

NODE2BITS: Compact Time- and Attribute-aware Node Representations for User Stitching

Di Jin¹, Mark Heimann¹, Ryan A. Rossi², and Danai Koutra¹

¹ University of Michigan, Ann Arbor

² Adobe Research

Abstract. Identity stitching, the task of identifying and matching various online references (*e.g.*, sessions over different devices and timespans) to the same user in real-world web services, is crucial for personalization and recommendations. However, traditional user stitching approaches, such as grouping or blocking, require pairwise comparisons between a massive number of user activities, thus posing both computational and storage challenges. Recent works, which are often application-specific, heuristically seek to reduce the amount of comparisons, but they suffer from low precision and recall. To solve the problem in an application-independent way, we take a heterogeneous network-based approach in which users (nodes) interact with content (*e.g.*, sessions, websites), and may have attributes (*e.g.*, location). We propose NODE2BITS, an efficient framework that represents multi-dimensional features of node contexts with binary hashcodes. NODE2BITS leverages *feature-based temporal* walks to encapsulate short- and long-term interactions between nodes in heterogeneous web networks, and adopts SimHash to obtain compact, binary representations and avoid the quadratic complexity for similarity search. Extensive experiments on large-scale real networks show that NODE2BITS outperforms traditional techniques and existing works that generate real-valued embeddings by up to 5.16% in $F1$ score on user stitching, while taking only up to 1.56% as much storage.

1 Introduction

Personalization and recommendations increase user satisfaction by providing relevant experiences and handling the online information overload in news, web search, entertainment, and more. Accurately modeling user behavior and preferences over time are at the core of personalization. However, tracking user activity online is challenging as users interact with tens of internet-enabled devices from different locations daily, leading to fragmented user profiles. Without unified profiles, the observed user data are sparse, non-representative of the population, and insufficient for accurate predictions that drive business success.

In this work, we tackle the problem of *identity or user stitching*, which aims to identify and group together logged-in and anonymous sessions that correspond to the *same* user despite taking place across different channels, platforms, devices and browsers [30]. This problem is a form of entity or identity resolution [13,2],

also known as entity linking, record linkage, and duplicate detection [6,21,2]. Unlike entity resolution where textual information per user (*e.g.*, name, address) is available, identity stitching relies solely on user interactions with online content and web metadata. Although cookies can help stitch several different sessions of the same user, many users have multiple cookies (*e.g.*, a cookie for each device or web browser) [8], and most cookies expire after a short time, and therefore cannot help to stitch users over time. Similarly, IP addresses change across locations resulting in fragmentation or even erroneous stitching between users who have the same IP address at different times (*e.g.*, airports). Meanwhile, fingerprinting approaches [12] capture user similarity based on device or browser configurations, not on behavioral patterns that remain consistent across devices or browsers. On the other hand, exhaustive solutions for entity resolution require quadratic number of comparisons between all pairs of entities, which is computationally intractable for large-scale web services. This can be partially handled via the heuristic of blocking [24], which groups similar entity descriptions into blocks, and only compares entities within the same block.

To overcome these challenges and better tailor to the user stitching setup, our solution is based on the idea that the same user accesses *similar content* across platforms and has *similar behavior* over time. We model the user interactions with different content and platforms over time in a dynamic heterogeneous network, where user stitching maps to the identification of *nodes* that correspond to the same real-world entity. Motivated by the success of node representation learning, we aim to find embeddings of time-evolving ‘user profiles’ over this rich network of interactions. For large-scale graphs, however, the customary dense node representations for each node can often impose a formidable memory requirement, on par with that of the original (sparse) adjacency matrices [19]. Thus, to efficiently find *sparse* binary representations and link entities based on similar activity while avoiding the pairwise comparison of all user profiles, we solve the following problem:

Problem 1 (Temporal, Hash-based Node Embeddings). *Given a graph $G(V, E)$, the goal of hash-based network embedding is to learn a function $\chi : V \rightarrow \{0, 1\}^d$ such that the derived binary d -dimensional embeddings (1) preserve similarities in interactions in G , (2) are space-efficient, and (3) accurately capture temporal information and the heterogeneity of the underlying network.*

We introduce a *general* framework called NODE2BITS that captures temporally-valid interactions between nodes in a network, and constructs the contexts based on topological features and (optional) side information in the interaction. These

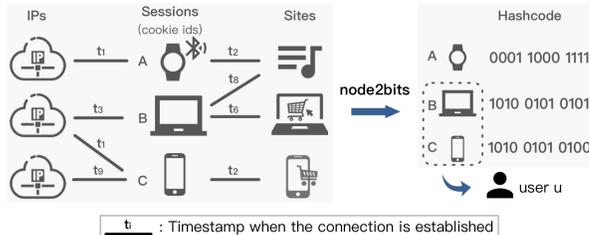


Fig. 1: NODE2BITS overview. NODE2BITS encodes the temporal, heterogeneous information of each node into binary hashcodes for efficient user stitching.

feature-based contexts are then turned into histograms that incorporate node type information at different *temporal distances*, and are mapped to binary hashcodes through SimHash [5]. Thanks to locality sensitive hashing [17], the hashcodes, which are time-, attribute- and structure-aware, preserve the similarities in temporal interaction patterns in the network, and achieve both space and computational efficiency for similarity search. Given these sparse, hash-based embeddings of all entities, we then cast user stitching as a supervised binary classification task or a hashing-based unsupervised task. As an example, in Fig. 1, devices B and C are associated with identical IPs and similar online sales websites visited afterwards, thus they are encoded similarly and could correspond to the same user.

Our contributions are:

- **Embedding-based Formulation:** Going beyond traditional blocking techniques, we formulate the problem of user stitching as the problem of finding temporal, hash-based embeddings in heterogeneous networks such that they maintain *similarities between user interactions over time*.
- **Space-efficient Embeddings:** We propose NODE2BITS, a practical, intuitive, and fast framework that generates *compact, binary embeddings* suitable for user stitching. Our method combines random walk-based sampling of contexts, their feature-based histogram representations, and locality sensitive hashing to preserve the heterogeneous equivalency of contexts over time.
- **Extensive Empirical Analysis:** Our experiments on real-world networks show that NODE2BITS outputs a space-efficient binary representation which uses between $63\times$ and $339\times$ less space than the baselines while achieving comparable or better performance in user stitching tasks. Moreover, NODE2BITS is scalable for large real-world temporal and heterogeneous networks.

For reproducibility, the code is at <https://github.com/GemsLab/node2bits>.

2 Preliminaries and Definitions

Before we introduce NODE2BITS, we discuss two key concepts that our method is based on: our dynamic heterogeneous network model, and temporal random walks. We give the main symbols and their definitions in Table 1.

2.1 Dynamic Heterogeneous Network Model

As we mentioned above, we model the user interactions with content, websites, devices etc. as a heterogeneous network, which is formally defined as:

Definition 1 (HETEROGENEOUS NETWORK). *A heterogeneous network $G = (V, E, \psi, \xi)$ is comprised of (i) a nodeset V and edgeset E , (ii) a mapping $\psi : V \rightarrow \mathcal{T}_V$ of nodes to their types, and (iii) a mapping $\xi : E \rightarrow \mathcal{T}_E$ to edge types.*

Many graph types are special cases of heterogeneous networks: **(1)** homogeneous graphs have $|\mathcal{T}_V| = |\mathcal{T}_E| = 1$ type; **(2)** k -partite graphs consist of $|\mathcal{T}_V| = k$ and $|\mathcal{T}_E| = k - 1$ types; **(3)** signed networks have $|\mathcal{T}_V| = 1$ and $|\mathcal{T}_E| = 2$ types; and **(4)** labeled graphs have a single label per node/edge.

Table 1: Summary of major symbols and their definitions.

Symbol	Definition
$G(V, E, \xi, \psi)$	(un)directed and (un)weighted heterogeneous network with nodeset V , edgeset E , a mapping ξ from nodes to node types, and an edge mapping ψ , resp.
$ V = N, E = M$	number of nodes and edges in G
$\mathcal{T}_V, \mathcal{T}_V ; \mathcal{T}_E, \mathcal{T}_E $	set of node/edge types in the heterogeneous graph and its size, resp.
\mathbf{F}	$N \times \mathcal{F} $ feature matrix including node attributes and derived features
$f_{ij}, f_{(j)}$	$(i, j)^{th}$ element of \mathbf{F} and index of its j^{th} feature, resp.
\mathcal{W}	set of random walks
$(\mathbf{w}_L)_{L \in \mathbb{N}}, \mathbf{w}_L[u]$	sequence of nodes in a random walk of length L , and the index of node u , resp.
L	the maximum length of a random walk
Δt	'temporal distance' in \mathcal{W} based on temporally ordered edge transitions
$\mathcal{C}_u^{\Delta t}, \mathcal{C}_u^{\Delta t} f$	context of node u at distance Δt , and feature-based context, resp.
$g_i : \mathcal{C} \rightarrow \{0, 1\}$	i^{th} LSH function that hashes a node context into a binary value
$K^{\Delta t}, K$	embedding dimension at distance Δt , and output dimension $K = \sum_{\Delta t=1}^{\text{MAX}} K^{\Delta t}$
$\mathbf{h}(\mathcal{S}), \mathbf{h}(\mathcal{S} \cdot)$	unconditional and conditional b -bin histogram of values in enclosed set \mathcal{S} , resp.
\mathbf{Z}	$N \times K$ output binary embeddings or hashcodes

Most real networks capture evolving processes (e.g., communication, browsing activity) and thus change over time. Instead of approximating a dynamic network as a sequence of lossy discrete static snapshots G_1, \dots, G_T , we model the *temporal interactions* in a *lossless* fashion as a *continuous-time dynamic network* [23].

Definition 2 (CONTINUOUS-TIME DYNAMIC NETWORK). *A continuous-time dynamic, heterogeneous network $G = (V, E_\tau, \psi, \xi, \tau)$ is a heterogeneous network with E_τ temporal edges between vertices V , where $\tau : E \rightarrow \mathbb{R}^+$ is a function that maps each edge to a corresponding timestamp.*

2.2 Temporal Random Walks

A walk on a graph is a sequence of nodes where each pair of successive nodes are connected by an edge. Popular network embedding methods generate walks using randomized procedures [25,14] to construct a corpus of node IDs or node context. In continuous-time dynamic networks, a *temporally valid* walk is defined as a sequence of nodes connected by edges with non-decreasing timestamps (e.g., representing the order that user-content interactions occurred) and they were first used for embeddings in [23].

Definition 3 (TEMPORAL WALK). *A temporal walk of length L from v_1 to v_L in graph $G = (V, E, \psi, \xi)$ is a sequence of vertices $\langle v_1, v_2, \dots, v_L \rangle$ such that $\langle v_i, v_{i+1} \rangle \in E_\tau$ for $1 \leq i < L$, and the timestamps are in valid temporal order: $\tau(v_i, v_{i+1}) \leq \tau(v_{i+1}, v_{i+2})$ for $1 \leq i < (L - 1)$.*

3 NODE2BITS: Hash-based Embedding Framework

Motivated by the task of user stitching, we aim to develop NODE2BITS to compactly describe each node/entity in the context of *realistic interactions* (Problem 1). Accordingly, NODE2BITS is required to: **(R1)** support heterogeneous

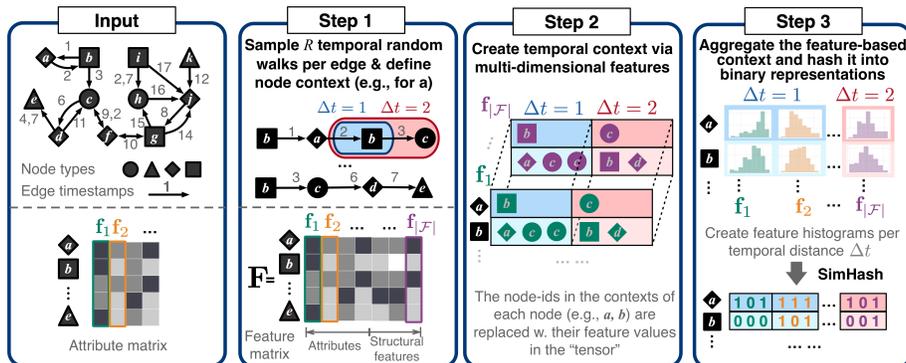


Fig. 2: NODE2BITS workflow. Given a graph and its attribute matrix (optional), NODE2BITS (1) samples *temporal* random walks to obtain sequences that respect time, derives contexts at different temporal distances (temporal contexts of a and b are derived from the walk $\{b, a, b, c\}$), as well as the feature matrix \mathbf{F} ; (2) creates temporal contexts based on multi-dimensional features in \mathbf{F} ; and (3) aggregates them into feature-based histograms to obtain sparse, binary, similarity-preserving embeddings via SimHash.

networks where the nodes and edges can be of any arbitrary type (*e.g.*, a user, web page, IP, tag, spatial location); (**R2**) preserve the temporal validity of the events and interactions in the data; (**R3**) scale in runtime to large networks with millions of nodes/edges; and (**R4**) scale in memory requirements with space-efficient yet powerful *binary* embeddings that capture ID-independent similarities. Next we detail the three main steps of NODE2BITS: (§ 3.1) Sampling temporal random walks and defining temporal contexts; (§ 3.2) Constructing temporal contexts based on multi-dimensional features; (§ 3.3) Aggregating and hashing contexts into sparse embeddings. We give the overview of NODE2BITS in Figure 2 and Algorithm 1.

3.1 Temporal Random Walk Sampling

The first step of NODE2BITS is to capture interactions in a node’s context, which is important for the user stitching task: instead of simple interactions corresponding to pairwise edges, it samples more complex interaction sequences via random walks. But unlike many existing representation learning approaches [25,14], our method samples *realistic* interactions in the order that they happen via L -step *temporal* random walks (Definition 3 [23]), thus satisfying requirement **R2**.

NODE2BITS defines the *temporal context* $\mathcal{C}_u^{\Delta t}$ of node u at temporal distance Δt as the *collection of entities* that are at Δt -hops away from node u in the sampled random walks. Formally:

$$\mathcal{C}_u^{\Delta t} = \{v : |\mathbf{w}_L[v] - \mathbf{w}_L[u]| = \Delta t, \forall \mathbf{w}_L \in \mathcal{W}\}, \quad (1)$$

where $\mathbf{w}_L[\cdot]$ is the index of the corresponding node in the random walk $(\mathbf{w}_L)_{L \in \mathbb{N}}$. For example, in Figure 2 (Step 1) the context of node a at temporal distance 2 is $\mathcal{C}_a^{\Delta t=2} = \{c\}$ (highlighted in red). Depending on the temporal context that we

want to capture, Δt can vary up to a *MAX* distance. Intuitively, small values of temporal distance capture more *direct* interactions and similarities between entities. In static graphs, Δt simply corresponds to the distance between nodes in the sampled sequences, without capturing any temporal information.

Temporal locality. The context that is defined above does not explicitly incorporate the *time elapsed* between consecutively sampled interactions. However, when modeling temporal user interactions, it is important to distinguish between short-term and long-term transitions. Inspired by [23], NODE2BITS accounts for the *closeness* or *locality* between consecutive contexts (*i.e.*, $\mathcal{C}_u^{\Delta t}$ and $\mathcal{C}_u^{\Delta t+1}$) through different biased temporal walk policies. For example, in the short-term policy, the transition probability from node u to v is given as the softmax function:

$$P[v|u] = \frac{\exp(-\tau(u,v)/d)}{\sum_{i \in \Gamma_\tau(u)} \exp(-\tau(u,i)/d)} \quad (2)$$

where $\tau()$ maps an edge to its timestamp, $d = \max_{e \in E_\tau} \tau(e) - \min_{e \in E_\tau} \tau(e)$ is the total duration of all timestamps, and $\Gamma_\tau(u)$ is the set of temporal neighbors reached from node u through temporally valid edges. Similarly, in the long-term policy, the transition probability from node u to v is given as in Equation (2) but with positive signs in the numerator and denominator.

3.2 Temporal Context based on Multi-dimensional Features

The context in Equation (1) depends on the node identities (IDs). However, in a multi-platform environment, a single entity may have multiple node IDs, thus contributing to seemingly different contexts. To generate ID-independent contexts that are appropriate for user stitching, we make the temporal contexts attribute- or feature-aware (**R1**), by building on the assumption that corresponding or similar entities have similar features. Formally, we assume that a network may have a set of input node attributes (*e.g.*, IP address, device type), as well as a set of derived topological features (*e.g.*, degree, PageRank), all of which are stored in a $N \times |\mathcal{F}|$ feature matrix \mathbf{F} (Figure 2, Step 1). We then generalize our random walks to not only respect time (**R2**) [23], but also capture this feature information using the notion of attributed/feature-based walks proposed in [1]:

Definition 4 (FEATURE-BASED TEMPORAL WALK). *A feature-based temporal walk of length L from node v_1 to v_L in graph G is defined as a sequence of feature values corresponding to the sequence of vertices in a valid temporal walk (Dfn. 3). For the j^{th} feature $f_{(j)}$, the corresponding feature-based temporal walk is*

$$\langle w_{L, f_{(j)}} \rangle_{L \in \mathbb{N}} = \langle f_{v_1, j}, f_{v_2, j}, \dots, f_{v_L, j} \rangle, \quad (3)$$

where $f_{v_i, j}$ is the value of the j^{th} feature for node v_i , stored in matrix \mathbf{F} .

Our definition is general as it allows walks to obey time while each node may have a d -dimensional vector of input attribute values and/or derived structural features, which can be discrete or real-valued [1].

Temporally-valid, multi-dimensional feature contexts. NODE2BITS extends the previously generated temporal contexts to incorporate node features and remove the dependency on node IDs. Following the definition of feature-based temporal walks, given $|\mathcal{F}|$ features, our method generates $|\mathcal{F}|$ -dimensional contexts per node u and temporal distance Δt by replacing the node IDs in Equation (1) with their corresponding feature values (Figure 2, Step 2). Formally, the temporally-valid, multi-dimensional feature contexts are defined as:

$$\mathcal{C}_u^{\Delta t}|f_{(j)} = \{f_{v,j} : \forall v \in \mathcal{C}_u^{\Delta t}\} \quad \forall \text{ feature } f_{(j)} \in \mathcal{F}, \quad (4)$$

where $f_{v,j}$ is the value of the j^{th} feature for node v .

3.3 Feature-based Context Aggregation and Hashing

The key insight in user stitching is that each user interacts with similarly ‘typed’ entities through similar relations over time: for example, in online-sales logs, a user likely browses similar types of goods in logged-in and anonymous sessions; and in online social networks, accounts sharing near-identical interaction patterns, such as replies or shares, are potentially from the same person. Based on this insight, NODE2BITS augments the previously generated temporal, multi-dimensional feature contexts with node types (and implicitly the corresponding relations or edge types), which is a key property of heterogeneous networks (**R1**). It subsequently aggregates them and derives similarity-preserving and space-efficient, binary entity representations (**R4**) via locality sensitive hashing.

Context Aggregation. Unlike existing works that aggregate contextual features into a single value such as mean or maximum [15,29], NODE2BITS aggregates them into *less lossy* representations: *histograms* tailored to heterogeneous networks by distinguishing between node types (**R1**). Specifically, it models the transitional dependency across node and relation types by further conditioning the derived contexts in Equation (4) on the node types $p_i \in \mathcal{T}_V$ (*i.e.*, each temporal context consists of the features of only one node type). We denote the temporal contexts conditioned on both features and node types as $\mathcal{C}_u^{\Delta t}|f, p$. The final histogram representation of node u at temporal distance Δt consists of the concatenation of the histograms over the conditional contexts at Δt (Figure 2, Step 3):

$$\mathbf{h}(\mathcal{C}_u^{\Delta t}) = [\mathbf{h}(\mathcal{C}_u^{\Delta t} | f_{(1)}, p_1), \mathbf{h}(\mathcal{C}_u^{\Delta t} | f_{(2)}, p_1), \dots, \mathbf{h}(\mathcal{C}_u^{\Delta t} | f_{(|\mathcal{F}|)}, p_{|\mathcal{T}_V|})]. \quad (5)$$

In this representation, the features are binned logarithmically to account for the often skewed distributions of structural features (*e.g.* degree). We note that the histograms can be further extended to edge types as shown in [19], for example by distinguishing pairs of nodes that are connected by multiple types of edges.

Similarity-preserving Representations via Hashing. Locality sensitive hashing (LSH) has been widely used for searching nearest neighbors in large-scale data mining [26]. In this work, we adopt SimHash [5] to obtain similarity-preserving and space-efficient representations (**R4**) for all the entities in the

heterogeneous network based on their aggregated time-, attribute-, and node type-aware contexts given by Equation (5).

Given a node-specific histogram $\mathbf{h}(\mathcal{C}_u^{\Delta t}) \in \mathbb{R}^d$ (with dimensionality $d = |\mathcal{F}||\mathcal{T}_V| \cdot b$, and b being the number of logarithmic bins for the features), SimHash generates a $K^{\Delta t}$ -dimensional³ binary hashcode or sketch $\mathbf{z}_u^{\Delta t}$ by projecting the histogram to $K^{\Delta t}$ random hyperplanes $\mathbf{r}_i \in \mathbb{R}^d$ as follows:

$$g_i(\mathbf{h}(\mathcal{C}_u^{\Delta t})) = \text{sign}(\mathbf{h}(\mathcal{C}_u^{\Delta t}) \cdot \mathbf{r}_i) \quad (6)$$

In practice, the hyperplanes do not need to be chosen uniformly at random from a multivariate normal distribution, but it suffices to choose them uniformly from $\{-1, 1\}^d$. The important property of locality sensitive hashing that guarantees that the similarities in the histogram space (which captures the temporal interactions between entities in G) are maintained is the following: for the SimHash family, the probability that a hash function agrees for two different vectors is equal to their cosine similarity. More formally, for two nodes u and v :

$$P(g_i(\mathbf{h}(\mathcal{C}_u^{\Delta t})) = g_i(\mathbf{h}(\mathcal{C}_v^{\Delta t}))) = 1 - \frac{\cos^{-1} \frac{\mathbf{h}(\mathcal{C}_u^{\Delta t}) \cdot \mathbf{h}(\mathcal{C}_v^{\Delta t})}{\|\mathbf{h}(\mathcal{C}_u^{\Delta t})\| \|\mathbf{h}(\mathcal{C}_v^{\Delta t})\|}}{180}. \quad (7)$$

In other words, the cosine similarity between nodes u and v in the context-space is projected to the sketch-space and can be measured by the cardinality of matching between $\mathbf{z}_u^{\Delta t}$ and $\mathbf{z}_v^{\Delta t}$, where $\mathbf{z}_\bullet^{\Delta t} = [g_1(\mathbf{h}(\mathcal{C}_\bullet^{\Delta t})), g_2(\mathbf{h}(\mathcal{C}_\bullet^{\Delta t})), \dots, g_{K^{\Delta t}}(\mathbf{h}(\mathcal{C}_\bullet^{\Delta t}))]$.

For each node u in G , the final binary representation is obtained by concatenating the hashcodes for contexts at different temporal distances Δt , resulting in a K -dimensional vector (since $K = \sum_{\Delta t=1}^{\text{MAX}} K^{\Delta t}$):

$$\mathbf{z}_u = [\mathbf{z}_u^{\Delta t=1} \ \mathbf{z}_u^{\Delta t=2} \ \dots \ \mathbf{z}_u^{\Delta t=\text{MAX}}] \quad (8)$$

where we replace the -1 bits with 0s to achieve a more space-efficient representation (**R4**). An example is shown in the second half of Step 3 in Figure 2, where the blue shades denote histograms and sketches for contexts in temporal distance $\Delta t = 1$, and red shades correspond to $\Delta t = 2$. Thus, the K -dimensional representation for each node, $\mathbf{z}_u \in \{0, 1\}^K$, captures the similarities between time-, feature- and node type-aware histograms across multiple temporal distances Δt . The similarity between two nodes' histograms can be quickly estimated as the proportion of common bits in their binary representations \mathbf{z}_\bullet .

Given these representations, we can perform user stitching by casting the problem as supervised binary classification or an unsupervised task based on the output of hashing (Equation (8)), which we discuss in § 4.1. Putting everything together, we give the pseudocode of NODE2BITS in Algorithm 1 and its detailed version (for reproducibility) in Appendix A. The runtime computational complexity of NODE2BITS is $\mathcal{O}(MRL + NK)$, which is linear to the number of edges when $M \gg N$ as K is relatively small (**R3**). The output space complexity is $\mathcal{O}(NK)$ -bit. NODE2BITS requires even less storage if the binary vectors are represented in the sparse format (see § 4.4 for empirical results). We provide detailed runtime complexity and space analysis in appendix B.

³ We assume that the length of each sketch at distance Δt is given as $K^{\Delta t} = \frac{K}{\text{MAX}}$.

Algorithm 1 NODE2BITS Framework

Require: (un)directed heterogeneous graph $G(V, E, \psi, \xi)$, # random walks R per edge, max walk length L , max temporal distance MAX, embedding dimensionality $K^{\Delta t}$ at dist. Δt

- 1 For each edge e , perform R temporal walks based on the short- or long-term policy (§ 3.1)
- 2 Obtain temporal contexts $\mathcal{C}_u^{\Delta t}$ for each node u at temporal distances $\Delta t \leq \text{MAX}$ via Eq. (1)
- 3 Construct feature matrix \mathbf{F} with node attributes (if avail.) and topological features (§ 3.2)
- 4 Derive feature-based temporal contexts $\mathcal{C}_u^{\Delta t}|f_{(j)}$ by replacing $v \in \mathcal{C}_u^{\Delta t}$ with the feature value $f_{v,j}$, as shown in Eq. (4)
- 5 **for each** temporal distance $\Delta t = 1, \dots, \text{MAX}$ and node $u \in V$ **do**
- 6 Obtain u 's final histogram $\mathbf{h}(\mathcal{C}_u^{\Delta t})$ over its contexts using Eq. (5)
- 7 Obtain a $K^{\Delta t}$ -dim, sparse, binary hashcode $z_u^{\Delta t}$ based on (modified) SimHash (§ 3.3)
- 8 Obtain the binary n2B embeddings \mathbf{z}_u of all nodes across temporal distances Δt via Eq. (8)
- 9 Perform (un)supervised user stitching via binary classification or hashing (§ 4.1,4.3)

4 Experiments

We perform extensive experiments on real-world heterogeneous networks to answer the following questions: **(Q1)** Is NODE2BITS effective in the user stitching task, and how does it compare to traditional stitching and embedding methods? (§ 4.2-4.3) **(Q2)** Does NODE2BITS have low space requirements, and is it more space-efficient than the baselines? (§ 4.4) **(Q3)** Is NODE2BITS scalable? (§ 4.5)

4.1 Experimental Setup

We ran our analysis on Mac OS platform with 2.5GHz Intel Core i7, 16GB RAM.

Data. We use five real-world heterogeneous networks from the Network Repository [28], as well as a real, proprietary user stitching dataset, ‘Company X’ web logs. The latter data form a temporal heterogeneous network consisting of web sessions of user devices and their IP addresses.

Table 2: Network statistics and properties for our six real-world datasets. ‘D’: directed; ‘W’: weighted; ‘H’: heterogeneous; ‘T’: temporal network.

Data	Nodes	Edges	$ T_V $	D	W	H	T
citeseer	4460	2892	2		✓	✓	
yahoo	100,058	1,057,050	2	✓	✓	✓	
bitcoin	3,783	24,186	1	✓			✓
digg	283,183	6,473,708	2			✓	✓
wiki	1,140,149	7,833,140	1	✓			✓
comp-X	5,500,802	5,291,270	2	✓	✓	✓	✓

High degree nodes representing anomalous behavior (*e.g.*, bots or public WiFi hotspots) are filtered out. Our framework is also capable of modeling domain-specific features, such as user-agent strings and geolocation [20], if this is available. Even without them, however, it achieves strong performance. We give the statistics of all the networks in Table 2, and additional details in Appendix C.

Task Setup. With the exception of § 4.3, we cast the user stitching task as a binary classification problem, where for each pair of nodes we aim to predict whether they correspond to the same entity (*i.e.*, we should stitch them). We use logistic regression with regularization strength 1.0 and stopping criterion 10^{-4} .

For the real user stitching data, we use the held-out, ground-truth information to evaluate our method. For the five real-world networks without known user

pairs, we introduce user replicas following a similar procedure to [2]: we sample 5% of the nodes with degrees larger than average, introduce a replica u' for each sampled node u , and distribute the original edges between u and u' . Specifically, each edge remains connected to u with probability p_1 , otherwise it connects to the replica node u' . Additionally, each edge that is incident to u has probability p_2 to also connect to u' . Unless otherwise specified, we use $p_1 = 0.6$ and $p_2 = 0.3$.

Given the positive replica pairs, we sample an equal number of negative pairs uniformly at random and include these in the training and testing sets. Comp-X, the dataset with ground-truth replicas, also has pre-defined approximately 50/50 training-testing splits that we use. Afterwards, embeddings are derived for each node pair by concatenation: $[\mathbf{z}(u), \mathbf{z}(u')]$. Using these node pair embeddings, we learn a logistic regression (LR) model and use it to predict the node pairs that should be stitched in the held-out test set. These are the nodes that correspond to the same entity. We measure the predictive performance of all the methods in terms of AUC, accuracy and $F1$ score.

Baselines. We compare to various methods that target different network types:

- **Homogeneous networks:** *Static*– (1) Spectral embeddings or SE [33], (2) LINE [32], (3) DeepWalk or DW [25], (4) node2vec or n2vec [14], (5) struc2vec or s2vec [27], and (6) DNGR [4]. *Temporal*– (7) CTDNE [23].
- **Heterogeneous networks:** (8) Common neighbors (CN) [2], (9) metapath2vec or m2vec [10], and (10) AspEm [31].

The baselines are configured to achieve the best performance, for $K = 128$ -dimensional embeddings, according to the respective papers. For reproducibility, we describe the settings in Appendix D.

NODE2BITS Setup & Variants. Similar to the baselines, NODE2BITS performs $R = 10$ walks per edge, with length up to $L = 20$, and we set the max temporal distance $\text{MAX} = 3$. We justify these decisions in Appendix E.2. On the largest dataset, Comp-X, we use a maximum walk length $L = 5$ and temporal distance $\text{MAX} = 2$. While various node attributes can be given as input to NODE2BITS, for consistency we derive and use the total, in-/out-degree of each node in \mathbf{F} .

We experiment with different variants of NODE2BITS (or N2B for short): (1) **NODE2BITS-0** applies to static networks; (2) **NODE2BITS-SH** uses the short-term tactic in the random walks (§ 3.1); (3) **NODE2BITS-LN** uses the long-term tactic; and (4) **NODE2BITS-U** targets unsupervised user stitching, which most baselines cannot handle (except for CN). To explore our method’s performance in unsupervised settings (§ 4.3), we directly ‘cluster’ the LSH-based, binary node representations \mathbf{z}_u generated by NODE2BITS-0. The idea is that nodes that hash to the same ‘bucket’ likely map to the same entity and should be stitched. To map entities to buckets we use the banding technique [26]: per band—one per representation $\mathbf{z}^{\Delta t}$ at temporal distance Δt —we apply AND-construction on the output of bit sampling, and then OR-construction across the bands.

4.2 Accuracy in Supervised User Stitching

We start by evaluating the predictive performance of NODE2BITS for supervised user stitching on both static and temporal networks.

Table 3: Entity resolution results for *static* networks. Our method outperforms all the baselines. — OOT = Out Of Time (6h); OOM = Out Of Memory (16GB). The asterisk * denotes statistically significant improvement over the best baseline at $p < 0.05$ in a two-sided t-test.

	Metric	CN	SE	LINE	DW	n2vec	s2vec	DNGR	m2vec	AspEm	N2B-0
citeseer	AUC	0.9141	0.4846	0.5481	0.5614	0.6188	0.9344	0.5015	0.5546	0.5049	0.9480*
	ACC	0.9141	0.5045	0.5372	0.5579	0.6211	0.8936	0.4688	0.5357	0.5223	0.9196*
	F1	0.9137	0.5028	0.5371	0.5547	0.6159	0.8926	0.4682	0.5348	0.5222	0.9192*
yahoo	AUC	0.6851	0.5378	0.8050	0.7640	0.7636			0.8233	0.4938	0.8088
	ACC	0.6851	0.4760	0.7771	0.7117	0.7233	OOT	OOM	0.7827	0.5018	0.8010
	F1	0.6505	0.4375	0.7764	0.7117	0.7231			0.7823	0.5018	0.7987

Static Networks. Here we evaluate the effectiveness of multi-dimensional feature contexts. Since static networks lack temporal information, NODE2BITS performs random walks similarly to existing works to collect nodes in structural contexts. The main difference lies in representing diverse feature histograms. We run NODE2BITS against both homogeneous and heterogeneous baselines as shown in Table 3, and observe that it performs the best in most evaluation metrics on both graphs. NODE2BITS outperforms existing random-walk based methods as expected: node IDs in the contexts is distorted by the replicas generated, thus feature-based methods should prevail. This is also validated by the results for struc2vec, which captures the equivalency of structural feature sequences in embeddings. metapath2vec and LINE achieve promising result on yahoo but not on citeseer, as the latter is an undirected bipartite graph, node distributions of the 2-order contexts explored by LINE are highly correlated and indistinguishable for stitching. On the contrary, CN (common-neighbors) yields promising result on citeseer but not yahoo. This is likely due to the graph structure, which we explain in more detail in Sec.4.3. We encountered out-of-memory errors for DNGR due to the algorithmic complexity and out-of-time-limit for struc2vec.

CONCLUSION 1. *On static graphs, NODE2BITS achieves comparable performance in AUC, and slightly better F1 score with 0.60% – 2.10% improvement over baselines in the stitching task.*

Temporal Networks. Table 4 depicts the stitching performance of NODE2BITS using both the short- and long-term tactics against the same set of baselines used in static graphs as well as CTDNE, an embedding framework designed for temporal graphs. We exclude metapath2vec, as metapaths are not meaningful in homogeneous networks, and the method ran out of time for the heterogeneous networks. We observe that NODE2BITS-SH outperforms NODE2BITS-LN in most cases, which is reasonable because NODE2BITS-LN derives shorter contexts constrained by temporal-order. We also justify the effectiveness of temporal random walk by comparing it with both NODE2BITS-0 and static baselines where we only make use of the graph structures without specifying edge timestamps. We observe that NODE2BITS-0 is the best-performing method for the digg dataset and Comp-X over the temporal variants of NODE2BITS. The reason behind this is that there is a tradeoff in constraining temporal walks to respect time: we more accurately model realistic sequences of events at the cost of restricting

Table 4: Entity resolution results for *temporal* networks: strong performance for NODE2BITS variants. — OOT = Out Of Time (6h); OOM = Out Of Memory (16GB); * denotes statistically significant improvement over the best baseline at $p < 0.05$ in a two-sided t-test.

	Metric	CN	SE	LINE	DW	n2vec	s2vec	DNGR	AspEm	CTDNE	N2B-0	N2B-SH	N2B-LN
bitcoin	AUC	0.7474	0.5828	0.6071	0.6306	0.6462	0.8025	0.5909	0.5344	0.6987	0.7584	0.7609	0.7380
	ACC	0.7174	0.5842	0.5842	0.6158	0.6158	0.7263	0.5526	0.5316	0.6000	0.7211	0.7268	0.6737
	F1	0.7001	0.5728	0.5828	0.6158	0.6157	0.7263	0.5525	0.5315	0.5964	0.7209	0.7271	0.6735
digg	AUC	0.6217	0.5171	0.7878	0.7398	0.7445			0.5105	0.6967	0.8185*	0.7611	0.7587
	ACC	0.6217	0.5152	0.7694	0.6971	0.7013	OOT	OOM	0.5088	0.5915	0.7982*	0.7418	0.7444
	F1	0.5585	0.3770	0.7683	0.6960	0.7003			0.5088	0.5884	0.7958*	0.7411	0.7433
wiki	AUC	0.6997		0.7854					0.5374	0.7707	0.8230	0.8259*	0.8214
	ACC	0.6997	OOT	0.7132	OOM	OOM	OOT	OOM	0.5141	0.6488	0.7145	0.7510*	0.7103
	F1	0.6699		0.7129					0.5141	0.6398	0.7088	0.7476*	0.7067
comp-X	AUC	0.5970		0.5000					0.5213		0.8095*	0.7496	0.7525
	ACC	0.5970	OOM	0.6757	OOM	OOM	OOT	OOM	0.5103	OOM	0.8414*	0.7959	0.7975
	F1	0.5189		0.4032					0.5103		0.8154*	0.7581	0.7606

the possible context. On these particular temporal graphs, walks may already be limited in length by the bipartite structure, so the latter cost becomes more appreciable. Nevertheless, both static and dynamic versions of NODE2BITS almost always outperform other baselines. In particular, across all datasets, NODE2BITS-SH still outperforms the temporal baseline, CTDNE in all cases, which further demonstrates the effectiveness of multi-feature aggregation.

NODE2BITS variants outperform the static methods in nearly all cases except the bitcoin dataset where NODE2BITS-SH achieves lower AUC than struc2vec but higher ACC and F1-score. This is because NODE2BITS loses some information when representing the node contexts as binary vectors comparing with real-value representation. However, we consider this loss mild as NODE2BITS still outperforms all the other static baselines. In addition, struc2vec ran out of time on the larger datasets while NODE2BITS achieves promising performance efficiently with 3.90% – 5.16% improvement in AUC and 3.58% – 4.87% improvement in F1 score than the best baseline method. At the same time, our approach uses much less information than the static methods, since the length of the temporal walks are typically shorter than random walks that do not have to respect time.

CONCLUSION 2. *Dynamic and static variants of NODE2BITS outperform the other baselines by up to 5.2% in AUC and 4.9% in F1 score. Between the two dynamic variants, the short-term tactic performs better than the long-term one.*

Restricting the Output Space Requirements. To evaluate the performance of stitching with explicit storage requirement, we hash the real-value embeddings given by baselines into binary and achieve output storage to be consistent with NODE2BITS. We observe that NODE2BITS still achieves the best performance (refer to Table 6 in Appendix E.1 for more details).

4.3 Accuracy in Unsupervised User Stitching

As mentioned in § 4.1, NODE2BITS can naturally perform unsupervised user stitching by leveraging the generated node representations as hashcodes. Only

Table 5: Unsupervised stitching performance between CN and NODE2BITS

Metric	citeseer		yahoo		bitcoin		digg		wiki	
	CN	N2B-U	CN	N2B-U	CN	N2B-U	CN	N2B-U	CN	N2B-U
ACC	0.9141	0.8661	0.6851	0.7553	0.7474	0.7684	0.6217	0.7157	0.6997	0.7350
F1	0.9137	0.8660	0.6505	0.7518	0.7301	0.7663	0.5585	0.7074	0.6699	0.7349

nodes mapped to the same ‘buckets’ are candidates for stitching together. This process allows us to stitch entities without involving quadratic comparisons between all pairs of nodes in the graph. Similarly, CN outputs a set of nodes sharing a certain amount of neighbors as the candidates to be stitched together. We evaluate the quality of hashing given by NODE2BITS-U against CN, and make use of the candidates to predict the testing set of node pairs given by following the same setup in § 4.2 in an unsupervised scheme.

Based on the results in Table 5, we observe that NODE2BITS-U outperforms CN on every dataset other than citeseer. The reason is that in this “author contributes to paper” dataset, author references appearing in the same set of papers have high probability to correspond to the same researcher in reality. Therefore the assumption made by CN suits well this scenario, whereas NODE2BITS hashes nodes with similar features in the context instead of those with similar neighbor identities (IDs). For datasets with less strict cross-type relationship, NODE2BITS achieves 2.81% – 15.12% improvement in accuracy ACC and 4.96% – 26.66% improvement in $F1$ score (including digg, another bipartite graph with inner connected components of the same node types).

CONCLUSION 3. *The unsupervised variant of NODE2BITS, NODE2BITS-U, outperforms CN on most graphs.*

4.4 Output Storage Efficiency

Next we evaluate space efficiency of our proposed method over baselines that output node embeddings. Instead of real-value matrices, the binary hashcodes generated by NODE2BITS can be stored in the sparse format so presumably it should take trivial storage. We visualize the storage requirements in Figure 3 and provide detailed storage usage in Table 7 in Appendix E.3.

CONCLUSION 4. *Compared to the other methods, NODE2BITS uses between $63\times$ and $339\times$ less space (while always achieving comparable or better stitching performance as shown in § 4.2).*

4.5 Scalability

To evaluate the scalability, we report the runtime of applying NODE2BITS to obtain node representations for the datasets shown in Table 2 versus their numbers of edges. We note that NODE2BITS-SH runs only on temporal networks, i.e., a subset of the datasets. We also visualize the runtime of node2vec as reference, as it is designed for large graphs and is implemented in the same language (Python).

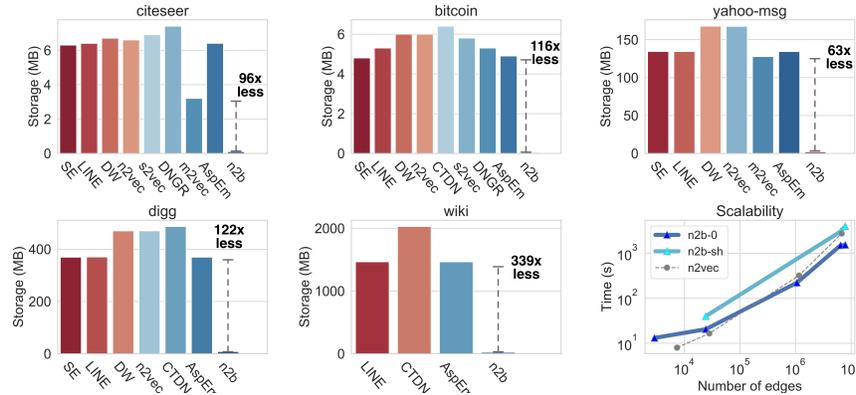


Fig. 3: First 5 plots: output storage in MB for all the methods that completed successfully in five datasets. NODE2BITS is also shown to be scalable for large graphs.

Based on the last subplot in Figure 3, we observe that NODE2BITS scales similarly as node2vec with less runtime space as node2vec ran out of memory on the largest dataset (wiki). As shown in Appendix B, the worst-case time complexity is linear in the edges. We give the exact runtimes in Table 8 in the Appendix E.3.

5 Related Work

Entity Resolution (the general problem under which user stitching falls) has been widely studied and applied in different domains such as databases and information retrieval [9,13]. Traditional methods that are based on distances can be broadly categorized into (1) pairwise-ER [7], which independently decide which pairs are same entity based on a distance threshold, and (2) clustering [8], which links nodes in the same cluster. However, these methods are computationally expensive and do not scale to large datasets. Other techniques range from supervised classification [30] to probabilistic soft logic [20] or fingerprinting [12] using side information (*e.g.*, user-agent strings, other web browser features, geo-location). These methods tend to be problem- or even data-specific. On the contrary, our method is general by modeling the data with a heterogeneous, dynamic network that uses both node features (optional) *and* graph structure.

Node embeddings aim to preserve a notion of node similarity into low-dimensional vector space. Most general methods [14,25,32] and the state-of-the-art for heterogeneous or dynamic networks [10,23], define node similarity based on co-occurrence or proximity in the original network (Appendix F). However, in the user stitching problem, it is possible that corresponding entities may not connect to the same entities, resulting in lower proximity-based similarity. Embeddings preserving *structural identity* [27,11,16,19,18,22,1] overcome this drawback. NODE2BITS additionally handles various graph settings (heterogeneous, dynamic) at greater space efficiency thanks to hashing.

Locality sensitivity hashing (LSH) was first introduced as a randomized hashing framework for efficient approximate nearest neighbor search in high dimensional space [17]. It specifies a family of hash functions, \mathcal{H} , that maps

similar items to the same bucket identified through hash codes with higher probability than dissimilar items [26]. LSH families for different distances have been widely studied, such as SimHash for cosine distance [5], min-hash for Jaccard similarity [3], and more (Appendix F). In our work, we leverage LSH to construct similarity-preserving and space-efficient node representations for user stitching.

6 Conclusion

We have proposed a hash-based network representation learning framework for identity stitching called NODE2BITS. It is both time- and attribute-aware, while also deriving space-efficient sparse binary embeddings of nodes in large temporal heterogeneous networks. NODE2BITS uses the notion of feature-based temporal walks to capture the temporal and feature-based information in the data. Feature-based temporal walks are a generalization of walks that obey time while also incorporating features (as opposed to node IDs). Using these walks, NODE2BITS generates contexts/sequences of temporally valid feature values. Experiments on real-world networks demonstrate the utility of NODE2BITS as it outputs space-efficient embeddings that use orders of magnitude less space compared to the baseline methods while achieving better performance in user stitching. An important practical consideration in the application of our work to user stitching is the balance of greater personalization with user privacy.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. IIS 1845491, Army Young Investigator Award No. W911NF1810397, an Adobe Digital Experience and an Amazon research faculty award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or other funding parties.

References

1. Ahmed, N.K., Rossi, R.A., Zhou, R., Lee, J.B., Kong, X., Willke, T.L., Eldardiry, H.: Learning role-based graph embeddings. In: StarAI workshop at IJCAI (2018)
2. Bhattacharya, I., Getoor, L.: Collective entity resolution in relational data. *TKDD* **1**(1), 1–36 (2007)
3. Broder, A.Z., Glassman, S.C., Manasse, M.S., Zweig, G.: Syntactic clustering of the web. *Computer Networks and ISDN Systems* **29**(8), 1157–1166 (1997)
4. Cao, S., Lu, W., Xu, Q.: Deep neural networks for learning graph representations. In: *AAAI*. pp. 1145–1152 (2016)
5. Charikar, M.S.: Similarity estimation techniques from rounding algorithms. In: *STOC*. pp. 380–388 (2002)
6. Christen, P.: *Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*. Springer (2012)
7. Cohen, W.W., Richman, J.: Learning to match and cluster large high-dimensional data sets for data integration. In: *KDD*. pp. 475–480 (2002)
8. Dasgupta, A., Gurevich, M., Zhang, L., Tseng, B., Thomas, A.O.: Overcoming browser cookie churn with clustering. In: *WSDM*. pp. 83–92 (2012)

9. Dong, X.L., Naumann, F.: Data fusion: resolving data conflicts for integration. *VLDB* **2**(2), 1654–1655 (2009)
10. Dong, Y., Chawla, N.V., Swami, A.: metapath2vec: Scalable representation learning for heterogeneous networks. In: *KDD*. pp. 135–144 (2017)
11. Donnat, C., Zitnik, M., Hallac, D., Leskovec, J.: Learning structural node embeddings via diffusion wavelets. In: *KDD*. pp. 1320–1329 (2018)
12. Eckersley, P.: How unique is your web browser? In: *International Symposium on Privacy Enhancing Technologies Symposium*. pp. 1–18. Springer (2010)
13. Getoor, L., Machanavajjhala, A.: Entity resolution for big data. In: *KDD*. pp. 1527–1527 (2013)
14. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: *KDD*. pp. 855–864 (2016)
15. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: *NeurIPS*. pp. 1024–1034 (2017)
16. Heimann, M., Shen, H., Safavi, T., Koutra, D.: Regal: Representation learning-based graph alignment. In: *CIKM*. pp. 117–126 (2018)
17. Indyk, P., Motwani, R.: Approximate nearest neighbors: towards removing the curse of dimensionality. In: *STOC*. pp. 604–613 (1998)
18. Jin*, D., Heimann*, M., Safavi, T., Wang, M., Lee, W., Snider, L., Koutra, D.: Smart roles: Inferring professional roles in email networks. In: *KDD* (2019)
19. Jin, D., Rossi, R.A., Koh, E., Kim, S., Rao, A., Koutra, D.: Latent network summarization: Bridging network embedding and summarization. In: *KDD* (2019)
20. Kim, S., Kini, N., Pujara, J., Koh, E., Getoor, L.: Probabilistic visitor stitching on cross-device web logs. In: *WWW*. pp. 1581–1589 (2017)
21. Kolb, L., Thor, A., Rahm, E.: Dedoop: Efficient deduplication with hadoop. *VLDB* **5**(12), 1878–1881 (Aug 2012)
22. Liu, Y., Zhu, L., Szekely, P., Galstyan, A., Koutra, D.: Coupled clustering of time-series and networks. In: *SDM* (2019)
23. Nguyen, G.H., Lee, J.B., Rossi, R.A., Ahmed, N.K., Koh, E., Kim, S.: Continuous-time dynamic network embeddings. In: *WWW BigNet* (2018)
24. Papadakis, G., Svirsky, J., Gal, A., Palpanas, T.: Comparative analysis of approximate blocking techniques for entity resolution. *VLDB* **9**(9), 684–695 (2016)
25. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: *KDD* (2014)
26. Rajaraman, A., Leskovec, J., Ullman, J.D.: *Mining Massive Datasets* (2014)
27. Ribeiro, L.F., Saverese, P.H., Figueiredo, D.R.: struc2vec: Learning node representations from structural identity. In: *KDD*. pp. 385–394 (2017)
28. Rossi, R.A., Ahmed, N.K.: The network data repository with interactive graph analytics and visualization. In: *AAAI* (2015), <http://networkrepository.com>
29. Rossi, R.A., Zhou, R., Ahmed, N.K.: Deep inductive network representation learning. In: *WWW*. pp. 953–960 (2018)
30. Saha Roy, R., Sinha, R., Chhaya, N., Saini, S.: Probabilistic deduplication of anonymous web traffic. In: *WWW*. pp. 103–104 (2015)
31. Shi, Y., Gui, H., Zhu, Q., Kaplan, L., Han, J.: Aspem: Embedding learning by aspects in heterogeneous information networks. In: *SDM*. pp. 144–152. SIAM (2018)
32. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: LINE: large-scale information network embedding. In: *WWW* (2015)
33. Von Luxburg, U.: A tutorial on spectral clustering. *Statistics and computing* **17**(4), 395–416 (2007)
34. Wang, Q., Wang, S., Gong, M., Wu, Y.: Feature hashing for network representation learning. In: *IJCAI*. pp. 2812–2818 (2018)

Appendix

A Detailed Algorithm

In § 3 we gave the overview of our proposed method, NODE2BITS. For reproducibility, here we also provide its more detailed pseudocode.

Algorithm 2 NODE2BITS framework in detail

Require:

(un)directed heterogeneous graph $G(V, E, \psi, \xi)$,
 number of random walks R per edge, max walk length L ,
 max temporal distance MAX ,
 embedding dimensionality $K^{\Delta t}$ at temporal dist. Δt (output dim. $K = \sum_{\Delta t=1}^{MAX} K^{\Delta t}$)

- 1: Construct $N \times |\mathcal{F}|$ feature matrix \mathbf{F} ▷ Matrix with node attributes and derived features
- 2: $\mathcal{C}_u^{\Delta t} \leftarrow \emptyset, \forall u \in V, \Delta t \in [1, \dots, MAX]$ ▷ Context of u : nodes at temporal distance Δt
- 3: **for** edge e and walk $= 1, \dots, R$ **do** ▷ Perform R feature-based temporal rand. walks per edge
- 4: $\langle w_L \rangle_{L \in \mathbb{N}} \leftarrow$ up to L -step temporal walk starting from edge e based on tactic ▷ Dfn. 3
- 5: $\mathcal{C}_u^{\Delta t} \leftarrow$ update the context of nodes $u \in \langle w_L \rangle_{L \in \mathbb{N}}$ and all temporal distances Δt
- 6: **for** $j = 1, \dots, |\mathcal{F}|$ **do** ▷ Iterate over all the features in matrix \mathbf{F}
- 7: ▷ Generate the feature-based context by replacing $v \in \mathcal{C}_u^{\Delta t}$ with the $f_{v,j}$.
- 8: ▷ Equivalent to context generation after a feature-based temporal walk (Dfn. 4).
- 9: $\mathcal{C}_u^{\Delta t}[f(j)] \leftarrow$ update the feature-based context of $u \in \langle w_L \rangle_{L \in \mathbb{N}}$
- 10:
- 11: **for** $\Delta t = 1, \dots, MAX$ **do**
- 12: Generate $K^{\Delta t}$ random hyperplanes
- 13: **for each** node $u \in V$ **do** ▷ For each node, summarize its context with
- 14: **for each** node $v \in \mathcal{C}_u^{\Delta t}$ **do** ▷ a histogram per feature and node type.
- 15: $\mathbf{h}(\mathcal{C}_u^{\Delta t}) = \text{concatenate}[\mathbf{h}(\mathcal{C}_u^{\Delta t} | f_{(1)}, p_1), \dots, \mathbf{h}(\mathcal{C}_u^{\Delta t} | f_{(|\mathcal{F}|)}, p_{|\mathcal{T}_V|})]$ ▷ (Eq. (4))
- 16: ▷ Obtain a sparse, binary hashcode per node based on (modified) SimHash.
- 17: $z_u^{\Delta t} \leftarrow$ SimHash of node histogram $\mathbf{h}(\mathcal{C}_u^{\Delta t})$ at distance Δt
- 18: $\mathbf{Z}^{\Delta t} \leftarrow N \times K^{\Delta t}$ matrix with each node's SimHash code per row
- 19: **return** Sparse node representation $\mathbf{Z} = \text{concatenate}[\mathbf{Z}^1, \dots, \mathbf{Z}^{MAX}]$

B Complexity Analysis

Time Complexity. The runtime complexity of NODE2BITS includes deriving (1) the set of R temporal random walks of length up to L , which is $\mathcal{O}(MRL)$ in the worst case; (2) the feature values of nodes in the walks from step (1); and (3) hashing the feature values of nodes in the context through random projection, which is $\mathcal{O}(NK)$. Thus, the total runtime complexity is $\mathcal{O}(MRL + NK)$, which is linear to the number of edges when $M \gg N$ as K is relatively small (**R3**).

Runtime Space Complexity. The space required in the runtime consists three parts: (1) the set of temporal random walks (represented as vectors) per edge with complexity $\mathcal{O}(MRL)$, (2) the histograms of feature contexts $N|\mathcal{F}||\mathcal{T}_V|$, and (3) the set of randomly-generated hyperplanes NK . Therefore, the total runtime space complexity is $\mathcal{O}(MRL + N(|\mathcal{F}||\mathcal{T}_V| + K))$.

Output Space Complexity. The output space complexity of NODE2BITS is $\mathcal{O}(NK)$ -bit. The space required to store binary vectors is guaranteed to be $32\times$ less than vectors represented with real-value floats (4 bytes) with the same dimension. In practice, NODE2BITS requires even less storage if the binary vectors are represented in the sparse format (see Section 4.4 for empirical results).

C Data Description

Below we provide a more detailed description of the network datasets that we use in our experiments (Table 2).

- **citeseer:** CiteSeerX is an undirected, heterogeneous network that contains the bipartite relations between authors and papers they contributed.
- **yahoo:** Yahoo! Messenger Logs is a heterogeneous network capturing message exchanges between users at different locations (node attribute).
- **bitcoin:** soc-bitcoinA is a who-trusts-whom network on the Bitcoin Alpha platform. The directed edges indicate user ratings.
- **digg:** This heterogeneous network consists of users voting stories that they like and forming friendships with other users.
- **wiki:** wiki-talk is a temporal homogeneous network capturing Wikipedia users editing each other’s Talk page over time.
- **comp-X:** A temporal heterogeneous network is derived from a company’s web logs and consists of web sessions of users and their IPs. In the stitching task, we predict the web session IDs that correspond to the same user.

D Configuration of Baselines

As we mentioned in § 4.1, we configured all the baselines to achieve the best performance according to the respective papers. For all the baselines that are based on random walks (*i.e.*, node2vec, struc2vec, DeepWalk, metapath2vec, CTDNE), we set the number of walks to 20 and the maximum walk length to $L = 20$. For node2vec, we perform grid search over $p, q \in \{0.25, 0.50, 1, 2, 4\}$ as mentioned in [14] and report the best performance. For metapath2vec, we adopt the recommended meta-path “Type 1-Type 2-Type 1” (*e.g.*, type 1 = author; type 2 = publication). In DNGR, we set the random surfing probability $\alpha = 0.98$ and use a 3-layer neural network model where the hidden layer has 1024 nodes. We use 2nd-LINE to incorporate 2nd-order proximity in the graph. For all the embedding methods, we set the embedding dimension to $K = 128$. Unlike those, CN outputs clusters, each of which corresponds to one entity.

E Additional Empirical Analysis

E.1 Justification of hashing

In this experiment we hash the outputs given by baseline embedding methods using SimHash [5] and then perform stitching on two temporal graphs to study their performance under the constraint of storage comparable to NODE2BITS. Based on Table 6, we observe fluctuation in the stitching performance of baseline methods, for example, almost all baselines got degenerated scores in all metrics on the bitcoin dataset,

Table 6: Justification of hashing

	Metric	SC*	LINE*	DW*	n2vec*	CTDNE	s2vec*	AspEm	N2B-0	N2B-SH	N2B-LN
bitcoin	AUC	0.5160	0.5807	0.5904	0.6265	0.6652	0.7703	0.5212	0.7584	0.7754	0.7380
	ACC	0.5158	0.5421	0.5632	0.5895	0.6632	0.7105	0.5211	0.7211	0.7368	0.6737
	F1	0.3757	0.5415	0.5611	0.5893	0.6608	0.7087	0.3334	0.7209	0.7361	0.6735
digg	AUC	0.5001	0.7909	0.7607	0.7599	0.7203		0.5000	0.8185	0.7611	0.7587
	ACC	0.5001	0.7751	0.7039	0.7030	0.6357	OOT	0.5000	0.7982	0.7418	0.7444
	F1	0.3338	0.7746	0.7039	0.7030	0.6228		0.3332	0.7958	0.7411	0.7433

especially for struc2vec. On the other hand, however, node2vec, LINE and CTDNE got slight increased scores on yahoo dataset, which is likely due to the fact that the small real-values are amplified when hashed into binary for logistic regression binary classification. It is also possibly due to the graph structure. We leave further discussion in the future work, but nevertheless, NODE2BITS outperforms these baselines in all cases. This empirical experiment demonstrates that NODE2BITS effectively preserves context information in the binary hashcodes.

E.2 Sensitivity analysis

We also perform sensitivity analysis of the hyperparameters used in this work on the bitcoin dataset. Particularly, we perform grid analysis by varying (1) max temporal distances, (2) the number of temporal walks per edge and (3) the length of walks. The results are given in Figure 4. Figure 4a indicates that when $MAX = 3$, NODE2BITS achieves the best performance. This implies that although it is potentially beneficial to incorporate nodes in temporally distant contexts, it will also incorporate information that is less relevant. Therefore, we set $MAX = 3$ by default for the experiments in this work. Figure 4b-4c imply that the performance of NODE2BITS is not significantly affected by the number of walks performed or the length of these temporal walks. This is reasonable because NODE2BITS leverages these temporal walks to collect node features into the context and normalizes their occurrences in the histograms. Thus, adding more nodes in the ordered temporal contexts does not provide extra useful information. We empirically set the number of walks per edge to be 10 and the lengths to be 20 in the experiments of this work.

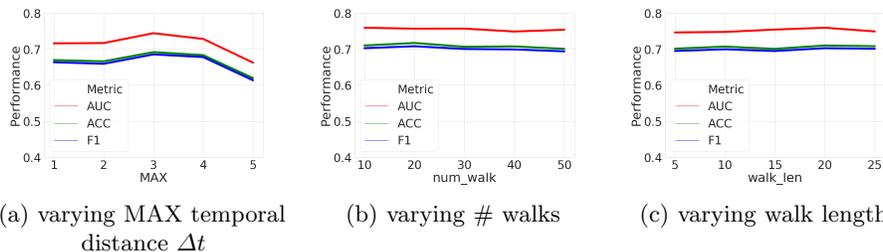


Fig. 4: Sensitivity Analysis on bitcoin dataset. NODE2BITS achieves highest scores in AUC, ACC and F -1 score when $MAX = 3$. Increasing the numbers of walks or increasing their lengths do not significantly affect the performance of NODE2BITS.

E.3 Output storage and runtime in detail

We report detailed output storage in Table 7 and the time elapsed when running all methods in Table 8. It can be seen that the node-wise sparse binary vectors generated by NODE2BITS take trivial amount of storage compared to the other methods, while its runtime is comparable to node2vec. NODE2BITS finished running on all datasets while most baselines fail to finish within the time limit on the large datasets, digg and wiki.

Table 7: Space required to store the output in MB. NODE2BITS requires $63\times$ - $339\times$ less space than other embedding methods. ‘-’ indicates that the method does not apply to that dataset, or encounters errors such as out-of-memory or out-of-time.

datasets	SC	LINE	DW	n2v	CTDNE	s2vec	DNGR	m2vec	AspEm	N2B
citeseer	6.3	6.4	6.7	6.6	-	6.9	7.4	3.2	6.4	0.033
yahoo	134.4	134.3	167.8	167.6	-	-	-	127.7	134.3	2.1
bitcoin	4.8	5.3	6.0	6.0	6.4	5.8	5.3	-	4.9	0.041
digg	369.1	370.0	469.8	469.8	486.6	-	-	-	369.3	2.9
wiki	-	1430	-	-	1980	-	-	-	1430	4.2

Table 8: Comparison between NODE2BITS and baselines in terms of runtime (in seconds). Note the runtime of dynamic NODE2BITS (short-term) for the *temporal* networks is shown in parentheses.

	citeseer	yahoo	bitcoin	digg	wiki
SC	23.72	766.42	4.80	8091.09	1
LINE*	144.94	223.87	134.48	227.28	415.00
DW*	8.90	209.72	16.99	2115.86	-
n2v*	7.99	222.14	15.91	2751.91	-
CTDNE	-	-	13.25	2227.66	4217.19
s2vec*	325.38	-	897.2	-	-
DNGR	128.63	-	97.09	-	-
m2vec	125.98	-	-	-	-
AspEm	0.62	4.70	0.71	15.318	386.24
CN	0.58	19.59	0.70	63.95	109.11
N2B	13.15	221.84	20.52 (39.97)	1507.95 (3062.13)	1537.24 (3997.85)

F Additional Related Work

In this section we provide additional related work, complementing our discussion in § 5.

Node Embeddings. Here we give some more details about proximity-based methods, which we employ in our experiments. DeepWalk [25] and node2vec [14] leverage vanilla and 2-order random walk, respectively, to explore the identities of the neighborhood; LINE [32] can be seen as a special case of DeepWalk by setting the context to be 1 [56]; metapath2vec [10] relies on predefined meta-schema to perform random walk in heterogeneous networks. In the field of temporal network embedding, most approaches [65, 43] approximate the dynamic network as discrete static snapshots overtime, which does not apply to user stitching tasks as sessions corresponding to the same user could occur in multiple timespans. CTDNE [23] first explores temporal proximity by learning temporally valid embeddings based on a corpus of temporal random walks. Another related field is hashing-based embedding, for example, node2hash [34] proposes to hash

the pairwise node proximity derived from random walk into low-dimensional hashcode as the embeddings. Due to the quadratic complexity in computing the pairwise proximity between nodes, node2hash does not apply to large-scale networks. One limitation of these methods is that training a skip-gram architecture on the entire corpus sampled by random walks can be memory-intensive. A further limitation of these approaches, as well as existing deep architectures [63, 40] is that for nodes to have similar embeddings, they must be in close proximity (e.g. neighbors or nodes with several common neighbors) in the network. This is not necessarily the case for user stitching, where corresponding entities may exhibit similar behavior (resulting in similar local topologies) but not connect to the same entities.

Compared with proximity-based methods, embedding works exploring structural equivalency or similarity [1, 27, 58, 11, 16, 19, 18, 22] are more suitable to handle user stitching. Representative examples include the following: struc2vec [27], xNetMf [16], and EMBER [18] define similarity in terms of degree sequences in node-centric subgraphs; DeepGL [29] learns deep relational functions applied to degree, triangle counts and other graph invariants in an inductive scheme. Role2vec [1] proposes a framework that inductively learns structural similarity by introducing attributed random walk atop relational operators, while MultiLENS [19] summarizes node embeddings obtained by recursive application of relational operators. CCTN [22] embeds and clusters nodes in a network that are not only well-connected but also share similar behavioral patterns (e.g., similar patterns in the degree or other structural properties over time).

Locality sensitivity hashing (LSH). More recently, LSH functions that are robust to distortion [35]; require less storage of the hash codes [49, 64]; generate codewords with balanced amounts of items [47] or compute hash functions efficiently [48, 50, 44, 61] have attracted much attention. LSH has been used in a variety of data mining applications, including network alignment [41], network inference [60], anomaly detection [53], and more. In addition, there are also works devoted to learning to hash [35] where the main idea is to learn hash codes through an optimization objective function, or intelligently probe multiple adjacent code words that are likely to contain query results in a hash table for similarity search [52]. But these methods do not apply to large-scale graphs directly.

Additional References

35. Aghazadeh, A., Lan, A., Shrivastava, A., Baraniuk, R.: Rhash: robust hashing via l-norm distortion. In: IJCAI. pp. 1386–1394 (2017)
36. Ahmed, N.K., Rossi, R.A., Zhou, R., Lee, J.B., Kong, X., Willke, T.L., Eldardiry, H.: Learning role-based graph embeddings. In: StarAI workshop at IJCAI (2018)
37. Dong, Y., Chawla, N.V., Swami, A.: metapath2vec: Scalable representation learning for heterogeneous networks. In: KDD. pp. 135–144 (2017)
38. Donnat, C., Zitnik, M., Hallac, D., Leskovec, J.: Learning structural node embeddings via diffusion wavelets. In: KDD. pp. 1320–1329 (2018)
39. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: KDD. pp. 855–864 (2016)
40. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: NeurIPS, pp. 1024–1034 (2017)
41. Heimann, M., Lee, W., Pan, S., Chen, K.Y., Koutra, D.: Hashalign: Hash-based alignment of multiple graphs. In: Pacific-Asia Conference on Knowledge Discovery and Data Mining. pp. 726–739. Springer (2018)

42. Heimann, M., Shen, H., Safavi, T., Koutra, D.: Regal: Representation learning-based graph alignment. In: CIKM. pp. 117–126 (2018)
43. Hisano, R.: Semi-supervised graph embedding approach to dynamic link prediction. In: International Workshop on Complex Networks. pp. 109–121. Springer (2018)
44. Ji, J., Li, J., Yan, S., Tian, Q., Zhang, B.: Min-max hash for jaccard similarity. In: ICDM. pp. 301–309 (2013)
45. Jin*, D., Heimann*, M., Safavi, T., Wang, M., Lee, W., Snider, L., Koutra, D.: Smart roles: Inferring professional roles in email networks. In: KDD (2019)
46. Jin, D., Rossi, R.A., Koh, E., Kim, S., Rao, A., Koutra, D.: Latent network summarization: Bridging network embedding and summarization. In: KDD (2019)
47. Kang, B., Jung, K.: Robust and efficient locality sensitive hashing for nearest neighbor search in large data sets. In: NeurIPS BigLearn Workshop. pp. 1–8 (2012)
48. Li, P., Hastie, T.J., Church, K.W.: Very sparse random projections. In: KDD. pp. 287–296 (2006)
49. Li, P., König, A.C.: Theory and applications of b-bit minwise hashing. *Comm. of the ACM* **54**(8), 101–109 (2011)
50. Li, P., Owen, A., Zhang, C.H.: One permutation hashing. In: NeurIPS. pp. 3113–3121 (2012)
51. Liu, Y., Zhu, L., Szekely, P., Galstyan, A., Koutra, D.: Coupled clustering of time-series and networks. In: SDM (2019)
52. Lv, Q., Josephson, W., Wang, Z., Charikar, M., Li, K.: Multi-probe lsh: efficient indexing for high-dimensional similarity search. In: VLDB. pp. 950–961 (2007)
53. Manzoor, E., Milajerdi, S.M., Akoglu, L.: Fast memory-efficient anomaly detection in streaming heterogeneous graphs. In: KDD. pp. 1035–1044 (2016)
54. Nguyen, G.H., Lee, J.B., Rossi, R.A., Ahmed, N.K., Koh, E., Kim, S.: Continuous-time dynamic network embeddings. In: WWW BigNet (2018)
55. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: KDD (2014)
56. Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., Tang, J.: Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In: WSDM. pp. 459–467 (2018)
57. Ribeiro, L.F., Saverese, P.H., Figueiredo, D.R.: struc2vec: Learning node representations from structural identity. In: KDD. pp. 385–394 (2017)
58. Rossi, R.A., Ahmed, N.K.: Role discovery in networks. *TKDE* **27**(4), 1112–1131 (2015)
59. Rossi, R.A., Zhou, R., Ahmed, N.K.: Deep inductive network representation learning. In: WWW. pp. 953–960 (2018)
60. Safavi, T., Sripada, C., Koutra, D.: Scalable hashing-based network discovery. In: ICDM. pp. 405–414 (Nov 2017)
61. Shrivastava, A., Li, P.: Densifying one permutation hashing via rotation for fast near neighbor search. In: ICML. pp. 557–565 (2014)
62. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: LINE: large-scale information network embedding. In: WWW (2015)
63. Wang, D., Cui, P., Zhu, W.: Structural deep network embedding. In: KDD. pp. 1225–1234 (2016)
64. Weiss, Y., Torralba, A., Fergus, R.: Spectral hashing. In: NeurIPS. pp. 1753–1760 (2009)
65. Zhu, L., Guo, D., Yin, J., Ver Steeg, G., Galstyan, A.: Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE TKDE* **28**(10), 2765–2777 (2016)