

Bridging Network Embedding and Graph Summarization

Di Jin
University of Michigan
dijin@umich.edu

Eunye Koh
Adobe Research
eunye@adobe.com

Ryan A. Rossi
Adobe Research
rossi@adobe.com

Sungchul Kim
Adobe Research
sukim@adobe.com

Danai Koutra
University of Michigan
dkoutra@umich.edu

Anup Rao
Adobe Research
anuprao@adobe.com

ABSTRACT

An important reason behind the prevalence of node representation learning is their superiority in downstream machine learning tasks on graphs. However, storing the vector-based node representation of massive real-world graphs often requires space that is orders of magnitude larger. To alleviate this issue, we introduce the problem of *latent network summarization* that is complementary to the problem of network embedding, and propose a general framework called MULTI-LENS. Instead of deriving dense node-wise representations, the goal of latent network summarization is to summarize the structural properties of the graph in a latent space with dimensionality that is independent of the nodes or edges in the input graph. The size-independent graph summary given by MULTI-LENS consists of (i) a set of *relational aggregators* with their compositions (*relational functions*) that captures heterogeneous features of multi-order node-centric subgraphs, and (ii) the low-rank approximations to matrices that incorporate captured structural features. In addition, MULTI-LENS is able to derive node embeddings on the fly from this latent summary due to its inductive properties. MULTI-LENS bridges advantages brought by both embeddings and graph summarization, and applies to graphs with or without directionality, weights, attributes or labels. Extensive experiments on both synthetic and real-world graphs show that MULTI-LENS achieves 2 – 89% improvement in AUC for link prediction, while requiring *less than 79×* space compared to existing representation learning approaches. We also show the effectiveness of MULTI-LENS summaries in anomaly and event detection on two real-world graphs.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; **Logical and relational learning**; • **Mathematics of computing** → **Graph algorithms**; • **Information systems** → **Data mining**;

KEYWORDS

graph mining, graph summarization, graph compression, heterogeneous networks, structural similarity

1 INTRODUCTION

Recent advances in representation learning for graphs have led to a variety of connectivity- and feature-based embeddings that achieve superior performance in specific downstream tasks, such as link prediction, node classification, and alignment [12, 15, 34]. At the same time though, the learned, dense K -dimensional node embeddings

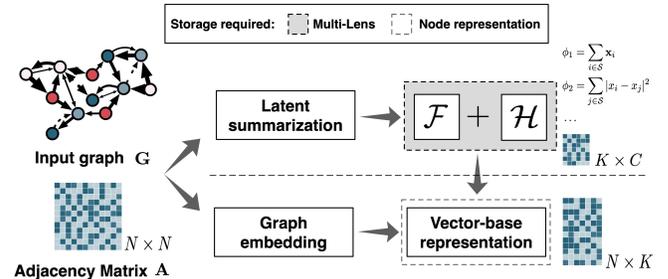


Figure 1: Our proposed approach to Latent Network Summarization called MULTI-LENS produces a summary \mathcal{J} consisting of relational functions \mathcal{F} and node-independent matrices \mathcal{H} . Notice that embedding methods output a $N \times K$ matrix of N node embeddings (and thus are graph size-dependent) whereas latent summarization methods output a $K \times C$ matrix where the dimensions are independent of the size of the graph. Nevertheless, MULTI-LENS can also derive node embeddings *on the fly* without having to store them.

(with real values) pose computational and storage challenges especially for massive graphs. By following the conventional setting of $K = 128$ for the dimensionality, a graph of 1 billion nodes requires roughly 1TB for its embeddings. Moreover, this dense representation often requires significantly more space than the original, sparse adjacency matrix of a graph. For example, for the datasets that we consider in our empirical analysis, the learned embeddings from existing representation learning techniques require 3 – 48× more space than the original edge files. To address these shortcomings, we introduce the problem of *latent network summarization*:

PROBLEM 1 (LATENT NETWORK SUMMARIZATION—INFORMAL). *Given an arbitrary network, the goal of latent network summarization is to find a low-dimensional representation in a latent space such that it is independent of the graph size, i.e., the number of nodes and edges. Among other tasks, the representation should support on-the-fly computation of specific node embeddings, which capture the structural properties of the nodes.*

Latent network summarization and network embedding are *complementary* learning tasks with fundamentally different goals and outputs, as shown in Fig. 1. In particular, the goal of network embedding is to derive N node embedding vectors of K dimensions each that capture the structural properties of the nodes. Thus, the output is a $N \times K$ matrix that is *dependent* on the size of the graph (number of nodes) [11, 29]. This is in contrast to the latent

network summarization problem, where the goal is to learn a *size-independent representation* of the graph. On the other hand, latent network summarization also differs from traditional summarization approaches that typically derive super-graphs. Instead of deriving a super-graph as done by most summarization methods [24], the goal of latent network summarization is to learn a size-independent latent representation of the graph.

To efficiently solve the proposed latent network summarization problem, we propose MULTI-LENS, a multi-level inductive approach based on graph function compositions. In a nutshell, the method begins with a set of arbitrary graph functions (e.g., degree) and iteratively uses relational aggregators over the neighborhoods to derive deeper function compositions that capture graph features at multiple levels (or distances). Low-rank approximation is then used to derive the latent network summarization at each level. We focus on heterogeneous graphs, a general class of graphs that models complex structural and semantic information for both entities and relationships, and includes the prevalent case of homogeneous graphs as a special case. In addition, MULTI-LENS has many important and useful properties including **(P1)** generality to handle arbitrary heterogeneous graphs, **(P2)** space-efficiency, **(P3)** naturally supporting inductive learning by deriving relational function compositions and **(P4)** being able to derive node embeddings on the fly to answer queries related to specific nodes in the network.

Our main contributions are summarized as follows:

- **Novel Problem Formulation.** We introduce and formulate the problem of *latent network summarization*, which is complementary yet fundamentally different from network embeddings.
- **Principled Approach.** We propose MULTI-LENS, which expresses a class of methods for latent network summarization. MULTI-LENS naturally supports inductive learning by leveraging relational function compositions; on-the-fly embedding computation for all or a subset of nodes; and heterogeneous input networks.
- **Time- and Space-efficiency.** MULTI-LENS is *scalable* with time complexity linear on the number of nodes and edges, and *space-efficient* with size independent of the graph size (i.e., the number of nodes and edges). Besides, MULTI-LENS is easily *parallelizable* as the computations for each node are independent of one another.
- **Empirical analysis on real datasets.** We show the superiority of our proposed method over state of the art baselines on a variety of real-world heterogeneous (bi-/multi-partite, weighted/unweighted) graphs via space and runtime analysis, and tasks such as event detection and link prediction. MULTI-LENS achieves at least 79x improvement in space efficiency and from about 2% to 89% improvement in accuracy.

2 LATENT NETWORK SUMMARIZATION

In this section, we introduce the problem of *latent network summarization*. Table 1 lists the main symbols and notations used in this work. The problem of latent network summarization aims to learn a compressed representation that captures the main structural information of the network and depends only on the complexity of the network instead of its size. More formally:

DEFINITION 1 (LATENT NETWORK SUMMARIZATION). *Given an arbitrary graph $G = (V, E)$ with V as the node set and E as the edge set ($|V| = N, |E| = M$), the goal of latent network summarization is*

Table 1: Summary of symbols and notations

Symbol	Definition
$G(V, E), \mathbf{A}$	(un)directed and (un)weighted heterogeneous network with node-set V and edge-set E ($ V = N, E = M$) and its adjacency matrix
$\mathcal{F} = \{f_i\}, F$	ordered set of relational functions and its size
$f_b \in \mathcal{F}_b$	the set of base graph functions (special relational functions)
l, L	index for level & total number of levels (i.e., max order of a rel. fns)
$\mathcal{B} = \{\mathbf{b}_i\}, B$	set of initial feature vectors in length N and its size
$\Phi = \{\phi_i\}, R$	set of relational aggregators and its size
S	set of nodes of interest
$\mathcal{T}_V, \mathcal{T}_V $	set of object types in the heterogeneous graph and its size
$\mathcal{T}_E, \mathcal{T}_E $	set of edge types in the heterogeneous graph and its size
$\Gamma_t(i), \Gamma(i)$	typed t /type-independent 1-neighborhood of node i , resp.
$\Gamma_t^+(i), \Gamma_t^-(i)$	out-/in-typed t neighborhood of node v
$q^{(l)}$	number of columns of matrix Y at level l
$\mathbf{X}^{(0)}$	$N \times p^{(0)}$ base feature matrix derived by \mathcal{F}_b
$\mathbf{X}^{(l)}$	$N \times p^{(l)}$ generated feature matrix for level l
\mathbf{x}, \mathbf{y}	N -dimensional feature vector (column of \mathbf{X} and \mathbf{Y} , respectively)
K_i, K	embedding dimension at level- i and the final output emb. dimension
$\mathbf{Y}^{(l)}$	$N \times q^{(l)}$ matrix induced by applying Ψ on feature matrix $\mathbf{X}^{(l)}$

to learn a function $G \rightarrow \mathbb{R}^{K \times C}$ that maps the graph G to a $K \times C$ matrix that captures the main structural information of G such that $K, C \ll N$ and $K, C \ll M$. Hence, the size of the output matrix is independent of the graph size. The output representations can be used directly in data mining tasks (e.g., anomaly detection), or can be used to derive all or a subset of node-specific embeddings on the fly for learning tasks such as link prediction or classification.

This problem differs from the network embedding problem that aims to derive an $N \times K$ embedding matrix consisting of N embedding vectors of K dimensions each, and thus the output is *dependent* on the size of the graph (number of nodes) [11, 29]. This is in contrast to the latent network summarization problem where the goal is to learn a size-independent representation of the graph. Hence, the network embedding problem is *graph size-dependent* whereas latent network summarization is *graph size-independent*.

3 MULTI-LENS APPROACH

To efficiently address the problem of latent network summarization introduced in Section 2, we propose MULTI-LENS, which expresses a class of latent network summarization methods that satisfies all desired properties **(P1-P4)**. MULTI-LENS leverages generally-defined relational operations to handle arbitrary types of networks (Sec. 3.1). The main idea is to explore node proximity in node-centric subgraphs at different distances (Sec. 3.2), and summarize resultant structural feature matrices through low-dimension approximation (Sec. 3.3-3.4). The overview of MULTI-LENS is shown in Fig. 2.

3.1 Handling Arbitrary Networks

Recall that our proposed problem definition (Sec. 2) applies to any arbitrary graph **(P1)**. As a general class, we refer to heterogeneous (information) networks or typed networks.

DEFINITION 2 (HETEROGENEOUS NETWORK). *A heterogeneous network is defined as $G = (V, E, \theta, \xi)$ with node-set V , edge-set E , a node type mapping function $\theta : V \rightarrow \mathcal{T}_V$, and an edge type mapping function defined as $\xi : E \rightarrow \mathcal{T}_E$. In the mapping functions, \mathcal{T}_V and \mathcal{T}_E denote the set of node object types and edge types, respectively.*

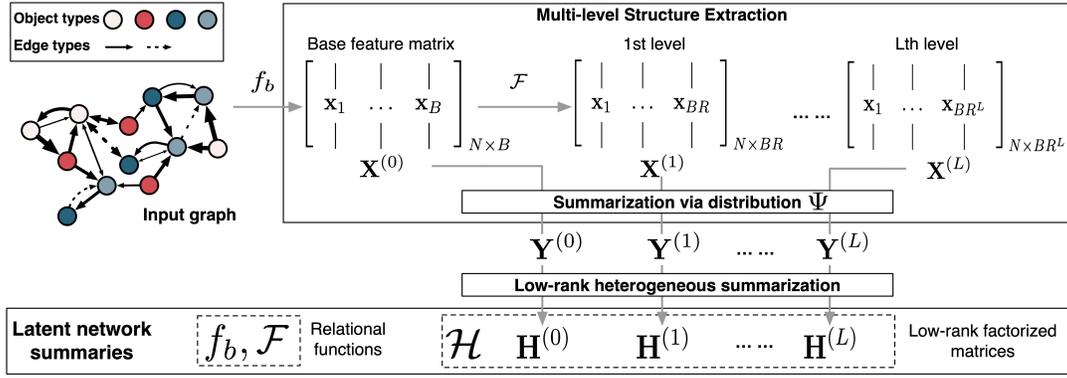


Figure 2: Overview of MULTI-LENS. It first applies base graph function f_b to derive the base feature matrix $X^{(0)}$ on which relational functions \mathcal{F} are iteratively applied to derive the feature matrices $X^{(l)}$ that capture proximity at different distances. MULTI-LENS applies $\Psi(X^{(l)})$ to represent column feature vectors $X^{(l)}$ through distributions denoted as $Y^{(l)}$, and then perform low-rank approximation to derive the corresponding factorized matrix $H^{(l)}$. The size of the resultant latent network summaries, $\mathcal{J} = \{\mathcal{F}, \mathcal{H}\}$, is independent of N, M .

We assume that the network is directed and weighted, since unweighted and undirected graphs are special cases. Within heterogeneous networks, the so-called typed 1-neighborhood or egonet¹ Γ_t of a node captures its local structural information:

DEFINITION 3 (TYPED 1-NEIGHBORHOOD Γ_t). Given an arbitrary node i in graph $G(V, E, \theta, \xi)$, the typed t -neighborhood $\Gamma_t(i)$ is the set of nodes with type t that are reachable by following directed edges $e \in E$ originating from i with 1-hop distance.

The 1-neighborhood of node i , $\Gamma(i)$, is a superset of the typed neighborhood $\Gamma_t(i)$, and includes nodes in the 1-neighborhood of i regardless of their types. Higher order neighborhoods are defined similarly, but more computationally expensive to explore. e.g., ℓ -neighborhood, $\Gamma^\ell(i)$ denotes the set of nodes reachable following directed edges $e \in E$ originating from i within ℓ -hop distance.

The goal of latent network summarization is to find a size-independent representation that captures the structure of the network and its underlying nodes. Capturing the structure depends on the semantics of the network (e.g., weighted, directed), and different ways have to be employed for different input networks types. To be able to generalize to arbitrary networks, we leverage relational aggregators and functions.

DEFINITION 4 (RELATIONAL AGGREGATOR). A relational aggregator $\phi(\mathbf{x}, \mathcal{S})$ is defined as a basic function (e.g., sum) that operates on the $N \times 1$ feature vector \mathbf{x} associated with a set of related graph elements \mathcal{S} (e.g., 1-neighborhood Γ) and returns a single value.

The computation applied over the graph or a subgraph (e.g., 1-hop neighborhood) generalizes for inductive/across-network transfer learning tasks [34]. Multiple relational aggregators iteratively applied over the same $(\mathbf{x}, \mathcal{S})$ constitute relational functions. A single relational aggregator can be seen as a special case of relational function.

DEFINITION 5 (RELATIONAL FUNCTION). A relational function $f \in \mathcal{F}$ is defined as a composition of relational aggregators $f =$

¹In this work we use 1-neighborhood and egonet interchangeably

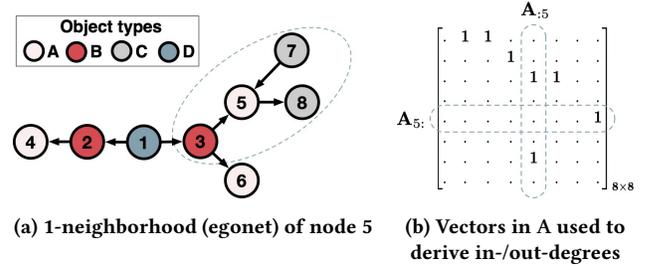


Figure 3: Derivation of graph statistics by applying order-1 relational functions on egonets. In this unweighted graph, $\text{out-deg}(5) = \sum(A_{5,:}, \Gamma)$ and $\text{in-deg}(5) = \sum(A_{:,5}, \Gamma)$ regardless object types.

$(\phi_1 \circ \dots \circ \phi_{h-1} \circ \phi_h)(\mathbf{x}, \mathcal{S})$ applied to feature values in \mathbf{x} associated with the set of related nodes \mathcal{S} . We say that f is order- h iff the feature vector \mathbf{x} is applied to h relational aggregators.

Figure 3 gives an example of applying order-1 relational functions on the egonet to derive graph statistics regardless of object types. One could also derive features related to graph heterogeneity by changing the context \mathcal{S} . For example in Figure 3a, the “number of neighbors in type B that point to node 5” can be capture as $\sum(A_{5,:}, \Gamma_B^{1-})$, where Γ_B^{1-} denotes the set of nodes with type B that are in the in-neighborhood. Together, relational aggregators and relational functions comprise the building blocks of our proposed method, MULTI-LENS.

3.2 Multi-level Structure Extraction

To extract subgraph features at different distances, the first step is to generate a set of node-level structural features via the so-called base graph functions \mathcal{F}_b . Based on \mathcal{F}_b , MULTI-LENS then composes new functions by iteratively applying a set of relational aggregators to generate new features, and describe them by the distribution to avoid overwhelming values. The output distribution-represented features form matrix Y as the input to summarization.

Base Graph Functions: As a special relational function, each base graph function $f_b \in \mathcal{F}_b$ is constituted by relational aggregators

that perform on an initial feature vector $\mathbf{b} \in \mathcal{B}$, where \mathbf{b} denotes the type of initial node-specific feature vector. The vector \mathbf{b} could be given as the row/column of the adjacency matrix corresponding to node i (as depicted in Fig. 3), or some other derived vector related to the node (e.g., its distance or influence to every node in the graph). Following [34], MULTI-LENS adopts $f_b = \sum$ to capture simple base features such as in/out/total degrees, but additional relational functions with arbitrary orders can also be leveraged. We denote applying the same base function to the egonets of all the nodes in graph G as follows:

$$f_b \langle \mathbf{b}, \Gamma^1 \rangle = [f_b(\mathbf{b}, \Gamma(1)), f_b(\mathbf{b}, \Gamma(2)), \dots, f_b(\mathbf{b}, \Gamma(N))]^T, \mathbf{b} \in \mathcal{B} \quad (1)$$

which forms an $N \times 1$ vector. For example, $f_b = \sum \langle \mathbf{A}_v, \Gamma \rangle$ enumerates the out-degree of all nodes in graph G . By applying f_b on each initial feature \mathbf{b} , we obtain the $N \times B$ base matrix $\mathbf{X}^{(0)}$:

$$\mathbf{X}^{(0)} = [f_b \langle \mathbf{b}_1, \Gamma \rangle, f_b \langle \mathbf{b}_2, \Gamma \rangle, \dots, f_b \langle \mathbf{b}_B, \Gamma \rangle], \mathbf{b}_{1 \dots B} \in \mathcal{B} \quad (2)$$

which intuitively captures different structural properties of the nodes within the 1-neighborhood Γ (shown in Fig. 2).

Relational Function Compositions: To capture high-order structural information at different distances in the graph, MULTI-LENS iteratively applies relational aggregators Φ on the columns of the base matrix $\mathbf{X}^{(0)}$. This process results in the multi-level structural representations of the graph as $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots, \mathbf{X}^{(L)}$, with j -th column defined as

$$\mathbf{X}_{:j}^{(l)} = \underbrace{\phi \circ \dots \circ \phi}_{l-1 \text{ times}} \langle \mathbf{X}_{:j}^{(0)}, \Gamma \rangle, \phi \in \Phi \quad (3)$$

At a given level, MULTI-LENS applies R aggregators Φ to the previously obtained $\mathbf{X}^{(l-1)}$, so the dimension of $\mathbf{X}^{(l)}$ is $N \times BR^l$. Although the dimension grows exponentially, [20] indicates that real-world graphs are dense with small diameters, which leads to the max composition with order $L \leq 5$. Intuitively, the l -order composition of relational functions captures graph features associated with the l -order node-centric subgraphs. As an example depicted in Fig. 4a, the \max aggregator captures the maximum of a specific feature \mathbf{x} in $\Gamma(i)$ for all nodes $i \in V$. Together these maximums form a new feature vector $\max \langle \mathbf{x}, \Gamma \rangle$. Subsequently in Fig. 4b, when \max is applied again to the resultant vector associated with nodes in $\Gamma(1)$, the maximum values from context $\Gamma(2)$ and $\Gamma(3)$ are aggregated and compared, which is equivalent to outputting the maximum in $\Gamma^2(1)$. Therefore, we can efficiently explore higher-order structural information for node i in its "broader" neighborhood by recursively applying the relational functions only over its egonet $\Gamma(i)$, without involving the exponential complexity of traversing higher-order subgraphs.

The particular order in which relational aggregators are applied is crucial in Eq. (3): the feature vector $\mathbf{x}^{(l)}$ derived from $\mathbf{x}^{(0)}$ will be semantically different if the same set of l relational aggregators are

Table 2: A few examples of relational aggregators.

Φ	Definition	Φ	Definition
\max/\min	$\max/\min_{i \in S} \mathbf{x}_i$	variance	$\frac{1}{ S } \sum_{i \in S} \mathbf{x}_i^2 - (\frac{1}{ S } \sum_{i \in S} \mathbf{x}_i)^2$
sum	$\sum_{i \in S} \mathbf{x}_i$	$l1\text{-distance}$	$\sum_{j \in S} x_i - x_j $
mean	$\frac{1}{ S } \sum_{i \in S} \mathbf{x}_i$	$l2\text{-distance}$	$\sum_{j \in S} (x_i - x_j)^2$

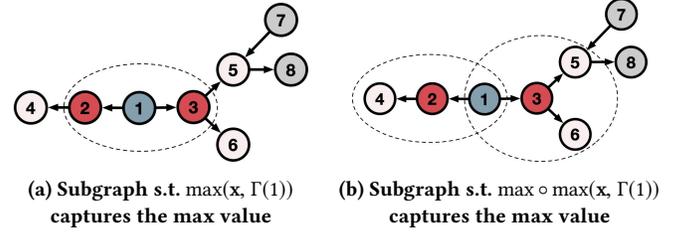


Figure 4: The composition of relational functions incorporates node features in expanded subgraphs (contexts). (a): $\max(\mathbf{x}, \Gamma(1))$ returns the maximum value of vector \mathbf{x} among nodes $\{1, 2, 3\}$. (b): $\max \circ \max(\mathbf{x}, \Gamma(1))$ returns the maximum value among the 2-neighborhood of node 1: $\{1, 2, 3, 4, 5, 6\}$.

applied in a different order. Examples of aggregators in Φ include *mean*, *variance*, *sum*, *max*, *min*, *l1-distance*, and *l2-distance*, etc. as shown in Table 2.

Handling skewed data: The l -order composed relational functions in \mathcal{F} gather l -order node context into the feature matrix $\mathbf{X}^{(l)}$, and MULTI-LENS aims to summarize $\mathbf{X}^{(l)}$ into a latent space. However, specific feature values in $\mathbf{X}^{(l)}$ could be overwhelming due to the power-law nature of real-world graphs, e.g., total degree, leading to under-representation of other features in the summary. To handle the skewness in our data, we describe the $N \times 1$ feature vector by the distribution of its feature values, on which we apply logarithmic binning [7, 16]. For feature vector \mathbf{x} , a set of nodes S' and c bins, logarithmic binning returns a vector of length c :

$$\Psi(\mathbf{x}, S', c) = [C(0), C(1), \dots, C(\log_a(c))] \quad (4)$$

where $C(z) = \sum_{i \in S'} \delta(z, \mathbf{x}_i)$, δ is the Kronecker delta (a.k.a indicator) function, a is the logarithm base, and $c = \max\{\max \mathbf{x}, c\}$. We set c to be the maximum between the explicitly given value (c) and the maximum feature value ($\max \mathbf{x}$) regardless of object types to make sure that the output bin counts remain the same across all features. We can explicitly fill in 0s in Eq. (4) in the case of $c > \max \mathbf{x}$.

Similar to Eq. (1), we use $\Psi(\mathbf{x}, S', c)$ to denote the process of applying Ψ function over all nodes in V (rows of \mathbf{X}) to get the $N \times c$ log-distribution feature matrix. Further, we denote the process of applying Ψ on all feature vectors (columns of \mathbf{X}) as $\mathbf{Y} = \Psi(\mathbf{X}, S', c)$. We use \mathbf{Y} to denote the resultant distribution-based feature matrix. In the next subsection, we will explain how to apply Ψ on different local subsets $S' \subseteq \Gamma$ in order to incorporate heterogeneity in the summary.

3.3 Summarizing Heterogeneity

So far values in \mathbf{X} are derived without considering node types, edge directionality, etc., which makes the features captured inaccurate. For example in an directed email exchange network, spammers with high out-degrees but trivial in-degrees could be incorrectly identified as important clients if edge directionality is not captured. Therefore, to get the compressed representation that captures the complexity to satisfy (P1) of the arbitrary input network G , we summarize individual node behaviors with respect to **a)** object types, **b)** edge directionality and **c)** edges types. Based on existing methods that describe a graph with multiple feature distributions [16],

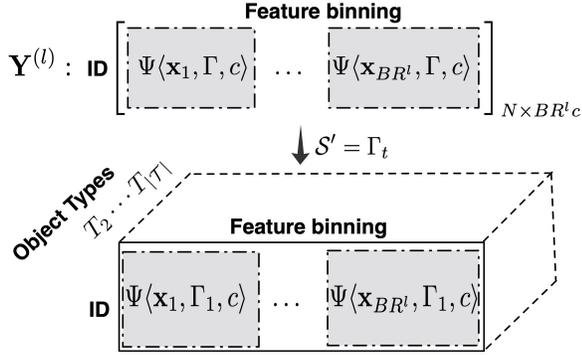


Figure 5: At level l , enumerating object types in Γ introduces one more dimension to the feature matrix and leads to a tensor. Note that $\Psi\langle\mathbf{x}, \Gamma, c\rangle$ is an $N \times c$ matrix (shaded area). Each layer of the tensor is also denoted as $\Psi\langle\mathbf{X}, \Gamma_t, c\rangle$ for brevity.

MULTI-LENS describes each individual node i by leveraging the distributions of feature values in \mathbf{X} associated with different contexts such as neighbors with the same type, or neighbors pointed by i , etc. to address heterogeneity **a-c**).

Object types: In heterogeneous graphs, the interaction patterns between a node and its neighbors of a specific type reveal important behavioral information. Intuitively, similar entities have similar interaction patterns with every single type of neighbor. For example, in the author-paper-venue networks, authors submitting papers to the same track at the same conference have higher similarity than authors submitting to different tracks at the same conference. To describe how a specific node i interacts with objects of type t , MULTI-LENS collects neighbors in that type by setting $\mathcal{S}' = \Gamma_t(i)$ and computes the “localized” distribution of a specific feature vector \mathbf{x} through $\Psi\langle\mathbf{x}, \Gamma_t(i), c\rangle$. Repeating this process for nodes $i \in V$ forms an $N \times c$ distribution matrix $\Psi\langle\mathbf{X}, \Gamma_t(i), c\rangle$.

MULTI-LENS enumerates all types of neighbors within Γ to incorporate complete interaction patterns for each node in the graph. This process can be seen as introducing one more dimension, the object types, to \mathbf{Y} to form a tensor, as shown in Fig. 5. We flatten the tensor through horizontal concatenation and denote it as \mathbf{Y}_{ot} :

$$\mathbf{Y}_{\text{ot}} = [\Psi\langle\mathbf{X}, \Gamma_{T_1}, c\rangle, \Psi\langle\mathbf{X}, \Gamma_{T_2}, c\rangle, \dots, \Psi\langle\mathbf{X}, \Gamma_{T_{|\mathcal{T}_V|}}, c\rangle] \quad (5)$$

Edge directionality: So far we assume the input graph is undirected by setting $\mathcal{S}' = \Gamma$ and search for neighbors in the 1-hop neighborhood regardless of edge directions. MULTI-LENS handles the directed input graphs by differentiating nodes from the out-neighborhood and in-neighborhood. The process is almost identical to the undirected case, but instead of setting $\mathcal{S}' = \Gamma$, we consider its two disjoint subsets Γ^+ and Γ^- with incoming and outgoing edges, respectively. The resultant distribution-based feature matrices are denoted as \mathbf{Y}_{ot}^+ and \mathbf{Y}_{ot}^- , respectively. Again, we horizontally concatenate them to get the feature matrix incorporating edge directionality \mathbf{Y}_{ed} as $\mathbf{Y}_{\text{ed}} = [\mathbf{Y}_{\text{ot}}^+, \mathbf{Y}_{\text{ot}}^-]$.

Edge types: Edge types in heterogeneous graphs play an important role in determining graph semantics and structure. The same connection between a pair of nodes with different edge types could convey entirely different meanings (e.g., an edge could indicate

“retweet” or “reply” in a Twitter-communication network). To handle this graph model, MULTI-LENS constructs subgraphs $g(V, E_t)$ restricted to a specific edge type $t \in \mathcal{T}_E$. For each subgraph, MULTI-LENS repeats the process to obtain the corresponding feature matrix \mathbf{Y}_{ed} per edge type that incorporates both node types and edge directionality, and concatenates them horizontally to obtain the final representation which we denote as:

$$\mathbf{Y}_{\text{et}} = [\mathbf{Y}_{\text{ed}1}, \mathbf{Y}_{\text{ed}2}, \dots, \mathbf{Y}_{\text{ed}|\mathcal{T}_E|}] \quad (6)$$

Therefore, \mathbf{Y}_{et} is of size $N \times q^{(l)}$ where $q^{(l)} = 2|\mathcal{T}_V||\mathcal{T}_E|c \cdot BR^l$.

In our proposed method, both object types and edge directionality are considered by default to construct \mathbf{Y} . We consider edge types only when the input is a multi-layer graph model or explicitly indicated.

3.4 MULTI-LENS Summarization

The important properties of MULTI-LENS that handles latent network summarization (Sec. 2) include space-efficiency (**P2**), being inductive (**P3**), and the ability to derive node-embeddings on the fly (**P4**). Although methods deriving compressed graph representation that satisfy (**P2**) alone are widely studied, none of them supports (**P4**). The main challenge is that the incurred information loss in the compression does not suffice to derive individual node behavioral representation. On the other hand, as indicated in [29], existing works that leverage the skip-gram model to derive node-embeddings can be seen as implicit low-rank factorization of a certain node-context matrix as $\mathbf{M} = \mathbf{U}\mathbf{H}^T$. A natural idea is to make use of the rank-context factor, \mathbf{H}^T as the summary, so that the node-rank factor \mathbf{U} can be derived as the embeddings. However, this requires storing \mathbf{M} in the first place, which violates (**P2**), and cannot handle transfer learning tasks (**P3**).

MULTI-LENS leverages $\mathbf{Y}_{\text{et}}^{(l)}$ as the node-context matrix capturing heterogeneity at level l and explicitly performs SVD to get a low-rank compressed graph representation $\mathbf{H}^{(l)}$:

$$\mathbf{H}^{(l)} = \sqrt{\Sigma_r^{(l)}} \mathbf{V}_r^{(l)T} \quad (7)$$

where $\sqrt{\Sigma_r^{(l)}}$ and $\mathbf{V}_r^{(l)T}$ are the square root of the singular values of $\mathbf{Y}_{\text{et}}^{(l)}$, and its right singular vectors, respectively. Note that any dimensionality reduction techniques can be adopted as well, such as non-negative matrix factorization. MULTI-LENS collects and stores $\mathbf{H}^{(l)}$ into \mathcal{H} as a part of the latent summary. Note that there is no need to maintain the dense embedding matrix $\mathbf{U}^{(l)}$, since we can estimate it directly given $\mathbf{Y}_{\text{et}}^{(l)}$ and the stored $\mathbf{H}^{(l)}$. Further, we do not need to maintain $\mathbf{Y}_{\text{et}}^{(l)}$, since we can derive it on the fly given the relational function compositions in \mathcal{F} . In practice, both terms can be derived efficiently through SVD because $\mathbf{Y}_{\text{et}}^{(l)}$ is sparse and $\mathbf{H}^{(l)}$ is of low-ranks (see Sec. 5 for more details). Finally, all the intermediate matrices derived do not need to be stored and can be estimated directly using \mathcal{F} .

Therefore, MULTI-LENS outputs a summary \mathcal{J} that consists of two elements: (1) a set of relational functions \mathcal{F} , which is used to capture node-wise structural information that applies to arbitrary networks and (2) a set of low-rank factorized matrices, \mathcal{H} , which captures the significant features in subgraphs at different distances

in a latent space. Since functions in \mathcal{F} and factorized matrices in \mathcal{H} are independent of the nodes or edges of the input graph, both require trivial storage and achieve space efficiency (P2). Furthermore, unlike existing embedding methods we do *not* need to store a node-wise representation matrix $\mathbf{U} \in \mathbb{R}^{N \times K}$, which is space-consuming and depends on the network size. As we will discuss next, elements in \mathcal{F} and \mathcal{H} together handle inductive learning tasks (P3) and the derivation of specific node embeddings (P4).

Inductive Summaries (P3): Given the set of relational functions \mathcal{F} , MULTI-LENS automatically composes new relational functions to capture structural features that are transferable. Thus, the factorized matrices \mathcal{H} learned on G can be transferred to another graph G' for inductive learning tasks. Suppose we learn the factorized matrix \mathcal{H} from G using \mathcal{F} . Equation (8) illustrates the process of inductively learning node embeddings on a new, previously unseen graph G' .

$$\mathbf{U}^{(l)} = \mathbf{Y}^{(l)}(\mathbf{H}^{(l)})^\dagger \quad (8)$$

where $\mathbf{H}^{(l)} \in \mathcal{H}$ is the multi-level factorized summary matrix learned on G , $(\mathbf{H}^{(l)})^\dagger$ denotes the pseudo-inverse of $\mathbf{H}^{(l)}$ and $\mathbf{Y}^{(l)}$ is obtained via applying \mathcal{F} to G' . The pseudo-inverse, $(\mathbf{H}^{(l)})^\dagger$ can be computed efficiently through SVD as long as the rank of $\mathbf{H}^{(l)}$ are limited (e.g., empirically setting $K_l \leq 128$).

To apply Eq. (8), one needs to ensure the dimensions of $\mathbf{H}^{(l)}$ “fit” in the equation by setting $r^{(l)} = r^{(l)}$ and $c = c'$ at each level l . In essence, this is equivalent to set the dimensions of $\mathbf{H}^{(l)}$ to be identical to $\mathbf{H}^{(l)}$, which is the output given by applying MULTI-LENS on G' . By setting $r^{(l)} = r^{(l)}$, MULTI-LENS captures the same number of major graph-level components at each level of the relational compositions, and by setting $c = c'$, the numbers of bins in the distribution-based feature matrices are set identical at all levels. Both criteria can be easily satisfied as we can record the order of ranks at each level, and manually set c large enough (e.g., maximum of all feature values) since all-0 columns can be filtered out automatically. The embeddings learned in the inductive scheme capture the node-wise structural difference between graphs, which can naturally be applied to domain-specific graph mining and time-evolving analysis. We showcase a temporal anomaly detection application in Sec. 6.4.

On-the-fly embedding derivation (P4): MULTI-LENS relies on relational functions \mathcal{F} to derive graph feature matrices $\mathbf{Y}^{(l)}$ at different levels. As indicated in [29], the embeddings matrix at each level, $\mathbf{U}^{(l)}$ can be computed explicitly through SVD such that $\mathbf{U}^{(l)} = \mathbf{U}_r^{(l)} \sqrt{\Sigma_r^{(l)}}$. We concatenate $\mathbf{U}^{(l)}$ given as output at each level to form the final node embeddings. The dimension of embeddings is determined by rank r at level l : we explicitly set $r = K_l$, so $\sum_{l=1}^L K_l = K$ determines the final embedding dimension.

Multi-level Summarization: MULTI-LENS explores node structural similarity based on the assumption that similar nodes should have similar structural behaviors in *node-centric subgraphs* at different distances. As mentioned in Sec. 3.2, matrix $\mathbf{X}^{(l)}$ captures feature values associated with the l -order node-centric subgraphs of G , and $\mathbf{Y}_{\text{et}}^{(l)}$ is its distribution-based representation that incorporates heterogeneity. By applying Eq. (7) to each $\mathbf{Y}^{(l)}$, we obtain the hierarchical summarization of G , as these subgraphs are explored differently

Algorithm 1 Latent Graph Summarization

Require:

(un)directed heterogeneous graph G ,
a set of relational aggregators $\Phi = \{\phi_1, \dots, \phi_R\}$,
Layer-wise embedding dimension constants $K_l, \sum_{l=1}^L K_l = K$
a constant c indicating the #bins

```

1:  $\mathcal{F} \leftarrow f_b$  ▷ Base graph functions: Eq. (1)
2: Initialize  $\mathbf{X}^{(0)}$  ▷ Eq. (2)
3: for  $l = 1, \dots, L$  do ▷ layers in the multi-level summarization
4:   Initialize  $\mathbf{X}^{(l)}$ 
5:   for  $r = 1, \dots, R$  do ▷ relational aggregators  $\{\phi_1, \dots, \phi_R\}$ 
6:     parallel for  $j = 1, \dots, BR^l$  do ▷ Columns in  $\mathbf{X}^{(l)}$ 
7:        $\mathbf{X}^{(l)} = \mathbf{X}^{(l)} \cup \phi_r(\mathbf{X}_j^{(l-1)}, \Gamma)$  ▷ Feature concatenation
8:       Solve Eq. (6) to derive  $\mathbf{Y}_{\text{et}}^{(l)}$  ▷ Capture heterogeneity
9:       Solve Eq. (7) by feeding  $\mathbf{Y}_{\text{et}}^{(l)}$  to obtain  $\mathbf{H}^{(l)}$ 
10:      Add the relational function compositions to  $\mathcal{F}$ 
11:       $\mathcal{H} \leftarrow \mathcal{H} \cup \mathbf{H}^{(l)}$ 
12: return summarization  $\mathcal{J} = \{\mathcal{F}, \mathcal{H}\}$ 

```

but coherently. Higher-order contains more “global” structural information at the cost of higher complexity, while low-order focuses on more “local” information with lower complexity. As a result, the derived embedding matrix \mathbf{U}_l also preserves node similarity at multiple levels.

We provide a general overview of MULTI-LENS in Algorithm 1.

4 DISCUSSION

Here we discuss the generalizations of our proposed approach to labeled and attributed graphs. It is straightforward to see that homogeneous, bipartite, signed, and labeled graphs are all special cases of heterogeneous graphs with $|\mathcal{T}_V| = |\mathcal{T}_E| = 1$ types, $|\mathcal{T}_V| = 2$ and $|\mathcal{T}_E| = 1$ types, $|\mathcal{T}_V| = 1$ and $|\mathcal{T}_E| = 2$ types, and $\{1, \dots, k\}$ node (or edge) types, respectively. Other special cases include k -partite and attributed graphs. Therefore, our approach naturally generalizes to all of these graphs. MULTI-LENS also naturally supports attributed graphs that have multiple attributes per node or edge (instead of a single label): Given an initial set of attributes organized in an attribute matrix $\mathbf{X}^{(b)}$, we can concatenate $\mathbf{X}^{(b)}$ with the base attribute matrix and apply our approach as before. Alternatively, we can transform the graph into a labeled one by applying a labeling function $\xi: \mathbf{x} \rightarrow y$ that maps every node’s attribute vector \mathbf{x} to a label y [1].

5 THEORETICAL ANALYSIS

Computational Complexity. The time complexity of MULTI-LENS includes deriving distribution-based matrix representation \mathbf{Y}_{et} and the rank- r approximation. Based on the sparsity of \mathbf{Y}_{et} , MULTI-LENS can perform SVD efficiently through fast Monte-Carlo Algorithm by extracting the most significant K singular values [10] with computational complexity $O(KN^2)$.

LEMMA 5.1. *The computational complexity of MULTI-LENS is*

$$O((cBR^L|\mathcal{T}_V||\mathcal{T}_E| + K^2)N + M).$$

PROOF. The computational unit of MULTI-LENS is the relational operation performed over the egonet of a specific node. Searching the neighbors for all node $i \in V$ has complexity $O(N + M)$ through BFS and subsequent queries can be handled in constant time. So the complexity to derive Y_{et} is linear to the number of binned features computed through relational operations. As indicated in the edge type subsection from Sec. 3.3, the number of binned features is $BR^l \cdot 2|\mathcal{T}_V||\mathcal{T}_E|c$. The second part, low-rank approximation, can be accomplished in $O(K_l^2 N)$ by extracting the most significant K_l singular values at level l . Further, deriving all K singular values has $O(K^2 N)$ complexity as $\sum_{l=1}^L K_l^2 \leq (\sum_{l=1}^L K_l)^2 = K^2$. Therefore, the overall computational complexity is $O(BR^L N |\mathcal{T}_V||\mathcal{T}_E|c + K^2 N + M)$. Note that both R and L are small constants in our proposed method (e.g., $R = 7$ and $L \leq 5$). Our proposed method scales linearly with the number of nodes and edges ($N + M$) in graph G . \square

Space Complexity. The space needed by MULTI-LENS includes the set of relational functions \mathcal{F} and factorized feature matrices \mathcal{H} .

LEMMA 5.2. *The space complexity of MULTI-LENS is*

$$O(cKBR^L |\mathcal{T}_V||\mathcal{T}_E| + F).$$

PROOF. Storing the set of ordered compositions of relational functions in the summary requires space complexity $O(F)$. For the set of matrices \mathcal{H} , we store $\mathbf{H}^{(l)}$ at all L levels. As shown in the time complexity analysis, the number of binned features (columns) in Y_{et} over all levels is $2BR^L |\mathcal{T}_V||\mathcal{T}_E|c$, which includes incorporating $|\mathcal{T}_V|$ object types with both in-/out- directionality and all edge types. The size of the output summarization matrices is thus $O(KBR^L |\mathcal{T}_V||\mathcal{T}_E|c)$, which is independent to the network size, instead only capturing graph complexity. This satisfies the crucial property we desire from latent summarization (Def. 1). \square

Our proposed method is easy to parallelize as the relational functions are applied to the subgraphs of each node independently, and the feature values are computed locally.

6 EXPERIMENTS

We evaluate MULTI-LENS by answering the following research questions to address the desired properties: **(P1)** Can MULTI-LENS be applied to arbitrary networks? **(P2)** Is MULTI-LENS space-efficient? **(P3)** Can MULTI-LENS handle inductive learning tasks, and **(P4)** Does MULTI-LENS derive node embeddings on-the-fly? How well does the resultant embeddings perform specific machine learning tasks?

6.1 Experimental Setup

Data: We use a variety of real-world heterogeneous network data to address **(P1)**, for which we present statistics in Table 3.

- **Yahoo! Messenger Logs:** This heterogeneous network represents Yahoo! messenger user communication patterns, where edges indicate message exchanges. Users are also associated with the locations from which they have sent messages.

- **DBpedia:** This unweighted, heterogeneous network contains multiple bipartite relations between entities, such as producer-occupation, producer-works, works-genre.

Table 3: Heterogeneous graph statistics

Data	#Nodes	#Edges	#Node Types	Graph Type
yahoo-msg	100 058	1 057 050	2	weighted
dbpedia	495 936	921 710	4	unweighted
digg	283 183	4 742 055	2	unweighted
bibsonomy	977 914	3 754 828	3	weighted

- **Digg:** This heterogeneous network records the voting behaviors of users to stories they like. Node types include users and stories. Each edge represents one vote or a friendship.

- **Bibsonomy:** This k-partite network represents the behaviors of users assign tags to publications. Each user assigns multiple tags, and a publication could be tagged multiple times.

Baselines: We compare the proposed method with methods commonly used in graph summarization, matrix factorization and representation learning over networks, namely, they are: **(1)** Node aggregation or NA for short [2, 40], **(2)** Spectral clustering or SC [39], **(3)** LINE [38], **(4)** DeepWalk or DW [28], **(5)** Node2vec or n2vec [13], **(6)** struc2vec or s2vec [31], **(7)** DNGR [5], **(8)** GraRep or GR [4], **(9)** Metapath2vec or m2vec [8], and **(10)** AspEm [36]. To run baselines that do not explicitly support heterogeneous graphs, we align nodes of the input graph according to their object types and re-order the IDs to form the homogeneous representation. For clarity, we denote them with suffix “-H”. In node aggregation, CoSum [40] ran out of memory due to the computation of pairwise node similarity. We chose Louvain [2] as an alternative that scales to large graphs and forms the basis of many node aggregation methods.

Configuration: We run all experiments on Mac OS platform with 2.5GHz Intel Core i7 and 16GB memory. For fairness, we do not employ parallelization and terminate processes exceeding 1 day. We evaluate MULTI-LENS against the optimal performance achieved by the baselines: we use 2nd-LINE to incorporate 2-order proximity in the graph; we run node2vec with grid searching over $p, q \in \{0.25, 0.50, 1, 2, 4\}$ as mentioned in [13] and report the best. For GraRep, we set $k = 2$ to incorporate 2-step relational information. For DNGR, we follow the paper to set the random surfing probability $\alpha = 0.98$ and use a 3-layer neural network model where the hidden layer has 1024 nodes. Metapath2vec is more related to our proposed method as it is designed to handle heterogeneous graphs. We retain the same settings (number of walks = 1000, walk length = 100) to generate walk paths and adopt a similar the meta-path “Type 1-Type 2-Type 1” as the “A-P-A” schema as suggested in the paper. For MULTI-LENS, although arbitrary relational functions can be used, we use order-1 $f_b = \sum$ as the base graph function for simplicity in our experiments. To begin, we derive in-/out- and total degrees to construct the $N \times 3$ base feature matrix $\mathbf{X}^{(0)}$ denoted as $[f_b(\mathbf{b}_1, \Gamma), f_b(\mathbf{b}_2, \Gamma), f_b(\mathbf{b}_3, \Gamma)]$ where $\mathbf{b}_1 = \mathbf{A}_{:,i}$, $\mathbf{b}_2 = \mathbf{A}_{:,i}$, and $\mathbf{b}_3 = (\mathbf{A} + \mathbf{A}^T)_{:,i}$, for $i \in V$. We set $L = 1$ to construct order-2 relational functions (the base function f_b we use is in order-1) to equivalently incorporate 2-order proximity as LINE does, but we do not limit other methods to incorporate higher order proximity. We set the output dimensions of all node representations to be $K = 128$. All other settings are kept default.

Table 4: Link prediction: node embeddings derived by MULTI-LENS outperforms all baselines measured by every evaluation metric. Specifically, MULTI-LENS outperforms embedding baselines by 1.88% ~ 88.49% in AUC and 0.90% ~ 70.25% in ACC. It outperforms even more over the aggregation-based methods. *OOT = Out Of Time (24 hours), OOM = Out Of Memory (16GB)

Data	Metric	NA-H	SC-H	LINE-H	DW-H	n2vec-H	GR-H	s2vec-H	DNGR-H	m2vec	AspEm	MULTI-LENS
yahoo-msg	AUC	0.7189	0.5375	0.6745	0.7715	0.7830	0.7535			0.6708	0.5587	0.8349
	ACC	0.2811	0.5224	0.6269	0.6927	0.7036	0.6825			0.6164	0.5379	0.7456
	F1 macro	0.2343	0.5221	0.6265	0.6897	0.7016	0.6821			0.6145	0.5377	0.7433
dbpedia	AUC	0.6002	0.5211	0.9632	0.8739	0.8774				0.6364	0.5869	0.9822
	ACC	0.3998	0.5399	0.9111	0.8436	0.8436				0.5869	0.5860	0.9192
	F1 macro	0.2968	0.4539	0.9110	0.8402	0.8402				0.5860	0.5860	0.9192
digg	AUC	0.7199	0.6625	0.9405	0.9664	0.9681				0.9552	0.5644	0.9863
	ACC	0.2801	0.6512	0.8709	0.9023	0.9049				0.8891	0.5459	0.9528
	F1 macro	0.2660	0.6223	0.8709	0.9019	0.9046				0.8890	0.5459	0.9528
bibsonomy	AUC	0.7836	0.6694	0.9750	0.6172	0.6173				0.6127	0.5790	0.9913
	ACC	0.2164	0.6532	0.9350	0.5814	0.5816				0.5790	0.5772	0.9493
	F1 macro	0.2070	0.6064	0.9349	0.5781	0.5782				0.5772	0.5772	0.9493

Table 5: Comparison between embedding methods and latent network summarization in terms of storage required (MB) for the outputs. MULTI-LENS requires at least 79 times less storage than embedding methods.

Data	NA-H	SC-H	LINE-H	n2vec-H	DW-H	m2vec-H	AspEm	MULTI-LENS
yahoo	0.85	116.01	116.55	150.19	143.26	49.51	122.36	0.62
dbpedia	5.10	575.11	578.51	807.05	806.93	-	606.91	0.81
digg	2.55	328.42	330.91	458.38	448.37	140.33	346.51	0.54
bibson.	9.23	1134.05	1142.28	1614.36	1614.37	-	1196.85	0.75

6.2 Space Efficiency

To evaluate MULTI-LENS in terms of space efficiency (P2), we measure how well it compresses the large-scale heterogeneous graph data (in MB) and report the results in Table 5. For comparison, we also report the space required to store the output of other existing methods. We observe that MULTI-LENS requires by at least 79 times less space when compared to embedding methods. This is because the space required by latent graph summarization methods depend entirely on the complexity and ability to compress the graph. On the other hand, the node-aggregation approach takes less space by outputting an $N \times 1$ vector indicating which supernode each individual node is grouped into, but MULTI-LENS still outperforms this graph size-dependent representation. Specifically, GraRep is only runnable on the smallest yahoo-msg dataset and takes 121.72 MB to store the embeddings. As an aside, metapath2vec takes 82.155 GB to store the preprocessed random-walk files for yahoo-msg, and even more for digg.

6.3 Link Prediction in Heterogeneous Graphs

We show the effectiveness of MULTI-LENS in a link prediction task. We use logistic regression with default settings (regularization strength = 1.0, stopping criteria = 10^{-4}). An edge e_{ij} is represented by the concatenating the embeddings of its source and destination: $emb(e_{ij}) = [emb(i), emb(j)]$ as used in [34]. We learn node embeddings as follows: for each dataset $G(V, E)$, we create the sub-graph $G'(V, E')$ by keeping all the nodes but randomly removing $\sim 40\%$ edges. We run all methods on G' to get node embeddings and randomly select $10\%|E|$ edges as the training data. Out of the

removed edges, 25% ($10\%|E|$) are used as missing links for testing. We also randomly create the same amount of “fake edges” for both training and testing. Table 4 illustrates the prediction performance measured with AUC, ACC, and F1 macro scores.

We observe that MULTI-LENS outperforms all baselines measured by every evaluation metric. MULTI-LENS outperforms embedding baselines by 1.88% ~ 88.49% in AUC and 0.90% ~ 70.25% in ACC. MULTI-LENS outperforms the aggregation-based methods even more, as the latter methods generate $N \times 1$ vectors as the embedding to map each node to a supernode. For runnable baselines designed for node embeddings in homogeneous graphs (baseline3 to baseline 8), the output result is expected as MULTI-LENS incorporates heterogeneous contextual information within 2-neighborhood in the node representation. Metapath2vec only finished running on the yahoo-msg (15.37 hrs) and digg (21.48 hrs) dataset within our time limit. However, it performs even worse than the other baselines. One possible reason is the mismatch between graph patterns and the meta-path we use: there are dense “inter type” interaction in the dataset, making the “crossing type” meta-path insufficient to capture the heterogeneity. This also applies to AspEm as one need to identify which aspect(s) to be selected in the embeddings. On the contrary, MULTI-LENS does not require the explicitly defined meta-schema, as it captures both directionality and object types automatically. A side note is that MULTI-LENS runs almost as fast as the methods designed for homogeneous graphs such as LINE and node2vec (122.509s for yahoo-message).

For GraRep, DNGR and Cosum, we encountered out-of-memory errors partially due to implementation (MATLAB), but more importantly due to the algorithmic complexity. GraRep requires expensive k -times matrix multiplication to obtain A^k , the k -step probability transition matrix. DNGR leverages multi-layer neural networks, which also involve significant amounts of computation to get non-linear relations among nodes. CoSum requires the similarity between nodes within and across object types as the input, rendering $O(N^2)$ computational complexity. For MULTI-LENS, however, the computational complexity of MULTI-LENS is linear to the number of nodes and edges $N + M$. Although learning higher levels could be expensive, empirical results show that when $L = 1$ MULTI-LENS could achieve promising results.

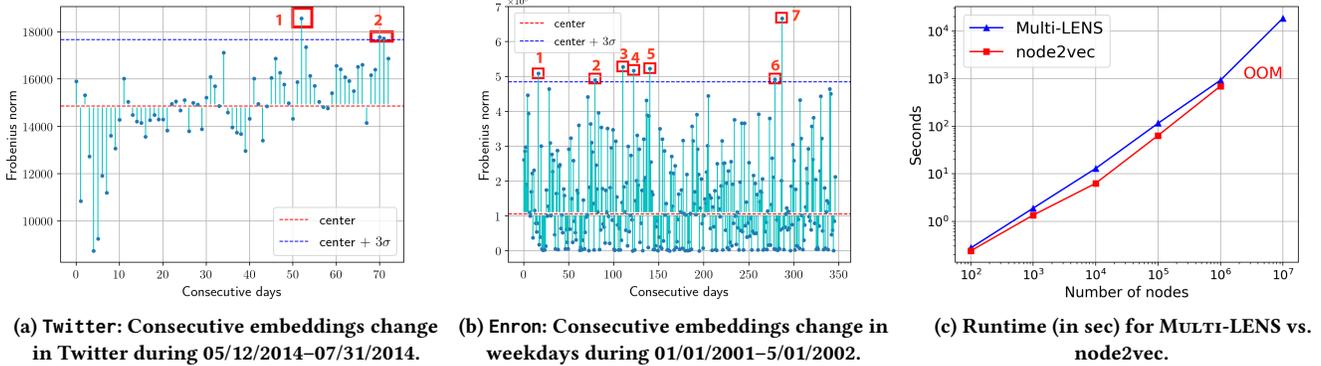


Figure 6: (a)-(b) Major event detection in real world datasets; (c) Runtime reported on Erdős-Rényi graphs with an average degree of 10. The proposed method scales similarly to node2vec with less memory requirement while node2vec runs out of memory on the graph with 10^7 nodes.

6.4 Inductive Anomaly Detection

We perform both anomalous subgraph detection and real-world graph event detection to evaluate how well MULTI-LENS handles inductive learning tasks (P3).

Anomalous Subgraph Detection: We first perform anomalous subgraph detection on both synthetic and real-world graphs. Following the literature [26], we first generate G_1, G_2 as the “background” graphs, and then induce an anomalous subgraph into G_2 by randomly selecting n nodes and adding $\binom{n}{2}p$ inter-connected edges to form an anomalous ER subgraph. We leverage the summarization \mathcal{J} learned from G_1 to detect the anomalous injection. Specifically in the synthetic setting, we generate two Erdős-Rényi (ER) graphs as G_1^{syn} and G_2^{syn} , each of which has 10^4 nodes and average degree 10 ($p^{\text{back}} = 10^{-3}$) as the background graphs. In the real-graph setting, we construct G_1^{real} and G_2^{real} using two consecutive daily data in the bibsonomy dataset. We leverage \mathcal{J} learned from G_1 to inductively learn the node embeddings of G_2 and compute their the node-wise Euclidean distances. We select the top- n deviating nodes as the reference to injected anomalies, and report detection precision in Table 6. It can be seen that MULTI-LENS gives promising results in almost every case except when the injection is very small. The reason for that these ER injections are so sparse and can be overwhelmed by the natural structural difference between G_1 and G_2 . When the injection becomes bigger, MULTI-LENS could effectively capture these anomalous behaviors.

Graph-based Event Detection: We also apply MULTI-LENS to handle event detection in real-world graphs. Our goal is to detect important events that are unfolding over time, which appear unusual or anomalous with respect to the global behavior of the complex network. The datasets we used are the Twitter² and Enron³ graphs. Twitter has totally 308 499 nodes and 2 601 834 edges lasting from 05/12/2014 to 07/31/2014, and Enron has totally 80848 nodes and 2 233 042 edges lasting from 01/01/2001 to 05/01/2002.

In Twitter, nodes represent entities such as keywords or hashtags appearing in Twitter and edges denote their coexistence. We

Table 6: Anomalous subgraphs (with n nodes and probability p) detection precision on both synthetic and real-world graphs. In the synthetic setting, Erdős-Rényi (ER) subgraphs are injected into the background ER graph G_2^{syn} with $n = 10^4$ and $p = 10^{-3}$. In the real-graph setting, anomalous ER graphs are injected into the background bibsonomy graph G_2^{real} .

$p \setminus n$	REAL GRAPH					SYNTHETIC GRAPH		
	100	200	300	400	500	50	75	100
0.1	0.200	0.780	0.950	0.973	0.980	0.06	0.3333	0.81
0.2	0.640	0.945	0.967	0.985	0.996	0.74	0.9867	1
0.3	0.870	0.960	0.990	0.995	0.996	1	1	1
0.4	0.890	0.975	0.993	0.995	1	1	1	1
0.5	0.920	0.990	0.993	1	1	1	1	1
0.6	0.920	0.990	0.997	1	1	1	1	1
0.7	0.940	0.990	1	1	1	1	1	1
0.8	0.960	0.990	1	1	1	1	1	1
0.9	0.980	1	1	1	1	1	1	1
1.0	0.980	1	1	1	1	1	1	1

capture serious events in terms of terrorism or domestic security happened between May to July, 2014. Enron denotes the message-exchange behavior between email addresses in the Enron corpus. We testify major events following the timeline of Enron scandal. Similar to the synthetic inductive analysis, we adopt the summary learned from graph G_{t-1} to get node embeddings of G_t . Fig. 6a shows the Frobenius norm between node embeddings in consecutive instances of the daily Twitter co-mentioning activity and the marked days are 3δ units away from the median value as indicated in [19]. Intuitively, huge distances between node embeddings of consecutive daily graphs indicate abrupt changes of graph structures, which can be seen as the sign of huge events. As shown in Fig. 6a, MULTI-LENS captures two events: (1) the Gaza-Israel conflict and (2) Ebola Virus Outbreak. Compared with other events in the same time period, the detected ones are most impactful in terms of the number of people affected, and the attraction they drew.

Similarly, by applying MULTI-LENS to Enron, we detect several events that are notable in the company’s history as studied in [19]. The events shown in Fig. 6b include (1) The quarterly conference

²<http://odds.cs.stonybrook.edu/twittersecurity-dataset/>

³<http://odds.cs.stonybrook.edu/enroninc-dataset/>

call where Jeffrey Skilling, Enron’s CEO, reports "outstanding" status of the company; (2) The infamous quarterly conference call; (3) FERC institutes price caps across the western United States; (4) The California energy crisis ends; (5) Skilling announces desire to resign to Kenneth Lay, founder of Enron; (6) Baxter, former Enron vice chairman, commits suicide, and (7) Enron executives Andrew Fastow and Michael Kopper invoke the Fifth Amendment before Congress.

6.5 Scalability & Runtime

In this experiment we evaluate MULTI-LENS in terms of its ability to derive node embeddings on the fly (P4).

Scalability: We first demonstrate the scalability of MULTI-LENS on the Erdős-Rényi graph with the number of nodes ranging from 10^2 to 10^7 . We compare its performance with node2vec, which is designed specifically for scalability. As shown in Fig. 6c, node2vec runs out of memory on the graph with 10^7 nodes, whereas MULTI-LENS scales almost as well as node2vec while using less space.

On-the-fly embeddings: We also measure the runtime of MULTI-LENS to derive specific node embeddings on the fly. We leverage the summary \mathcal{J} generated from graph G to derive the embeddings of 50 randomly sampled nodes on the same graph. In Table 7, we report the average runtime. GraRep takes 451.3s on the runnable yahoo-msg dataset, and fails to run on the remaining datasets. Thus, we omit it from Table 7. MULTI-LENS requires only a few seconds per node embedding, while *none* of the baselines can handle this task without generating the embeddings for *all* the nodes.

Table 7: Runtime to derive node embeddings on the fly. We report the average runtime to derive embeddings of 50 randomly sampled nodes using MULTI-LENS; all baselines must derive the embeddings of all nodes to get specific ones.

Data	NA-H	SC-H	LINE-H	n2vec-H	m2vec-H	AspEm	MULTI-LENS
Yahoo-msg	62.98	828.41	163.41	97.39	1941.50	11.25	1.18
dbpedia	162.75	368.26	218.84	160.43	-	25.14	4.81
digg	662.93	2006.83	201.53	676.04	43200.97	22.94	2.83
bibsonomy	347.79	832.62	252.98	1132.73	-	56.60	8.91

7 RELATED WORK

We qualitatively compare MULTI-LENS to summarization and embedding methods in Table 8.

Node embeddings. Node embedding or representation learning has been an active area which aims to construct node representations that preserve a notion of similarity over the graph [12, 34]. For instance, [5, 8, 13, 28, 38] define node similarity in terms of proximity (based on the adjacency or positive pointwise mutual information matrix) using random walks (RW) or deep neural networks (DNN). More relevant to MULTI-LENS are approaches to capture similar node behavioral patterns (roles) or structural similarity [1, 15, 33, 34]. For instance, struc2vec and xNetMF [15, 31] define similarity based on node degrees, while DeepGL [34] learns deep inductive relational functions applied to degree, triangle counts, or any other graph invariants. [21, 29] investigate theoretical connection between matrix factorization and the skip-gram architecture.

Table 8: Qualitative comparison of MULTI-LENS to existing summarization and embedding methods. Does the method: handle heterogeneous graphs; yield an output that is size-independent, but node-specific, and representations that are independent of node proximity; support inductive learning and scale well (i.e., it is subquadratic on the network size)?

	INPUT	REPRESENTATIONS / OUTPUT			METHOD	
	Heterogeneity	Size indep.	Node specific	Proximity indep.	Scalable	Inductive
Aggregation [2]	✓	✗	✗	✗	✓	✓
Cosum [40]	✗	✗	✗	✓	✗	✗
AspEm [36]	✓	✗	✓	✗	✓	✓
metapath2vec [8]	✓	✗	✓	✗	✓	✓
node2vec [13]	✗	✗	✓	✗	✓	✓
LINE [38]	✗	✗	✓	✗	✓	✓
struc2vec [31]	✗	✗	✓	✓	✗	✗
DNGR [5]	✗	✗	✓	✗	✗	✗
MULTI-LENS	✓	✓	✓	✓	✓	✓

To handle heterogeneous graphs, metapath2vec [8] captures semantic and structural information by performing RW on predefined metapaths. role2vec [1] proposes a framework for inductive learning by defining attributed RW atop relational aggregators. There are also works based on specific characteristics in heterogeneous graphs. For example, HEBE [14] proposes to embed heterogeneous networks with hyperedges that represent interactions among involving objects in events. Another work, AspEm, proposes to represent underlying semantic facets as multiple “aspects” and select a subset to embed based on datasetwide statistics. Unlike node embedding methods that generate dense embeddings of fixed dimensionality, MULTI-LENS provides general, compact and multi-level latent summaries for *any type* of graph, which can also be used to generate node embeddings effectively in an *inductive* manner without specifying extra heterogeneous characteristics.

Summarization. We give an overview of graph summarization methods, and refer the interested reader to tutorials [17, 23] and surveys [24]. Most graph summarization works fall into 3 categories: (1) aggregation-based which group nodes [27, 30] or edges [3, 25]) into super-nodes/edges based on application-oriented criteria or existing clustering algorithms; (2) abstraction-based which remove less informative nodes or edges; and (3) compression-based [18, 22, 35] which aim to minimize the number of bits required to store the input graph. Summarization methods have a variety of goals, including query efficiency [32, 37], pattern understanding [9, 18], storage reduction [22], interactive visualization [9], and domain-specific feature selection [16]. Most summarization works handling heterogeneous networks [6, 40] target specific graph mining tasks such as visualization. The only work relevant is CoSum [40], which tackles entity resolution by aggregating nodes into supernodes based on their labels and structural similarity. Unlike the aforementioned approaches, MULTI-LENS focuses on *inductive* summarization of *any type* of graph, and summarizes in a *latent* graph space (i.e., the final representation is independent of nodes/edges). Moreover, it is general and not tailored to specific ML tasks.

8 CONCLUSION

This work introduced the problem of *latent network summarization* and described a general computational framework, MULTI-LENS to learn such space-efficient latent node summaries of the graph that are completely independent of the size of the network. The

output (size) of latent network summarization depends only on the complexity and heterogeneity of the network, and captures its key structural behavior. Compared to embedding methods, the latent summaries generated by our proposed method require $79\times$ less storage, while achieving 2–89% improvement in AUC and up to 70% improvement in accuracy for the link prediction task. Overall, the experiments demonstrate the effectiveness of MULTI-LENS for link prediction, anomaly and event detection, as well as its scalability and space efficiency.

REFERENCES

- [1] Nesreen K. Ahmed, Ryan A. Rossi, Rong Zhou, John Boaz Lee, Xiangnan Kong, Theodore L. Willke, and Hoda Eldardiry. 2018. Learning Role-based Graph Embeddings. In *IJCAI StarAI*.
- [2] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefevre. 2008. Fast unfolding of communities in large networks. *JSTAT* 2008, 10 (2008), P10008.
- [3] Gregory Buehrer and Kumar Chellapilla. 2008. A scalable pattern mining approach to web graph compression with communities. In *WSDM*. ACM, 95–106.
- [4] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2015. Grarep: Learning graph representations with global structural information. In *CIKM*. 891–900.
- [5] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2016. Deep Neural Networks for Learning Graph Representations. In *AAAI*. 1145–1152.
- [6] Chen Chen, Cindy X Lin, Matt Fredrikson, Mihai Christodorescu, Xifeng Yan, and Jiawei Han. 2009. Mining graph patterns efficiently via randomized summaries. *VLDB* 2, 1 (2009), 742–753.
- [7] G. Cormode and S. Muthukrishnan. 2005. Space efficient mining of multigraph streams. In *Proceedings of the 24th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 271–282.
- [8] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *KDD*. 135–144.
- [9] Cody Dunne and Ben Shneiderman. 2013. Motif simplification: improving network visualization readability with fan, connector, and clique glyphs. In *CHI*.
- [10] Alan Frieze, Ravi Kannan, and Santosh Vempala. 2004. Fast Monte-Carlo algorithms for finding low-rank approximations. *Journal of the ACM (JACM)* 51, 6 (2004), 1025–1041.
- [11] Palash Goyal and Emilio Ferrara. 2017. Graph Embedding Techniques, Applications, and Performance: A Survey. *arXiv preprint arXiv:1705.02801* (2017).
- [12] Palash Goyal and Emilio Ferrara. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* 151 (2018), 78–94.
- [13] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *KDD*. 855–864.
- [14] Huan Gui, Jialu Liu, Fangbo Tao, Meng Jiang, Brandon Norick, and Jiawei Han. 2016. Large-scale embedding learning in heterogeneous event data. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*. IEEE, 907–912.
- [15] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. 2018. REGAL: Representation Learning-based Graph Alignment. In *CIKM*.
- [16] Di Jin and Danai Koutra. 2017. Exploratory analysis of graph data by leveraging domain knowledge. In *ICDM*. 187–196.
- [17] Arijit Khan, Sourav S Bhowmick, and Francesco Bonchi. 2017. Summarizing static and dynamic big graphs. *VLDB* 10, 12 (2017), 1981–1984.
- [18] Danai Koutra, U Kang, Jilles Vreeken, and Christos Faloutsos. 2015. Summarizing and understanding large graphs. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 8, 3 (2015), 183–202.
- [19] Danai Koutra, Joshua T Vogelstein, and Christos Faloutsos. 2013. Deltacon: A principled massive-graph similarity function. In *Proceedings of the 2013 SIAM International Conference on Data Mining*. SIAM, 162–170.
- [20] J. Leskovec, J. Kleinberg, and C. Faloutsos. 2005. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [21] Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *NIPS*. 2177–2185.
- [22] Yongsub Lim, U Kang, and Christos Faloutsos. 2014. Slashburn: Graph compression and mining beyond caveman communities. *TKDE* 26, 12 (2014), 3077–3089.
- [23] Shou-De Lin, Mi-Yen Yeh, and Cheng-Te Li. 2013. Sampling and Summarization for Social Networks. In *SDM*.
- [24] Yike Liu, Tara Safavi, Abhilash Dighe, and Danai Koutra. 2018. Graph Summarization Methods and Applications: A Survey. *CSUR* 51, 3 (2018), 62.
- [25] Antonio Maccioni and Daniel J Abadi. 2016. Scalable pattern matching over compressed graphs via dedensification. In *SIGKDD*. ACM, 1755–1764.
- [26] Benjamin A Miller, Michelle S Beard, Patrick J Wolfe, and Nadya T Bliss. 2015. A spectral framework for anomalous subgraph detection. *IEEE Transactions on Signal Processing* 63, 16 (2015), 4191–4206.
- [27] Saket Navlakha, Rajeev Rastogi, and Nisheeth Shrivastava. 2008. Graph summarization with bounded error. In *SIGMOD*. 419–432.
- [28] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *KDD*. 701–710.
- [29] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. 2018. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *WSDM*. 459–467.
- [30] Sriram Raghavan and Hector Garcia-Molina. 2003. Representing web graphs. In *ICDE*. 405–416.
- [31] Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. 2017. struc2vec: Learning node representations from structural identity. 385–394.
- [32] Matteo Riondato, David Garcia-Soriano, and Francesco Bonchi. 2017. Graph summarization with quality guarantees. *DMKD* 31, 2 (2017), 314–349.
- [33] Ryan A. Rossi and Nesreen K. Ahmed. 2015. Role Discovery in Networks. *Transactions on Knowledge and Data Engineering* 27, 4 (April 2015), 1112–1131.
- [34] Ryan A. Rossi, Rong Zhou, and Nesreen K. Ahmed. 2018. Deep Inductive Network Representation Learning. In *WWW BigNet*.
- [35] Neil Shah, Danai Koutra, Tianmin Zou, Brian Gallagher, and Christos Faloutsos. 2015. Timecrunch: Interpretable dynamic graph summarization. In *KDD*. 1055–1064.
- [36] Yu Shi, Huan Gui, Qi Zhu, Lance Kaplan, and Jiawei Han. 2018. AspEm: Embedding Learning by Aspects in Heterogeneous Information Networks. In *Proceedings of the 2018 SIAM International Conference on Data Mining*. SIAM, 144–152.
- [37] Qi Song, Yinghui Wu, and Xin Luna Dong. 2016. Mining summaries for knowledge graph search. In *ICDM*. 1215–1220.
- [38] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. LINE: Large-scale Information Network Embedding. In *WWW*. 1067–1077.
- [39] Lei Tang and Huan Liu. 2011. Leveraging social media networks for classification. *Data Mining and Knowledge Discovery* 23, 3 (2011), 447–478.
- [40] Linhong Zhu, Majid Ghasemi-Gol, Pedro Szekely, Aram Galstyan, and Craig A Knoblock. 2016. Unsupervised entity resolution on multi-type graphs. In *ISWC*. 649–667.