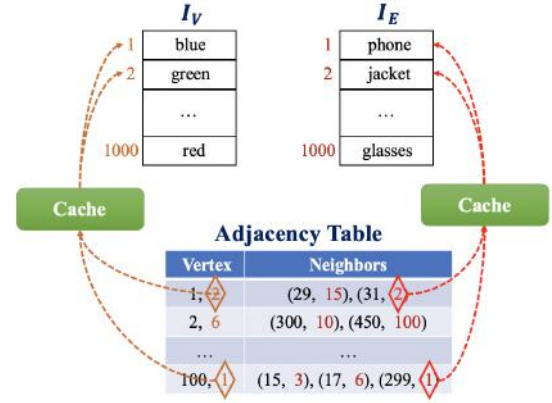


Figure 4: Index structure of graph storage.

I_V (I_E resp.) is a unique attribute associated on vertex (edge resp.). As illustrated in Figure 4, in the adjacency table, for each vertex u , we store the index of attribute $\mathcal{A}_V(u)$ in I_V , and for each edge (u, v) , we also store the index of attribute $\mathcal{A}_E(u, v)$ in I_E . Let N_D and N_L be the average number of neighbors and average length of attributes. Let N_A be the number of distinct attributes on vertices and edges. Obviously, our separate storage strategy decreases the space cost from $O(nN_DN_L)$ to $O(nN_D + N_A N_L)$.

Undoubtedly, separate storage of the attributes will increase the access time for retrieving the attributes. On average, each vertex will need to access the index I_E at most N_D times to collect the attributes of all of its neighbors. To mitigate this, we add two cache components to reside the frequently accessed items in I_V and I_E , respectively. We adopt the least recently used (LRU) replacing strategy [12] in each cache.

Caching Neighbors of Important Vertices. In each worker, we further propose a method to locally cache the neighbors of some important vertices to reduce the communication cost. The intuitive idea is that if a vertex v is frequently accessed by other vertices, we can store v ’s out-neighbors in each partition it occurs. By doing this, the visiting cost of other vertices to their neighbors via v can be greatly reduced. However, if the number of neighbors of v is large, storing multiple copies of v ’s neighbors will also incur huge storage cost. To make a better trade-off, we define a metric to evaluate the *importance* of each vertex, which decides whether a vertex is worth to cache or not.

Let $D_i^{(k)}(v)$ and $D_o^{(k)}(v)$ denote the number of k -hop in and out-neighbors of the vertex v , respectively. Certainly, $D_i^{(k)}(v)$ and $D_o^{(k)}(v)$ can measure the benefit and cost of caching the out-neighbors of v , respectively. Thus, the k -th *importance* of v , denoted as $Imp^{(k)}(v)$, is defined as

$$Imp^{(k)}(v) = \frac{D_i^{(k)}(v)}{D_o^{(k)}(v)}. \quad (1)$$

We only cache the out-neighbors of a vertex v if its importance value $Imp^{(k)}(v)$ is sufficiently large. In Algorithm 2, lines 5–9 present the process of caching neighbors of important vertices. Let h denote the maximum depth of neighbors we consider. For each vertex v , we cache the 1 to k -hop out-neighbors of v if $Imp^{(k)}(v) \geq \tau_k$, where τ_k is a user-specified threshold. Notice that, setting h to a small number, usually 2, is enough to support a series of practical GNN algorithms. Practically, we find that τ_k is not a sensitive parameter. By experimental evaluation, setting τ_k to a small value around 0.2 can make the best trade-off between cache cost and benefit.

Interestingly, we find that the vertices to be cached is only a very small part of the whole graph. As analyzed in [51], the direct in

and out-degree of vertices in real-world graphs, i.e., $D_i^{(1)}(v)$ and $D_o^{(1)}(v)$, often obey the power-law distribution. That is, only a very few vertices in the graph have large in and out-degree. Based on this, we derive the following two theorems. The proof can be found in the appendix.

THEOREM 1. *If the in and out-degree distribution of the graph obey the power-law distribution, for any $k \geq 1$, the number of k -hop in and out-neighbors of the vertices in the graph also obey the power-law distribution.*

THEOREM 2. *If the in and out-degree distribution of the graph obey the power-law distribution, the importance value of the vertices in the graph also obey the power-law distribution.*

Theorem 2 indicates that only a very few vertices in the graph have large importance values. That means, we only need to cache a small number of important vertices to achieve a significant cost decrease of graph traversals.

Algorithm 2: Partition and Caching

Input: graph \mathcal{G} , partition number p , cache depth h , threshold $\tau_1, \tau_2, \dots, \tau_h$
Output: p subgraphs

- 1 Initialize p graph servers
- 2 **for** each edge $e = (u, v) \in \mathcal{E}$ **do**
- 3 $j = \text{ASSIGN}(u)$
- 4 Send edge e to the j -th partition
- 5 **for** each vertex $v \in V$ **do**
- 6 **for** $k \leftarrow 1$ to h **do**
- 7 Compute $D_i^{(k)}(v)$ and $D_o^{(k)}(v)$
- 8 **if** $\frac{D_i^{(k)}(v)}{D_o^{(k)}(v)} \geq \tau_k$ **then**
- 9 Cache the 1 to k -hop out-neighbors of v on each partition where v exists

3.3 Sampling

Recall that, GNN algorithms rely on aggregating neighborhood information to generate embeddings of each vertex. However, the degree distribution of real-world graphs is often skewed [51], which makes the convolution operation hard to operate. To tackle this, existing GNNs usually adopt various sampling strategies to sample a subset of neighbors with aligned sizes. Due to its importance, in our system, we abstract a sampling layer specified to optimize the sampling strategies.

Abstraction. Formally, the sampling function takes input a vertex subset V_T and extracts a small subset $V_S \subseteq V_T$ such that $|V_S| \ll |V_T|$. By taking a thorough overview of current GNN models, we abstract three kinds of different samplers, namely TRAVERSE, NEIGHBORHOOD and NEGATIVE.

- TRAVERSE: is used to sampling a batch of vertices or edges from the whole partitioned subgraphs.
- NEIGHBORHOOD: will generate the context for a vertex. The context of this vertex may be one or multi hop neighbors, which are used to encode this vertex.
- NEGATIVE: is used to generate negative samples to accelerate the convergence of the training process.

Implementation. In the literature, the sampling method plays an important role to enhance the efficiency and accuracy of the GNN algorithms [25, 24, 5, 21]. In our system, we treat all samplers as plugins. Each of them can be implemented independently. The three types of samplers can be implemented as follows.

For TRAVERSE samplers, they get data from the local subgraphs. For NEIGHBORHOOD samplers, they can get one-hop neighbors from local storage as well as multi-hop neighbors from local cache. If the neighbors of a vertex are not cached, a call to remote graph server is needed. When getting the context of a batch of vertices, we first partition the vertices into sub-batches, and the context of each sub-batch will be stitched together after being returned from the corresponding graph server. NEGATIVE samplers usually generate samples from local graph server. For some special cases, negative sampling from other graph server may be needed. Negative sampling is flexible in algorithm, and we do not need to call all graph servers in a batch. In summary, a typical sampling stage can be achieved as illustrated in Figure 5.

```

Define a TRAVERSE sampler as s1
Define a NEIGHBORHOOD sampler as s2
Define a NEGATIVE sampler as s3
...
def sampling(s1, s2, s3, batch_size):
    vertex = s1.sample(edge_type, batch_size)
    # hop_nums contains neighbor count at each hop
    context = s2.sample(edge_type, vertex, hop_nums)
    neg = s3.sample(edge_type, vertex, neg_num)
    return vertex, context, neg

```

Figure 5: Sampling stage using three kinds of samplers

We can accelerate training by adopting several efficient sampling strategies with dynamic weights. We implement the update operation in a sampler’s backward computation, just like gradient back propagation [29] of an operator. So when updating needed, what we should do is to register a gradient function for the sampler. The updating mode, synchronous or asynchronous, is due to the training algorithm.

Till now, both reading and updating will be operated on the graph storage in memory, which may lead to weak performance. According to the neighborhood requirement, the graph is partitioned by source vertices. Based on this, we split the vertices on a graph server into groups. Each group will be related with a request-flow bucket, in which the operations, including reading and updating, are all about the vertices in this group. The bucket is a lock-free queue. As shown in Figure 6, we bind each bucket to a CPU core and then each operation in the bucket will be processed sequentially without locking, which will further enhance the efficiency of the system.

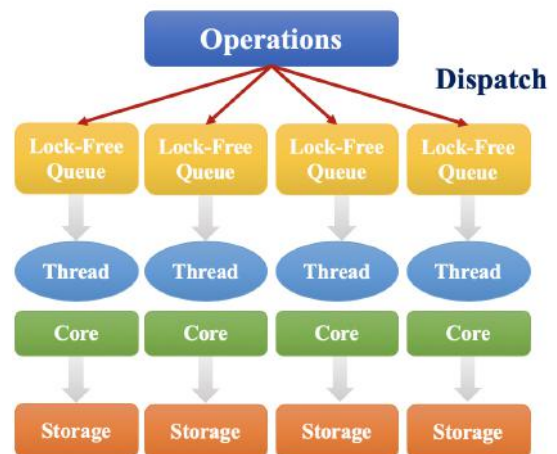


Figure 6: Lock-free graph operations

3.4 Operator

Abstraction. After sampling, the output data is aligned, and we can process it easily. Upon samplers, we need some GNN-like operators to consume them. In our system, we abstract two kinds of operators, namely AGGREGATE and COMBINE [32, 11, 27, 25]. Their roles are as follows.

- **AGGREGATE:** collects the information of each vertex’s neighbors to produce a unified result. For example, the AGGREGATE function in Algorithm 1 maps a series of vectors $\mathbf{h}_u \in \mathbb{R}^d$ to a single vector $\mathbf{h}'_v \in \mathbb{R}^d$, where u belongs to sampled neighborhood nodes of v . \mathbf{h}'_v is an intermediate result to further generate $\mathbf{h}_v^{(k)}$. The AGGREGATE function acts as the *convolution* operation since it collects the information from its surrounding neighborhoods. In different GNN methods, a variety of aggregating methods are applied, such as element-wise mean, max-pooling neural network and long short-term memory (LSTMs) [25, 24].
- **COMBINE:** takes care of how to use neighbors of a vertex to describe the vertex. In Algorithm 1, the COMBINE function maps the two vectors $\mathbf{h}_v^{(k-1)}$ and $\mathbf{h}'_v \in \mathbb{R}^d$ into a single vector $\mathbf{h}_v^{(k)} \in \mathbb{R}^d$. The COMBINE function can integrate the information of the previous hop and the neighborhoods into an unified space. Usually, in existing GNN methods, $\mathbf{h}_v^{(k-1)}$ and \mathbf{h}'_v are summed together to fed into a deep neural network.

Implementation. Notice that, both samplers and GNN-like operators not only do computations forward, but also take charge of parameters updating backward if needed. So that we can make the whole model as a network for an end-to-end training. Considering the characteristics of graph data, a lot of optimization can be taken into account to achieve better performances. Similar to SAMPLE, AGGREGATE and COMBINE are plugins of *AliGraph*, which can be implemented independently. A typical operator is made up of forward and backward computations to be easy to be involved in a deep network. Based on operators, users can set up a GNN algorithm quickly.

To further accelerate the computation of the two operators, we apply strategies by materialization of intermediate vectors $\mathbf{h}_v^{(k)}$. Notice that, as shown in [11], in each mini-batch during the training process, we can share the set of sampled neighbors for all vertices in the mini-batch. As well, we can also share the vectors $\mathbf{h}_v^{(k)}$ for all $1 \leq k \leq k_{max}$ among vertices in the same mini-batch. To this end, we store k_{max} vectors $\hat{\mathbf{h}}_v^{(1)}, \hat{\mathbf{h}}_v^{(2)}, \dots, \hat{\mathbf{h}}_v^{(k_{max})}$ to be the newest vectors of all vertices in the mini-batch. Then, in the AGGREGATE function, we apply vectors in $\hat{\mathbf{h}}_v^{(1)}, \hat{\mathbf{h}}_v^{(2)}, \dots, \hat{\mathbf{h}}_v^{(k_{max})}$ to obtain $\hat{\mathbf{h}}'_v$. After that, we apply $\hat{\mathbf{h}}'_v$ and $\mathbf{h}_v^{(k-1)}$ to compute $\hat{\mathbf{h}}_v^{(k)}$ by the COMBINE function. Finally, the stored vector $\hat{\mathbf{h}}_v^{(k)}$ is updated by $\hat{\mathbf{h}}_v^{(k)}$. By this strategy, the computation cost on the operators can be greatly reduced.

4. METHODOLOGY

On top of the system, we discuss the design of algorithms in this section. We show that existing GNNs can be easily built on *AliGraph*. Besides, we also propose a bunch of new GNN algorithms to tackle the four newly arisen challenges of embedding real-world graph data as summarized in Section 1. All of them are plugins in the *algorithm* layer of the *AliGraph* platform.

4.1 State-of-the-Art GNNs

As our *AliGraph* platform is abstracted from upon the general GNN algorithms, existing GNNs can be easily implemented on this platform. Specifically, the GNNs listed in Table 1 can all be built in

AliGraph by following the framework in Algorithm 1. Here we take the GraphSAGE as an example. Other GNNs can be implemented in a similar way. We omit them due to space limitations. Notice that, for GraphSAGE, it applies a simple node-wise sampling to extract a small subset from the neighbor set of each vertex. Obviously, its sampling strategy can be easily implemented by using our SAMPLING operator. Then, we need to instantiate the AGGREGATE and COMBINE functions in Algorithm 1. The GraphSAGE can apply the weighted element-wise mean in the AGGREGATE function in line 4. Besides, other more complex functions such as the max-pooling neural network and LSTM can also be used. In other GNN methods such as GCN, FastGCN and AS-GCN, we can replace different strategies on SAMPLING, AGGREGATE and COMBINE.

4.2 In-House Developed GNNs

Our in-house developed GNNs focus on various aspects, e.g., sampling (AHEP), multiplex (GATNE[7]), multimode (Mixture GNN[41]), hierarchy (Hierarchical GNN[35]), dynamic (Evolving GNN[61]) and multi-sourced information (Bayesian GNN).

AHEP Algorithm. This algorithm is designed to mitigate the heavy computation and storage costs of the traditional embedding propagation (EP) algorithm [16] on heterogeneous networks, HEP [59]. HEP follows the general framework of GNN with minor modifications adapted to AHG. Specifically, in HEP, the embeddings of all vertices are generated in an iterative manner. In the k -th hop, for each vertex v and each node-type c , all neighbors u of v in type c propagate its embedding $\mathbf{h}_{u,c}$ to v to reconstruct an embedding $\mathbf{h}'_{v,c}$. The embedding of v is then updated by concatenating $\mathbf{h}'_{v,c}$ across all node types. Whereas, in AHEP (HEP with adaptive sampling), we sample important neighbors instead of considering the whole set of neighbors. During this process, we design a metric to evaluate the importance of each vertex by incorporating its structural information and features. After that, all neighbors in different types are separately sampled according to their corresponding probability distributions. We carefully design the probability distributions to minimize the sampling variance. In a specific task, to optimize the AHEP algorithm, the loss function can be generally described as

$$L = L_{SL} + \alpha L_{EP} + \beta \Omega(\Theta), \quad (2)$$

where L_{SL} is the loss from supervised learning in the batch, L_{EP} is the embedding propagation loss with sampling in the batch, $\Omega(\Theta)$ is the regularizer of all trainable parameters, and α, β are two hyper-parameters. As verified by experimental results in Section 5, AHEP runs much faster than HEP while achieving comparable accuracies.

GATNE Algorithm. This algorithm is designed to cope with graphs with heterogeneous and attribute information on both vertices and edges. To address the above challenges, we propose a novel approach to capture both rich attributed information and to utilize multiplex topological structures from different node types, namely General Attributed Multiplex HeTeroogeneous Network Embedding, or abbreviated as GATNE. The overall embedding result of each vertex consists of three parts: the general embedding, the specific embedding and the attribute embedding, which correspondingly characterize the structural information, the heterogeneous information and the attribute information, respectively. For each vertex v and any node type c , the general embedding \mathbf{b}_v and the attribute embedding \mathbf{f}_v keep the same. Let t be an adjustable hyper-parameter and $\mathbf{g}_{v,t'}$ where $1 \leq t' \leq t$ be meta-specific embeddings. The specific embedding \mathbf{g}_v is obtained by concatenating all $\mathbf{g}_{v,t'}$. Then, for each type c , the overall embedding of $\mathbf{h}_{v,c}$ w.r.t. c can be written as

$$\mathbf{h}_{v,c} = \mathbf{b}_v + \alpha_c \mathbf{M}_c^T \mathbf{g}_v \mathbf{a}_c + \beta_c \mathbf{D}^T \mathbf{x}_v, \quad (3)$$

where α_c and β_c are two adjustable coefficients reflecting the importance of the specific embedding and the attribute embedding; the matrix $\mathbf{a}_c \in \mathbb{R}^m$ of coefficients are computed by using the self-attention mechanism in [39]; and \mathbf{M}_c and \mathbf{D} are two trainable transformation matrices. The final embedding result \mathbf{h}_v can then be obtained by concatenating all $\mathbf{h}_{v,c}$.

The embeddings can be learned by applying the random walk based methods similar to [43, 22]. Specifically, given a vertex v in type c in a random walk and the window size p , let $v^{-p}, v^{-p+1}, \dots, v, v^1, \dots, v^p$ denote its context. We need to minimize the negative log-likelihood as

$$-\log P_{\theta_c}(v^{-p}, \dots, v^p | v) = \sum_{1 \leq |p'| \leq p} -\log P_{\theta_c}(v_{p'} | v), \quad (4)$$

where θ_c denotes all the parameters w.r.t. type c and $P_{\theta_c}(v_{p'} | v)$ is defined by the softmax function. The objective function $-\log P_{\theta_c}(u | v)$ for each pair of vertices u and v can be easily approximated by the negative sampling method.

Mixture GNN. This model is a mixture GNN model to tackle with the heterogeneous graphs with multi-modes. In this model, we extend the skip-gram model on homogeneous graphs [43] to fit the polysemous situation on heterogeneous graphs. In the traditional skip-gram model, we try to find the embedding of graphs with parameters θ through maximizing the likelihood as

$$L_\theta = \log \Pr_\theta(Nb(v) | v) = \sum_{u \in Nb(v)} \log \Pr_\theta(u | v), \quad (5)$$

where $Nb(v)$ denotes the neighbors of v and $\Pr_\theta(u | v)$ is a softmax function. In our setting on heterogeneous graphs, each node owns multiple senses. To differentiate them, let P be the known distribution of node senses. We can rewrite the objective function as

$$L_{poly,\theta} = \log \Pr_{P,\theta}(Nb(v) | v) = \sum_{u \in Nb(v)} \log \Pr_{P,\theta}(u | v). \quad (6)$$

At this time, it is hard to incorporate the negative sampling method to directly optimize Equation (6). Alternatively, we derive a novel lower bound L_{low} of Equation (6) and try to maximize L_{low} . We find that the terms in the lower bound L_{low} can be approximated by the negative sampling. As a result, the training process can be easily implemented by slightly modifying the sampling process in existing work such as Deepwalk [43] and node2vec [22].

Hierarchical GNN. Current GNN methods are inherently flat and do not learn hierarchical representations of graphs: a limitation that is especially problematic to explicitly investigate such similarities of various types of user behaviors. This model combines the hierarchical structure to strengthen the expression power of GNN. Let $\mathbf{H}^{(k)} \in \mathbb{R}^{|V| \times d}$ denote the matrix of node embeddings computed after k steps of the GNN and \mathcal{A} be the adjacency matrix of the graph \mathcal{G} . In Algorithm 1, traditional GNN iteratively learns $\mathbf{H}^{(k)}$ by combining $\mathbf{H}^{(k-1)}$, \mathcal{A} and some trainable parameters $\theta^{(k)}$. Initially, we have $\mathbf{H}^{(0)} = \mathbf{X}$, where \mathbf{X} represent matrix of the node features.

In our hierarchical GNN, we learn the embedding result in a layer-to-layer fashion. Specifically, let $\mathcal{A}^{(l)}$ and $\mathbf{X}^{(l)}$ denote the adjacency matrix and the node feature matrix in the l -th layer, respectively. The vertex embedding result matrix $\mathbf{Z}^{(l)}$ in the l -th layer is learned by feeding $\mathcal{A}^{(l)}$ and $\mathbf{X}^{(l)}$ into the single-layer GNN method. After that, we cluster some vertices in the graph and update the adjacency matrix \mathcal{A}^l to \mathcal{A}^{l+1} . Let $\mathbf{S}^{(l)}$ denote the learned assignment matrix in the l -th layer. Each row and column in $\mathbf{S}^{(l)}$ corresponds to a cluster in the l -th and $(l+1)$ -th layer, respectively. $\mathbf{S}^{(l)}$ can be

obtained by a softmax function applied on another pooling GNN upon $\mathcal{A}^{(l)}$ and $\mathbf{X}^{(l)}$. Taking $\mathbf{Z}^{(l)}$ and $\mathbf{S}^{(l)}$ in hand, we can obtain the new coarsened adjacency matrix $\mathcal{A}^{(l+1)} = \mathbf{S}^{(l)T} \mathcal{A}^{(l)} \mathbf{S}^{(l)}$ and the new feature matrix $\mathbf{X}^{(l+1)} = \mathbf{S}^{(l)T} \mathbf{Z}^{(l)}$ for the next $(l+1)$ -th layer. As verified in Section 5, the multi-layer hierarchical GNN is more effective than the single-layer traditional GNNs.

Evolving GNN. This model is proposed to embedding vertices in the dynamic network setting. Our goal is to learn the representations of vertices in a sequential of graphs $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots, \mathcal{G}^{(T)}$. To capture the evolving nature of dynamic graphs, we divide the evolving links into two types: 1) the normal evolution representing the majority of reasonable changes of edges; and 2) burst links representing rare and abnormal evolving edges. Based on them, the embedding of all vertices in the dynamic graphs are learned in an interleave manner. Specifically, at timestamp t , the normal and burst links found on graph $\mathcal{G}^{(t)}$ are integrated with the GraphSAGE model [25] to generate embedding results \mathbf{h}_v of each vertex v in $\mathcal{G}^{(t)}$. Then, we apply a method to predict the normal and burst information on the graph $\mathcal{G}^{(t+1)}$ by using Variational Autoencoder and RNN model [31]. This process is executed in iterations to output the embedding results of each vertex v at each timestamp t .

Bayesian GNN. This model integrates two sources of information, knowledge graph embedding (e.g., symbolic) or behavior graph embedding (e.g., relations), through the Bayesian framework. To be more specific, it mimics the human understanding process in the cognitive science, in which each cognition is driven by adjusting the prior knowledge under a specific task. Specifically, given a knowledge graph \mathcal{G} and an entity (vertex) v in \mathcal{G} , its basic embedding \mathbf{h}_v is learned by purely considering \mathcal{G} itself, which characterizes the prior knowledge in \mathcal{G} . Then, a task-specific embedding \mathbf{z}_v is generated according to \mathbf{h}_v and a correction term δ_v respect to the task. That is,

$$\mathbf{z}_v \approx \mathbf{f}(\mathbf{h}_v + \delta_v), \quad (7)$$

where \mathbf{f} is a non-linear function that projects $\mathbf{h}_v + \delta_v$ as \mathbf{z}_v .

Notice that, learning exact δ_v and \mathbf{f} seems infeasible since each entity v has a different δ_v and the \mathbf{f} function is very complex. To address this problem, we apply a generation model from \mathbf{h}_v to \mathbf{z}_v by considering the second-order information. Specifically, for each entity v , we sample its correction variable δ_v from a Gaussian distribution $N(\mathbf{0}, s_v^2)$, where s_v is determined by the coefficients of \mathbf{h}_v . Then, for each pair of entities v_1 and v_2 , we sample $\mathbf{z}_{v_1} - \mathbf{z}_{v_2}$ according to another Gaussian distribution

$$N(\mathbf{f}_\phi(\mathbf{h}_{v_1} + \delta_{v_1}) - \mathbf{f}_\phi(\mathbf{h}_{v_2} + \delta_{v_2}), \text{diag}(\sigma_{v_1}^2 + \sigma_{v_2}^2)),$$

where ϕ representing the trainable parameters for the function \mathbf{f} . Let the posterior mean of δ_v be $\hat{\mu}_v$ and $\hat{\phi}$ be the resulting parameters, we finally apply $\mathbf{h}_v + \hat{\mu}_v$ as the corrected embedding for the knowledge graph, and $\mathbf{f}_{\hat{\phi}}(\mathbf{h}_v + \hat{\mu}_v)$ as the corrected task-specific embedding.

5. EXPERIMENTS

We conducted extensive experiments to evaluate our *AliGraph* platform, including both system and algorithms.

5.1 System Evaluation

In this subsection, we evaluate the performance of the underlying system in the *AliGraph* platform from the perspectives of *storage* (graph building and caching neighbors), *sampling* and *operator*. All experiments are carried on two datasets *Taobao-small* and *Taobao-large* described in Table 3, where the storage size of the latter is six times larger. Both of them represent the subgraphs of users and items extracted from the Taobao e-commerce Platform.

Table 3: Datasets used in system experiments. *Taobao-large* is six times large than *Taobao-small*.

| Dataset | # user vertices | # item vertices | # user-item edges | # item-item edges | # attributes of user | # attributes of item |
|---------------------|-----------------|-----------------|-------------------|-------------------|----------------------|----------------------|
| <i>Taobao-small</i> | 147,970,118 | 9,017,903 | 442,068,516 | 224,129,155 | 27 | 32 |
| <i>Taobao-large</i> | 483,214,916 | 9,683,310 | 6,587,662,098 | 231,085,487 | 27 | 32 |

Graph Building. The performance of graph building plays a central role in a graph computation platform. *AliGraph* supports various kinds of raw data from different file systems, partitioned or not. Figure 7 presents the time cost of graph building w.r.t. the number of workers on the two datasets. We have the following two observations: 1) the graph building time explicitly decreases w.r.t. the number of workers; 2) *AliGraph* can build large-scale graphs in minutes, e.g. 5 minutes for *Taobao-large*. This is much more efficient than most state-of-the-arts (e.g., PowerGraph [20]) that usually takes several hours).




Figure 7: Graph building time w.r.t. number of workers. The graph could be built in several minutes on two datasets.

Effects of Caching Neighbors. We examine the effects of caching k -hop neighbors of important vertices. In our caching algorithm, we need to set threshold for $Imp(v)$ as defined in Equation (1) with the analysis of $D_i^{(k)}$ and $D_o^{(k)}$. In the experiments, we locally cache the 1-hop (direct) neighbors of all vertices and vary the threshold controlling for caching the 2-hop neighbors. We gradually increase the threshold from 0.05 to 0.45 to test its sensitivity and effectiveness. Figure 8 illustrates the percentage of vertices being cached w.r.t. the threshold. We observe that the percentage of cached vertices decreases w.r.t. the threshold. When the threshold is smaller than 0.2, it decreases drastically and becomes relatively stable after that. This is because the importance of vertices obey the power-law distribution as we prove in Theorem 2. To make a good trade-off between the cache cost and the benefit, we set the threshold as 0.2 based on Figure 9 and only need to cache around 20% of extra vertices. We also compare our importance-based caching strategy w.r.t. two other strategies, namely the random strategy which caches the neighbors of a fraction of vertices selected at random and the LRU replacing strategy [12]. Figure 9 illustrates the cost time w.r.t. the percentage of cached vertices. We find that our method saves about 40%–50% time w.r.t. the random strategy and about 50%–60% time w.r.t. the LRU strategy, respectively. This is simply due to: 1) the randomly selected vertices are less likely to be accessed; and 2) the LRU strategy incurs additional cost since it frequently replaces cached vertices. Whereas, our importance-based cached vertices are more likely to be accessed by others.

Effects of Sampling. We test the effects of our optimized implementation on sampling with the batch size of 512 and cache rate 20%. Table 4 shows the time cost of the three types of sampling methods. We find that: 1) Sampling methods are very efficient which finish between a few milliseconds to no more than 60ms; 2) The sampling time grows slowly w.r.t. the graph size. Although the storage size of *Taobao-large* is six times larger compared to *Taobao-small*, the sampling time on the two datasets is quite close. These observations verify that our implementations of the sampling methods are efficient and scalable.




Figure 8: Cache rate w.r.t. threshold. Setting the threshold near 0.15 makes the best trade-off.




Figure 9: Cost time w.r.t. percentage of cached vertices. Our method saves much time than other cache strategies.

Effects of Operators. We further examine the effects of our implementation on the operators AGGREGATE and COMBINE. Table 5 shows the time cost of the two operators and the time costs can speed up by an order of magnitude with our proposed implementations. This is simply because we apply the caching strategy to eliminate the redundant computation of intermediate embedding vectors. Once again, this verifies the superiority of our *AliGraph* platform.

5.2 Algorithm Evaluations

In this subsection, we evaluate the performance of our proposed GNNs compared to state-of-the-arts. We first describe the experimental settings including the datasets, competitors and evaluation metrics. Then, we examine the efficiency and effectiveness of each proposed GNN.

5.2.1 Experimental Settings

Datasets. We employ two datasets in our experiments, including a public dataset from *Amazon* and *Taobao-small*. We choose *Taobao-small* due to the reason of the scalability of several competitors.

The statistics of the datasets are summarized in Table 6. Both of them are AHGs. The public dataset *Amazon* extracted from [42, 26] is the product metadata under the electronics category of the *Amazon* company. In this graph, each vertex represents a product with its attributes and each edge connects two products co-viewed or co-bought by the same user. It has two types of vertices, namely user and item, and four types of edges between users and items, namely click, add-to-preference, add-to-cart and buy.

Algorithms. We implement all of our proposed algorithms in this paper. For comparison, we also implement some representative graph embedding algorithms in different categories as follows:

C1: Homogeneous GE Methods. The compared methods include DeepWalk [43], LINE [50], and Node2Vec [22]. These methods can only be applied on plain graphs with purely structural information.

C2: Attributed GE Methods. The compared method includes

Table 4: Effects of optimized Sampling. All sampling methods can finish in no more than 60ms.

| Dataset | Setting | | Time (ms) | | |
|--------------|--------------|------------|-----------|--------------|----------|
| | # of workers | Cache Rate | TRAVERSE | NEIGHBORHOOD | NEGATIVE |
| Taobao-small | 25 | 18.46% | 2.59 | 45.31 | 6.22 |
| Taobao-large | 100 | 17.68% | 2.62 | 52.53 | 7.52 |

Table 5: Effects of optimized Operators with an order of magnitude of time speed up.

| Dataset | W/O Our Implementation (ms) | Our Implementation (ms) | Speedup Ratio |
|--------------|-----------------------------|-------------------------|---------------|
| Taobao-small | 7.33 | 0.57 | 12.9 |
| Taobao-large | 17.21 | 1.26 | 13.7 |

ANRL [58], which can generate embeddings capturing both structural and attributed information.

C3: Heterogeneous GE Methods. The compared methods include Methpath2Vec [14], PMNE [40], MVE [45] and MNE [57]. Methpath2Vec can only process graphs with multiple types of vertices while the other three methods can only process graphs with multiple types of edges. The PMNE involves three different kinds of approaches to extend the Node2Vec method, which are denoted as PMNE-n, PMNE-r and PMNE-c, respectively.

C4: GNN Based Methods. The comparison methods include Structural2Vec [46], GCN [32], Fast-GCN [11], AS-GCN [27], GraphSAGE [25] and HEP [59].

For fairness, all algorithms are implemented by applying the optimized operators on our system. If a method cannot process attributes and/or multiple types of vertices, we simply ignore these information in the embedding. We generate the embedding for each subgraph with the same type of edges and concatenate them together to be the final result for homogeneous based GNN. Notice that, in our examination, we do not compare each of our proposed GNN algorithms w.r.t. all competitors. This is because each algorithm is designed with different focus. We will detail the competitors of each GNN algorithm in reporting its experimental results.

Metrics. We evaluate both the efficiency and effectiveness of the proposed methods. The efficiency can be simply measured by the execution time of the algorithm. To measure effectiveness, following previous work [5, 13, 24], we apply the algorithm on the widely adopted link prediction task, which plays important roles in real-world scenarios such as recommendation. We randomly extract a portion of the data as the training data and reserve the remaining part as test data. To measure the quality of the results, four commonly used metrics are applied, namely the area under ROC curve (ROC-AUC), the PR curve (PR-AUC), the F_1 -score and the hit recall rate (HR Rate). Notably, each metric is averaged among different types of edges.

Parameters. We set d , the dimension of embedding vectors, to be 200 for all algorithms.

5.2.2 Experimental Results

AHEP Algorithm. The goal of the AHEP algorithm is to fast obtain the embedding result while does not sacrifice too much accuracy. In Table 7, we show the comparison results on result quality of AHEP w.r.t. its competitors on the Taobao-small dataset. In Figure 10, we illustrate time and space cost of different algorithms. Obviously, we have the following observations: 1) On the large Taobao-small dataset, HEP and AHEP are the only two algorithms that can produce results in reasonable time and space limits. However, AHEP is about 2-3 faster than HEP and uses much less memory than HEP. 2) In terms of the result quality, the ROC-AUC and F_1 -score of AHEP is comparable to HEP. These verify that AHEP can produce similar results of HEP by using much less time and space.

GATNE Algorithm. The goal of GATNE is designed to process graphs with heterogeneous and attributed information on both ver-

Table 6: Statistics of datasets used in experiments.

| Dataset | # of vertices | # of edges | # of vertex type | # of edge type |
|--------------|---------------|-------------|------------------|----------------|
| Amazon | 10,166 | 148,865 | 1 | 2 |
| Taobao-small | 156,988,021 | 666,197,671 | 2 | 4 |




Figure 10: Average memory cost and running time of per batch. × indicates the algorithm can not terminate in reasonable time. AHEP is 2-3 faster than HEP and uses much less memory on Taobao-small.

tices and edges. We show the comparison results of the GATNE algorithm w.r.t. its competitors in Table 8. Obviously we find that GATNE outperforms all existing methods in terms of all metrics. For example, on the Taobao-small dataset, GATNE improves the ROC-AUC, PR-AUC and F_1 -score by 4.6%, 1.84% and 5.08%, respectively. This is simply due to that GATNE simultaneously captures both the heterogenous information of vertices and edges and the attributes information. Meanwhile, we find that training time of GATNE decreases almost linearly w.r.t. the number of workers. The GATNE model converges in less than 2 hours with 150 distributed workers. The verifies the high efficiency and scalability of the GATNE method.

Mixture GNN. We compare our Mixture GNN method w.r.t. DAE [52] and β^* -VAE [36] methods. The hit recall rate of applying the embedding results into the recommendation task is shown in Table 9. Notice that, by applying our model, the hit recall rates have been improved by around 2%. Similarly, this improvement also makes significant contributions in a large network.

Hierarchical GNN. We compare our Hierarchical GNN method w.r.t. GraphSAGE. The results is shown in Table 10. The F_1 -score is significantly improved by around 7.5%. This indicates that our Hierarchical GNN can generate more promising embedding results.

Evolving GNN. We compare our Evolving GNN method w.r.t. other methods on the multi-class link prediction task. The competitors include the representative algorithms DeepWalk, DANE, DNE, TNE and GraphSAGE. These competitor algorithms can not handle dynamic graphs, thus we run the algorithm on each snapshot of the dynamic graphs and report the average performance over all timestamps. The comparison results on the Taobao-small dataset are shown in Table 11. We easily find that, Evolving GNN outperforms all other methods in terms of all metrics. For example, with burst change, Evolving GNN improves the micro and macro F_1 -score by 4.2% and 3.6%. This is simply because our proposed method can better capture the dynamic changes of real-world networks, thus can produce more promising results.

Bayesian GNN. The goal of this model is to combine Bayesian method with the traditional GNN model. We use GraphSAGE as the baseline and compare the results with and without incorporating the proposed Bayesian model. We present the hit recall rate of the recommendation result in Table 12. Notice that, we consider the granularity of both item brands and categories. Obviously, when applying our Bayesian model, the hit recall rates have been increased by 1% to 3% respectively. Notice that, this improvement can bring significant benefits on our network containing 9 million items.

Table 7: Effectiveness comparison of AHEP w.r.t. its competitors. AHEP is close to HEP on Taobao-small.

| Method | ROC-AUC (%) | F_1 -score (%) |
|----------------|-------------|------------------|
| Structural2Vec | N.A. | N.A. |
| GCN | N.A. | N.A. |
| FastGCN | N.A. | N.A. |
| GraphSAGE | N.A. | N.A. |
| AS-GCN | O.O.M | O.O.M |
| HEP | 77.77 | 57.93 |
| AHEP | 75.51 | 50.97 |

“N.A.” indicates the algorithm can not terminate in reasonable time.

“O.O.M.” indicates that the algorithm terminates due to out of memory.

6. RELATED WORK

In this section, we briefly review the state-of-the-arts on GE and GNN methods. Based on the four challenges summarized in Section 1, we categorize existing methods as follows.

Homogeneous. DeepWalk [43] first generates a corpus on graphs by random walk and then trains a skip-gram model on the corpus. LINE [50] learns node presentations by preserving both first-order and second-order proximities. NetMF [44] is a unified matrix factorization framework for theoretically understanding and improving DeepWalk and LINE. Node2Vec [22] adds two parameters to control the random walk process while SDNE [53] proposes a structure-preserving embedding method. GCN [32] incorporates neighbors’ feature representations using convolutional operations. GraphSAGE [25] provides an inductive approach to combine structural information with node features.

Heterogeneous. For graph with multiple types of vertices and/or edges, PMNE [40] proposes three methods to project a multiplex network into a continuous vector space. MVE [45] embeds networks with multiple views in a single collaborated embedding using the attention mechanism. MNE [57] uses one common embedding and several additional embeddings of each edge-type for each node, which are jointly learned by a unified network embedding model. Mvn2Vec [55] explores the embedding results by simultaneously modeling preservation and collaboration. HNE [6] jointly considers the contents and topological structures to be unified vector representations. PTE [49] constructs large-scale heterogeneous text network from labeled information, which is then embedded into a low-dimensional space. Metapath2Vec [14] and HERec [47] formalize meta-path based random-walks to construct the heterogeneous neighborhood of a node and then leverage skip-gram models to perform node embeddings.

Attributed. Attributed network embedding aims to seek for low-dimensional vector representations to preserve both topological and attribute information. TADW [56] incorporates text features of vertices into network representation learning by matrix factorization. LANE [28] smoothly incorporates label information into the attributed network embedding while preserving their correlations. AANE [27] enables joint learning process to be done in a distributed manner for accelerated attributed network embedding. SNE [37] proposes a generic framework for embedding social networks by capturing both the structural proximity and attribute proximity. DANE [19] can capture the high nonlinearity and preserve various proximities in both topological structure and node attributes. ANRL [58] uses a neighbor enhancement autoencoder to model the node attribute information and a skip-gram model to capture the network structure.

Dynamic. Actually, some static methods [43, 50] can also handle dynamic network by updating the new vertices based on static embedding. Considering the new vertices’ influence on the original networks, [15] extends the skip-gram methods to update the original

vertices’ embedding. [60] focuses on capturing the triadic structure properties for learning network embedding. Considering both the network structure and node attributes, [34] focuses on updating the top eigenvectors and eigenvalues for the streaming network.

7. CONCLUSIONS AND FUTURE WORK

We summarize four challenges from the current practical graph data problems, namely *large-scale*, *heterogeneous*, *attributed* and *dynamic*. Based on these challenges, we design and implement a platform, *AliGraph*, which provides both system and algorithms to tackle more practical problems. In the future, we will focus on but not limited to the following directions: 1) GNN for edge-level and subgraph-level embeddings; 2) More execution optimizations, such as co-location of computation variables in GNN with graph data to reduce the cross network traffic, introduction of new gradient optimization to leverage the trait of GNN to speed up the distributed training without accuracy loss, and better assignment of the workers in multi-GPU architectures; 3) Early-stop mechanism, which can help to terminate training tasks earlier when no promising results can generate.

8. REFERENCES

- [1] P. Battaglia, R. Pascanu, M. Lai, and D. J. Rezende. Interaction networks for learning about objects, relations and physics. In *NIPS*, pages 4502–4510, 2016.
- [2] S. Bhagat, G. Cormode, and S. Muthukrishnan. Node classification in social networks. *Computer Science*, 16(3):115–148, 2011.
- [3] E. G. Boman, K. D. Devine, and S. Rajamanickam. Scalable matrix computations on large scale-free graphs using 2d graph partitioning. 2013.
- [4] U. Brandes, M. Gaertler, and D. Wagner. Experiments on graph clustering algorithms. *LNCS*, 2832:568–579, 2003.
- [5] H. Cai, V. W. Zheng, C. C. Chang, H. Cai, V. W. Zheng, and C. C. Chang. A comprehensive survey of graph embedding: Problems, techniques and applications. *TKDE*, 30(9):1616–1637, 2017.
- [6] S. Chang, W. Han, J. Tang, G.-J. Qi, C. C. Aggarwal, and T. S. Huang. Heterogeneous network embedding via deep architectures. In *KDD*, pages 119–128, 2015.
- [7] Cen, Y., Zou, X., Zhang, J., Yang, H., Zhou, J., Tang, J. Representation Learning for Attributed Multiplex Heterogeneous Network. In *KDD*, 2019.
- [8] Liu, N., Tan, Q., Li, Y., Yang, H., Zhou, J., Hu, X. Is a Single Vector Enough? Exploring Node Polysemy for Network Embedding. In *KDD*, 2019.
- [9] Li, C., Shen, D., Jia, K., Yang, H. Hierarchical Representation Learning for Bipartite Graphs. In *IJCAI*, 2019.
- [10] Zhao, Y., Wang, X., Yang, H., Song, L. and Tang, J. Large Scale Evolving Graphs with Burst Detection. In *IJCAI*, 2019.
- [11] J. Chen, T. Ma, and C. Xiao. Fastgcn: fast learning with graph convolutional networks via importance sampling. *arXiv:1801.10247*, 2018.
- [12] M. Chrobak and J. Noga. Lru is better than fifo. In *Acm-siam Symposium on Discrete Algorithms*, 1998.
- [13] P. Cui, X. Wang, J. Pei, and W. Zhu. A survey on network embedding. *TKDE*, 2018.
- [14] Y. Dong, N. V. Chawla, and A. Swami. metapath2vec: Scalable representation learning for heterogeneous networks. In *KDD*, pages 135–144, 2017.
- [15] L. Du, Y. Wang, G. Song, Z. Lu, and J. Wang. Dynamic network embedding: An extended approach for skip-gram based network embedding. In *IJCAI*, pages 2086–2092, 2018.

Table 8: Effectiveness comparison of GATNE w.r.t. its competitors. GATNE outperforms all competitors in terms of all metrics on both Amazon and Taobao-small. GATNE lifts the F_1 -score by 16.43% on the Amazon dataset.

| Method | Amazon | | | Taobao-small | | |
|--------------|--------------|--------------|------------------|--------------|--------------|------------------|
| | ROC-AUC (%) | PR-AUC (%) | F_1 -score (%) | ROC-AUC (%) | PR-AUC (%) | F_1 -score (%) |
| DeepWalk | 94.20 | 94.03 | 87.38 | 65.58 | 78.13 | 70.14 |
| Node2Vec | 94.47 | 94.30 | 87.88 | N.A. | N.A. | N.A. |
| LINE | 81.45 | 74.97 | 76.35 | N.A. | N.A. | N.A. |
| ANRL | 95.41 | 94.19 | 89.60 | N.A. | N.A. | N.A. |
| Metapath2Vec | 94.15 | 94.01 | 87.48 | N.A. | N.A. | N.A. |
| PMNE-n | 95.59 | 95.48 | 89.37 | N.A. | N.A. | N.A. |
| PMNE-r | 88.38 | 88.56 | 79.67 | N.A. | N.A. | N.A. |
| PMNE-c | 93.55 | 93.46 | 86.42 | N.A. | N.A. | N.A. |
| MVE | 92.98 | 93.05 | 87.80 | 66.32 | 80.12 | 72.14 |
| MNE | 91.62 | 92.46 | 84.44 | 79.60 | 93.01 | 84.86 |
| GATNE | 96.25 | 94.77 | 91.36 | 84.20 | 95.04 | 89.94 |

Table 9: Effectiveness comparison of Mixture GNN w.r.t. its competitors. Mixture GNN improves the hit recall rate by around 2% on Taobao-small.

| Method | HR Rate@20 | HR Rate@50 |
|--------------------|----------------|----------------|
| DAE | 0.12622 | 0.21619 |
| β^* -VAE | 0.11767 | 0.19997 |
| Mixture GNN | 0.14317 | 0.23680 |

Table 10: Effectiveness comparison of Hierarchical GNN w.r.t. its competitors. Hierarchical GNN improves the hit recall rate by 7.5% on Taobao-small.

| Method | ROC-AUC(%) | PR-AUC(%) | F_1 -score(%) |
|-------------------------|--------------|--------------|-----------------|
| GraphSAGE | 82.89 | 44.45 | 45.76 |
| Hierarchical GNN | 87.34 | 54.87 | 53.20 |

- [16] A. G. Duran and M. Niepert. Learning graph representations with embedding propagation. In *Advances in Neural Information Processing Systems*, pages 5119–5130, 2017.
- [17] W. Fan, J. Xu, Y. Wu, W. Yu, J. Jiang, Z. Zheng, B. Zhang, Y. Cao, and C. Tian. Parallelizing sequential graph computations. In *SIGMOD*, pages 495–510, 2017.
- [18] A. Fout, J. Byrd, B. Shariat, and A. Ben-Hur. Protein interface prediction using graph convolutional networks. In *NIPS*, pages 6530–6539, 2017.
- [19] H. Gao and H. Huang. Deep attributed network embedding. In *IJCAI*, pages 3364–3370, 2018.
- [20] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. Powergraph: Distributed graph-parallel computation on natural graphs. In *OSDI*, 2012.
- [21] P. Goyal and E. Ferrara. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 2018.
- [22] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, pages 855–864, 2016.
- [23] T. Hamaguchi, H. Oiwa, M. Shimbo, and Y. Matsumoto. Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach. In *IJCAI*, 2017.
- [24] W. L. Hamilton, R. Ying, and J. Leskovec. Representation learning on graphs: Methods and applications. 2017.
- [25] W. L. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *NIPS*, pages 1025–1035, 2017.
- [26] R. He and J. McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*, pages 507–517.
- [27] X. Huang, J. Li, and X. Hu. Accelerated attributed network embedding. In *SDM*, pages 633–641. SIAM, 2017.
- [28] X. Huang, J. Li, and X. Hu. Label informed attributed network embedding. In *WSDM*, pages 731–739, 2017.
- [29] D. R. Hush and J. M. Salas. Improving the learning rate of back-propagation with the gradient reuse algorithm. In *IEEE International Conference on Neural Networks*, 1988.
- [30] G. Karypis and V. Kumar. Metis—unstructured graph partitioning and sparse matrix ordering system. *Technical Report*.
- [31] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013.
- [32] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [33] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. In *Nature*, pages 521–436. 2015.
- [34] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu. Attributed network embedding for learning in a dynamic environment. In *CIKM*, pages 387–396. ACM, 2017.
- [35] Li, C., Shen, D., Jia, K., Yang, H. Hierarchical Representation Learning for Bipartite Graphs. In *IJCAI*, 2019.
- [36] D. Liang, R. G. Krishnan, M. D. Hoffman, and T. Jebara. Variational autoencoders for collaborative filtering. 2018.
- [37] L. Liao, X. He, H. Zhang, and T.-S. Chua. Attributed social network embedding. *TKDE*, 30(12):2257–2270, 2018.
- [38] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. 2003.
- [39] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio. A structured self-attentive sentence embedding. *arXiv:1703.03130*, 2017.
- [40] W. Liu, P.-Y. Chen, S. Yeung, T. Suzumura, and L. Chen. Principled multilayer network embedding. In *ICDM*, pages 134–141. IEEE, 2017.
- [41] Liu, N., Tan, Q., Li, Y., Yang, H., Zhou, J., Hu, X. Is a Single Vector Enough? Exploring Node Polysemy for Network Embedding. In *KDD*, 2019.
- [42] J. McAuley, C. Targett, Q. Shi, and A. Van Den Hengel. Image-based recommendations on styles and substitutes. In *SIGIR*, pages 43–52. ACM, 2015.
- [43] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *KDD*, pages 701–710. ACM, 2014.
- [44] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *WSDM*, pages 459–467, 2018.
- [45] M. Qu, J. Tang, J. Shang, X. Ren, M. Zhang, and J. Han. An attention-based collaboration framework for multi-view network representation learning. In *CIKM*, pages 1767–1776. ACM, 2017.
- [46] L. F. R. Ribeiro, P. H. P. Saverese, and D. R. Figueiredo. struc2vec: Learning node representations from structural identity. 2017.

Table 11: Effectiveness comparison of Evolving GNN w.r.t. its competitors. Evolving GNN improves the F_1 -score by about 4% on Taobao-small.

| Method | Normal Evolution | | burst Change | |
|---------------------|------------------------|------------------------|------------------------|------------------------|
| | Micro F_1 -score (%) | Macro F_1 -score (%) | Micro F_1 -score (%) | Macro F_1 -score (%) |
| DeepWalk | N.A. | N.A. | N.A. | N.A. |
| DANE | N.A. | N.A. | N.A. | N.A. |
| TNE | 79.9 | 71.9 | 69.1 | 67.2 |
| GraphSAGE | 71.4 | 70.4 | 60.7 | 60.5 |
| Evolving GNN | 81.4 | 77.7 | 73.3 | 70.8 |

Table 12: Effectiveness comparison of Bayesian GNN w.r.t. its competitors. Bayesian GNN improves the hit recall rate by 1%–3% on Taobao-small.

| Granularity | HR Rate@ | Click | | Buy | |
|-------------|----------|--------------|----------------------|-----------|----------------------|
| | | GraphSAGE | GraphSAGE + Bayesian | GraphSAGE | GraphSAGE + Bayesian |
| Brand | 10 | 15.97 | 16.14 | 24.87 | 25.10 |
| | 30 | 16.65 | 17.12 | 25.70 | 26.57 |
| | 50 | 17.26 | 17.90 | 26.39 | 27.33 |
| Category | 10 | 27.46 | 27.49 | 27.85 | 27.91 |
| | 30 | 28.43 | 29.99 | 28.50 | 29.45 |
| | 50 | 29.58 | 32.88 | 26.26 | 31.47 |

- [47] C. Shi, B. Hu, X. Zhao, and P. Yu. Heterogeneous information network embedding for recommendation. *TKDE*, 2018.
- [48] I. Stanton and G. Kliot. Streaming graph partitioning for large distributed graphs. In *KDD*, 2013.
- [49] J. Tang, M. Qu, and Q. Mei. Pte: Predictive text embedding through large-scale heterogeneous text networks. In *KDD*, pages 1165–1174. ACM, 2015.
- [50] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. Line: Large-scale information network embedding. In *WWW*, pages 1067–1077, 2015.
- [51] S. Tanimoto. Power laws of the in-degree and out-degree distributions of complex networks. *Physics*, 2009.
- [52] P. Vincent, H. Larochelle, Y. Bengio, and P. A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008.
- [53] D. Wang, P. Cui, and W. Zhu. Structural deep network embedding. In *KDD*, pages 1225–1234, 2016.
- [54] Z. Wang, Y. Tan, and Z. Ming. Graph-based recommendation on social networks. In *APWeb*, 2010.
- [55] W. Xiong, M. Yu, S. Chang, X. Guo, and W. Y. Wang. One-shot relational learning for knowledge graphs. In *EMNLP*, pages 1980–1990, 2018.
- [56] C. Yang, Z. Liu, D. Zhao, M. Sun, and E. Y. Chang. Network representation learning with rich text information. In *IJCAI*, 2015.
- [57] H. Zhang, L. Qiu, L. Yi, and Y. Song. Scalable multiplex network embedding. In *IJCAI*, pages 3082–3088, 2018.
- [58] Z. Zhang, H. Yang, J. Bu, S. Zhou, P. Yu, J. Zhang, M. Ester, and C. Wang. Anrl: Attributed network representation learning via deep neural networks. In *IJCAI*, pages 3155–3161, 2018.
- [59] V. W. Zheng, M. Sha, Y. Li, H. Yang, Z. Zhang, and K.-L. Tan. Heterogeneous embedding propagation for large-scale e-commerce user alignment. In *ICDM*, 2018.
- [60] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang. Dynamic network embedding by modeling triadic closure process. In *AAAI*, 2018.
- [61] Zhao, Y., Wang, X., Yang, H., Song, L. and Tang, J. Large Scale Evolving Graphs with Burst Detection. In *IJCAI*, 2019.

Appendix

Proof of Theorem 1 Let $D_i^{(k)}$ and $D_o^{(k)}$ be two random variables representing the number of k -hop in and out-neighbors of a randomly chosen vertex from the graph, respectively. We derive the

probability distribution of $D_i^{(k)}$ and $D_o^{(k)}$ for each $k \geq 1$ by induction.

1. Following previous work [51], when $k = 1$, the in-degree $D_i^{(1)}$ and out-degree $D_o^{(1)}$ both obey the power-law distribution. Specifically, let $\gamma_o^{(1)}$ and $\gamma_i^{(1)}$ denote the exponent, we have $\Pr(D_i^{(1)} = q) \propto q^{-\gamma_i^{(1)}}$ and $\Pr(D_o^{(1)} = q) \propto q^{-\gamma_o^{(1)}}$.

2. Then, we consider the probability distribution of $D_i^{(k)}$ and $D_o^{(k)}$ where $k \geq 2$. Assume that $D_i^{(k-1)}$ and $D_o^{(k-1)}$ obey the power-law distribution with exponent $\gamma_i^{(k-1)}$ and $\gamma_o^{(k-1)}$, respectively. Let u be a randomly chosen vertex from the graph. If we randomly chose a $(k-1)$ -hop in-neighbor s of u and a one-hop out-neighbor v of u , v is obviously a k -hop out-neighbor of s . Since s is chosen randomly, we have

$$\begin{aligned} \Pr(D_o^{(k)} = q) &= \sum_{j=1}^n \Pr(D_o^{(1)} = j) \cdot \Pr(D_i^{(k)} = \frac{q}{j}) \\ &\propto \sum_{j=1}^n j^{-\gamma_o^{(1)}} \left(\frac{q}{j}\right)^{-\gamma_i^{(k-1)}} \propto q^{-\gamma_i^{(k-1)}} \sum_{j=1}^n j^{\gamma_i^{(k-1)} - \gamma_o^{(1)}}. \end{aligned}$$

Since n , $\gamma_i^{(k-1)}$ and $\gamma_o^{(1)}$ are all fixed, the last term is a constant value. As a result, we have $\Pr(D_o^{(k)} = q) \propto q^{-\gamma_o^{(k)}}$, where $\gamma_o^{(k)}$ is a term determined by $\gamma_i^{(k-1)}$ and the last term. In similar, we also have $\Pr(D_i^{(k)} = q) \propto q^{-\gamma_i^{(k)}}$. This indicates both $D_o^{(k)}$ and $D_i^{(k)}$ obey the power-law distribution.

By summarizing 1 and 2, we find that both $D_i^{(k)}$ and $D_o^{(k)}$ obey the power-law distribution for each $k \geq 1$.

Proof of Theorem 2 Let $Imp^{(k)} = D_i^{(k)}/D_o^{(k)}$ be a random variable denoting the importance of a randomly chosen vertex. We have

$$\begin{aligned} \Pr(Imp^{(k)} = \ell) &= \sum_{j=1}^n \Pr(D_o^{(k)} = j) \Pr(D_i^{(k)} = \ell j) \\ &= \sum_{j=1}^n j^{-\gamma_o^{(k)}} (\ell j)^{-\gamma_i^{(k)}} = \ell^{-\gamma_i^{(k)}} \sum_{j=1}^n j^{-(\gamma_i^{(k)} + \gamma_o^{(k)})}. \end{aligned}$$

Since the last term is also a constant value, we find that $Imp^{(k)}$ also obeys the power-law distribution. This analysis result indicates that the importance $Imp^{(k)}$ of most vertices is very small. Thus, we only need to cache a small number of vertices with large importance values. Intuitively, for any vertex whose importance is large, it has a large number of in-neighbors and a small number of out-neighbors.