

From random-walks to graph-sprints: a low-latency node embedding framework on continuous-time dynamic graphs

This is a preprint, for the final version please refer to the ACM ICAIF 2023 conference proceedings.

Ahmad Naser Eddin*
Feedzai
Coimbra, Portugal
Departamento de Ciência de
Computadores, Faculdade de Ciências,
Universidade do Porto
Porto, Portugal
ahmad.eddin@feedzai.com

Hugo Ferreira
Feedzai
Coimbra, Portugal
hugo.ferreira@feedzai.com

Jacopo Bono*
Feedzai
Coimbra, Portugal
jacopo.bono@feedzai.com

João Ascensão
Feedzai
Coimbra, Portugal
joao.ascensao@feedzai.com

Pedro Bizarro
Feedzai
Coimbra, Portugal
pedro.bizarro@feedzai.com

David Aparício
Departamento de Ciência de
Computadores, Faculdade de Ciências,
Universidade do Porto
Porto, Portugal
daparicio@dcc.fc.up.pt

Pedro Ribeiro
Departamento de Ciência de
Computadores, Faculdade de Ciências,
Universidade do Porto
Porto, Portugal
pribeiro@dcc.fc.up.pt

ABSTRACT

Many real-world datasets have an underlying dynamic graph structure, where entities and their interactions evolve over time. Machine learning models should consider these dynamics in order to harness their full potential in downstream tasks. Previous approaches for graph representation learning have focused on either sampling k -hop neighborhoods, akin to breadth-first search, or random walks, akin to depth-first search. However, these methods are computationally expensive and unsuitable for real-time, low-latency inference on dynamic graphs. To overcome these limitations, we propose *graph-sprints* a general purpose feature extraction framework for continuous-time-dynamic-graphs (CTDGs) that has low latency and is competitive with state-of-the-art, higher latency models. To achieve this, a streaming, low latency approximation to the random-walk based features is proposed. In our framework, time-aware node embeddings summarizing multi-hop information are computed using only single-hop operations on the incoming edges. We evaluate our proposed approach on three open-source datasets and two in-house datasets, and compare with three state-of-the-art algorithms (TGN-attn, TGN-ID, Jodie). We demonstrate that our *graph-sprints* features, combined with a machine learning classifier, achieve competitive performance (outperforming all baselines for the node classification tasks in five datasets). Simultaneously, *graph-sprints* significantly reduce inference latencies, achieving close to an order of magnitude speed-up in our experimental setting.

CCS CONCEPTS

• **Computer systems organization** → *Real-time system architecture*; • **Information systems** → *Extraction, transformation and*

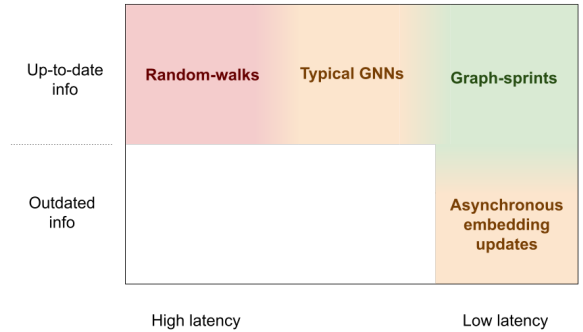


Figure 1: CTDGs Approaches Overview. Current approaches either compute embeddings using real-time information but sacrificing latency, or compute low-latency embeddings but use outdated information in the computation. Our proposed method, *graph-sprints*, computes embeddings in low latency using real-time information.

loading; • **Applied computing** → *Enterprise applications*; • **Computing methodologies** → *Temporal reasoning*.

KEYWORDS

Continuous-time dynamic graphs (CTDGs), Streaming graphs, Graph representation learning, Graph feature engineering, Graph Neural Networks

*These authors contributed equally to this work




Figure 2: AUC vs. Runtime trade-off. Our proposed methods based on *graph-sprints* (GS or GS+Raw), allow for low latency inference while outperforming state-of-the-art methods in terms of AUC on node classification tasks. X-axis represents time in seconds to process 200 batches of size 200, Y-axis represents test AUC. Error bars denote the standard deviation over 10 random seeds.

1 INTRODUCTION

Many real-world datasets have an underlying graph structure. In other words, they are characterized not only by their individual data points but also by the *relationships* between them. Moreover, they are typically dynamic in nature, meaning that the entities and their interactions change over time. Examples of such systems are social networks, financial datasets, and biological systems [3, 20, 29]. Dealing with dynamic graphs is more challenging compared to static graphs, especially if the graphs evolve in continuous time (also known as continuous-time dynamic graphs or CTDGs). The majority of machine learning models on graph datasets are based on graph neural networks (GNNs), achieving state-of-the-art performance [25]. A few deep neural network architectures emerged to deal with CTDGs in the past years [7, 8, 11, 12, 16, 19, 26]. One drawback of these approaches is that one either needs to sample k-hop neighborhoods to compute the embeddings (e.g. [16]) or perform random-walks (e.g. [11]). Both cases are computationally costly, resulting in high inference latencies. One solution is to decouple the inference from the expensive graph computations, like in APAN [22]. Performing the graph aggregations asynchronously results in inference using outdated information. However, in large data and high-frequency use-cases, such as our own focus on detecting fraud in financial transactions, there is a crucial need for low-latency solutions capable of processing up to 1000 transactions per second. Moreover, it is desirable for these solutions to leverage the latest information available in order to enhance detection capabilities.

In this work, we propose a graph feature extraction framework, which computes node embeddings in the form of features that characterize a node’s neighborhood in dynamic graphs. Importantly, our proposed framework is designed for low-latency settings while still using the most up-to-date information during the embedding calculations. Since we derive our framework starting from random-walk based methods, we term our method *graph-sprints*. The learned *graph-sprints* features can then be used in any downstream system,

for example in a machine learning model or a rule-based system. We show how the proposed *graph-sprints* features, combined with a neural network classifier, are faster to run while not sacrificing in predictive performance compared with the higher-latency GNNs (Figure 2).

Our **contributions** include: First, we propose *graph-sprints*, a graph feature extraction framework for CTDGs. Next, we introduce two strategies to minimize the memory requirements of the generated features. Lastly, we benchmark the node embedding quality against state-of-the-art GNNs using five datasets and two tasks: node classification and link prediction.

The remainder of the paper is organized as follows. We first discuss the proposed *graph-sprints* framework in Section 2.2. In Section 3, we evaluate our framework using three open-source datasets from different domains, and two in-house datasets from the money laundering domain. We discuss related work in Section 4 and in Section 5 we put forward our main conclusions.

2 METHODS

2.1 Random-walk based features

Before describing our *graph-sprints* framework, we briefly summarize a random-walk based feature extraction framework. We will then derive efficient *graph-sprints* computations in the next section. The random-walk based feature extraction framework requires the following steps, to generate a node embedding for a seed node.

- (1) **Select the seed node.** This selection depends on the use-case, and for CTDGs typically one considers entities involved in new activity, for instance if the change on the graph is adding a new edge between two nodes, then each of these two nodes could be a candidate for a seed node.
- (2) **Perform random-walks starting from the seed nodes.** During the random-walks, relevant data such as node or edge features of the traversed path are collected. The type of random-walks influences what neighborhood is summarized

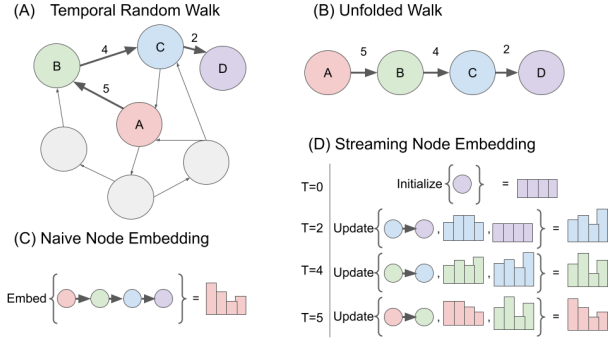


Figure 3: From random-walks to graph-sprints. Edges have a timestamp feature (numbers) representing the time that a relationship was created. (A) A temporal random-walk is traversed from the most recent interaction A-B towards older interactions. (B) The same random-walk can be seen as a time-series of edges. (C) Based on a full temporal random-walk, one can compute embeddings by aggregating encountered feature values (section 2.1). (D) One can compute similar embeddings in a streaming setting, from only the new edge and the existing embeddings of the involved nodes.

in the extracted features. Walks can be (un)directed, biased, and/or temporal.

- (3) **Summarize collected data.** The data collected over walks is aggregated into a fixed set of features, characterizing each seed node’s neighborhood. Examples of such aggregations are the average of encountered numerical node or edge features, the maximum of encountered out-degree, etc.

The computation of these features is costly, because multiple random-walks need to be generated for each seed node. For CTDGs, one would have to compute such features each time an edge arrives. This is infeasible for high-frequency use-cases such as fraud detection in financial transactions, where a decision about a transaction needs to be made in a few milliseconds. In the next section, we derive an efficient approximation to the above random-walk based features.

2.2 Graph-sprints: streaming graph features

In this section, we propose approximations to random-walk based features described in Section 2.1. Our aim in this section is to optimize the computation of such features by exploiting recurrence and abolishing the need to execute full random-walks (Figure 3).

2.2.1 Assumptions. For our approximations to be reliable, we make the following assumptions: the input graph is a CTDG with directed edges (we will relax this assumption later), edges have timestamps and the temporal walks respect time, in the sense that the next explored edge is older than the current edge. With these assumptions, one can unfold any directed temporal walk as a time-series (Figure 3A and 3B).

2.2.2 Streaming histograms as node embeddings. Given the above assumptions, we now formalize the approximation of random-walk based aggregations described in section 2.1.

In this framework, we do not consider random-walks with a fixed number of hops, and instead consider infinite walks, on top of which we compute embeddings analogously to exponential moving averages. The importance of older information compared to newer is controlled by a factor α between 0 and 1. A larger α gives more weight to features further away in the walk (or in the past), and we can therefore consider α the parameter that replaces the number of hops. The $(1 - \alpha)$ factor is used to ensure the weights of the weighted average sum to 1. Formally, let \vec{s}_i be a histogram with L bins, represented as an L -dimensional vector and characterizing the distribution of a feature f in the neighborhood of node i . A full infinite walk starting at node 0 computes the histogram \vec{s}_0 as:

$$\vec{s}_0 = \sum_{i=0}^{\infty} \alpha^i (1 - \alpha) \vec{\delta}(f_i) \quad (1)$$

where \sum is adding vectors, α is a discount factor between 0 and 1, controlling the importance of distant information in the summary \vec{s}_0 , and i denotes the hops of the walk ($i = 0$ being the newest node, or in other words the seed node of the infinite walk). f_i is the feature value at node i or edge i , and $\vec{\delta}(f_i)$ is an L -dimensional vector with element $\vec{\delta}_j = 1$ if the feature value f_i falls within bin j and $\vec{\delta}_j = 0$ for all other elements. Equation 1 then implements a streaming counts per bin, where older information is gradually forgotten. If the feature f_i is a node feature, then the value is taken from the current node. If it is an edge feature, then the feature value is taken from the edge connecting the current node and the chosen neighbor.

One could compute multiple such summaries per node, one for each node or edge feature of interest, and together they would summarize a neighborhood. The key idea is that we can now approximate the infinite random-walks, i.e., the infinite sum of equation 1, by performing only a finite number of $k \geq 1$ hops, followed by choosing a random neighbor of the last encountered node and choosing an available summary \vec{s}_k of that neighbor randomly, where \vec{s}_k is defined as

$$\vec{s}_k = \sum_{i=0}^{\infty} \alpha^i (1 - \alpha) \vec{\delta}(f_{i+k}) \quad (2)$$

With this strategy, we can approximate the summary \vec{s}_0 from equation 1 recurrently using

$$\vec{s}_0 \approx \sum_{i=0}^{k-1} \alpha^i (1 - \alpha) \vec{\delta}(f_i) + \alpha^k \vec{s}_k \quad (3)$$

Compared with equation 1, one now truncates the sum after k terms. Note that whenever the last histogram \vec{s}_k is normalized such that the bins sum to 1, e.g. using a uniform initialization for terminal nodes, equation 3 guarantees that all subsequent histograms will be normalized in the same way. Since we are interested in low-latency methods, we take the limit of $k = 1$ and Equation 3 becomes a streaming histogram:

$$\vec{s}_0 \leftarrow (1 - \alpha) \vec{\delta}(f_0) + \alpha \vec{s}_1 \quad (4)$$

The hyperparameter α can be chosen to depend on the number of hops or on time. When discounting by hops, this discount factor α is a fixed number between 0 and 1. When discounting by time,

the factor is made dependent on the difference in edge timestamps, for example exponentially or hyperbolically.

Using equation 4, one could approximate N (biased) random-walks by sampling N neighbors (non-uniformly), and subsequently combining the resulting histograms, e.g., by averaging. This would require performing N 1-hop look-ups each time.

Instead of that, we can increase efficiency even further by removing any stochasticity and updating a node's histogram at each edge arrival, combining the histograms of the two nodes involved in the arriving edge, as shown in equation 5:

$$\vec{s}_0 \leftarrow \beta \vec{s}_0 + (1 - \beta) \left((1 - \alpha) \vec{\delta}(f_0) + \alpha \vec{s}_1 \right) \quad (5)$$

In this way we combine all neighbours' information implicitly using a moving average over time.

Hyperparameter β is another discount factor between 0 and 1, controlling how much to focus on recent neighbor information in contrast to older information and which can optionally depend on time. In this way, we can update histograms in a fully streaming setting, using only information of each arriving edge. Nodes' summaries are initialized uniformly guaranteeing that the sum of all bins that belong to a certain edge is 1, e.g., if a feature has 10 bins then every bin is initialized with the value 0.1. We term this procedure *graph-sprints* and summarize it in algorithm 1.

Compared to equation 4, one can observe that the remaining sampling over single-hop neighbors is abolished, at the cost of imposing a more strict dependence on time. The advantage of algorithm 1 is that no list of neighbors needs to be stored. Moreover, algorithm 1 can be applied in parallel to both the source node and the destination node, and therefore edges are not required to be directed. In fact, while we derived equation 5 from random-walks, the attentive reader can notice that it can be interpreted as a special case of message passing where all neighbor summaries are aggregated using a weighted average, with weights that are biased by recency, and where the average is computed in a streaming fashion over time.

One special type of feature are the ones that have monotonically increasing values, exemplified by node degree features in a graph under conditions where edge deletion is non-existent. To avoid accumulating degrees over time, we propose to implement a streaming count of degrees per node. Every time an edge involving node u arrives, we compute

$$d_u = d_u \exp(-\Delta t / \tau_d) + 1 \quad (6)$$

where d_u denotes either in- or out-degree of node u , Δt denotes the time differences between the current edge involving node u and the previous one, and τ_d is a timescale for the streaming counts. This approach prevents values from growing too large and provides information about not only the structure but also the velocity and recency of edge additions. In financial crime detection, for example, it's important to know if an account starts having more transactions than usual suddenly.

2.2.3 Choosing histogram bins. Essential hyperparameters of this method are the choices of the boundaries of the histograms bins. We propose to use one bin per category for categorical features. If the cardinality of a certain feature is too high, we propose to form bins using groups of categories. For numerical features, one can

Algorithm 1 Graph-sprints (eq. 5)

Require: *EdgeStream* {Stream of arriving edges $e_{i,j}$ }

Require: \mathcal{F} {Set of features for GS (e.g., node degree)}

for $e_{v,u} \in \text{EdgeStream}$ **do**

Get \vec{s}_u, \vec{s}_v {Current summaries of nodes u, v }

$\vec{s}_v^* \leftarrow \alpha \vec{s}_v$ {Multiply all bins by α }

for $f \in \mathcal{F}$ **do**

if value(f) in bin j **then**

$\vec{s}_{vj}^* \leftarrow \vec{s}_{vj}^* + (1 - \alpha)$ {Add $(1 - \alpha)$ to bin j }

end if

end for

$\vec{s}_u \leftarrow \beta \vec{s}_u + (1 - \beta) \vec{s}_v^*$ {Updated summary of node u .}

end for

plot the distribution in the training data and choose sensible bin edges, for example on every 10th percentile of the distribution. The framework is not constrained by one choice of bins, as long as they can be updated in a streaming way.

2.3 Reducing space complexity

The space complexity of the graph-sprints approach (algorithm 1) is

$$M = |\mathcal{V}| \sum_{f \in \mathcal{F}} L_f \quad (7)$$

where $|\mathcal{V}|$ stands for the number of nodes, L_f stands for the number of bins of the histogram for feature f , and \mathcal{F} stands for the set of features chosen to collect in histograms. In case this memory is too high, we propose the following methods to reduce memory further.

Reducing histogram size using similarity hashing Following the similarity hashing approach proposed in Jin et al. [8], we extend the method to the streaming setting. All histograms as defined in the previous sections are normalized (in the sense that bin values sum to 1), and we can concatenate them into one vector \vec{s}_{tot} . We can now define a hash mapping by choosing k random hyperplanes in \mathbb{R}^M defined by unit vectors $\vec{h}_j, j = 1, \dots, k$.

The inner product between the histograms vector and the k unit vectors results in a vector of k values, each value θ can be calculated using the dot product of the unit vector \vec{h}_j and the histogram vector \vec{s}_{tot} , as illustrated in Equation 8. We use the superscript t to denote the current time step.

$$\theta_j^t = \vec{h}_j \cdot \vec{s}_{tot}^t \quad (8)$$

One can binarize the representation of the hashed vector using by taking the *sign* of the above θ_j^t .

Therefore, the resulting space complexity per node is k , replacing the number of bins in the memory M by the number of hash vectors k .

Importantly, the hashed histograms can be updated without storing any of the original histograms. Combining equations 4 and equation 8 and denoting $\vec{\delta}(f)$ the concatenation of the $\vec{\delta}$ vectors for all collected features, we get

$$\theta_j^{t+1} = \theta_j^t \cdot \alpha + \vec{h}_j \cdot \vec{\delta}(f) \cdot (1 - \alpha) \quad (9)$$

Therefore, we can compute the next hash θ_j^{t+1} or $\text{sign}(\theta_j^{t+1})$ directly from the previous θ_j^t and the new incoming features $\vec{\delta}(\vec{f})$. It is also important to note that this hashing scheme is preserved when averaging.

Reducing embedding size using feature importance One can reduce the needed memory by relying on feature importance techniques. One possibility is to train a classifier on the raw node and/or edge features and determine feature importances, after which only the top important features are used in the *graph-sprints* framework. Or similarly train on all bins and decide the bins to be used based on their importance in the classification task. Thus, either reducing the number of features, or the number of bins within the features, or both.

3 EXPERIMENTS

3.1 Experimental setup

We assess the quality of the graph based features generated by the **graph-sprints** framework on two different tasks, namely, node classification and link prediction. We use three publicly available datasets from the social and education domains and two proprietary datasets from money laundering domain. We detail their main characteristics in Table 1 and Table 7, respectively. All datasets are CTDGs and are labeled. Each dataset is split into train, validation, and test sets respecting time (i.e., all events in the train are older than the events in validation, and all events in validation are older than the events in the test set). We use Optuna [1] to optimize the hyperparameters of all models, training 100 models using the TPE sampler and with 40 warmup trials. Each model trains using early stopping with a patience of 10 epochs, where the early stopping metric computed on the validation set as area under ROC curve (AUC) for node classification and average precision (AP) for link prediction. All models were trained using a batch size of 200 edges. Hyperparameter ranges used during optimization are reported in Table 2, and the optimized values are reported in Table 5.

3.1.1 Baselines. As a first baseline, we reproduce a state-of-the-art GNN model for CTDGs, the temporal-graph network (TGN) [16], which leverages a combination of memory modules and graph-based operators to obtain node representations. As an important note, we mention that the pytorch geometric [4] implementation of TGN was used, for which the sampling of neighbors uses a different strategy than the original TGN implementation. Indeed, the original paper allowed to sample from interactions within the same batch as

Table 1: Information about public data [12]. We adopt the identical data partitioning strategy employed by the baseline methods we compare against, which also utilized these datasets.

	Wikipedia	Mooc	Reddit
#Nodes	9,227	7,047	10,984
#Edges	157,474	411,749	672,447
Label type	editing ban	student drop-out	posting ban
Positive labels	0.14%	0.98%	0.05%
Used split	75%-15%-15%	60%-20%-20%	75%-15%-15%

long as they are older, while the pytorch geometric implementation does not allow within-batch information to be used. As also noted in the pytorch geometric documentation, we believe the latter to be more realistic. As a consequence, our TGN results are not directly comparable with the originally published TGN performances. In any case, the graph-sprints embeddings were computed using the same batch size and therefore also do not have access to within-batch information, allowing a fair comparison between the algorithms.

Two variations of the TGN architecture were used. First, TGN-attn was implemented, which was the most powerful variation in the original paper but is expected to be slower due to the graph-attention operations. Second, TGN-ID was implemented, which is a variation of the TGN where no graph-embedding operators are used, and only the embedding resulting from the memory module is passed to the classification layers.

A third baseline we use is Jodie [12]. We use the TGN implementation of Jodie, where instead of using Graph attention embeddings on top of the memory embedding, a time projection embedding module is used and where the loss function is otherwise identical to the TGN setting. For a fair comparison with TGN we use the same memory updater module, namely, gated recurrent units (GRUs).

The TGN-ID and Jodie baselines do not require sampling of neighbors, and were therefore chosen as lower-latency baselines compared to TGN-attn.

3.1.2 Graph-sprints and classifier. For each arriving edge, we apply the graph-sprints feature update (algorithm 1) to both the source node and the destination node in parallel. All edge features are used for the computation of the graph-sprints features, and for each feature bin edges are chosen as the 10 quantiles computed on the training data. Since the graph-sprints framework only creates features, a classifier is implemented for the classification tasks. We chose to implement a neural network consisting of dense layers, normalization layers, and skip-connections across every two dense layers. Hyperparameter optimization proceeds in two steps. First, default parameters for the classifier are used to optimize the discount factors of the *graph-sprints* framework, α and β . For this step, 50 models are trained. Subsequently, hyperparameter optimization of the classifier follows same approach as TGN, training 100 models.

In all experiments, we test the following three cases. Firstly, we train the classifier using only raw features (Raw). We then train the classifier using only the graph-sprint features (GS). Finally, we train the classifier using both raw and graph-sprint features (GS+Raw).

3.1.3 Node Classification vs Link Prediction. For the node classification task on the Wikipedia, Reddit and Mooc datasets, we concatenate the source and destination node embeddings and feed the concatenated vector to the classifier, as is usual for these datasets since labels are on the edge level, for instance, in Mooc dataset where edges connect students to courses, labels indicating whether the student will drop-out are on the edge level.

For the link prediction task, all existing edges are considered positive edges, and negative edges are generated following the same approach as the original TGN paper [16], a negative edge is sampled for every positive one. We perform the link prediction task both in the *transductive* and *inductive* settings. In the *transductive* setting, negative edges are sampled on the same graph used for training. Conversely, in the *inductive* setting, the sampled negative edges are

constrained to include at least one new node which was not used in the training graph.

Table 2: Hyperparameters ranges for the Graph-Sprints (GS) method and the other baselines (GNN)

Method	Hyperparameter	min	max
GS	α	0.1	1
GS	β	0.1	1
GNN/GS	Learning rate	0.0001	0.01
GNN/GS	Dropout perc	0.1	0.3
GNN/GS	Weight decay	0.000000001	0.001
GNN/GS	Num of dense layers	1	3
GNN/GS	Size of dense layer	32	256
GNN/GS	Batch size	128	2024
GNN	Memory size	32	256
GNN	Neighbors per node	5	10
GNN	Num GNN layers	1	3
GNN	Size GNN layer	32	256

3.2 Public datasets experiments

3.2.1 Task performance. In Table 3 we report the average test AUC \pm std for the **Node classification** task. Our approach involved retraining the best model obtained after hyperparameter optimization, using 10 different random seeds. We can observe that on all datasets, the best model for node classification uses a variation of our graph-sprint framework (either GS or GS+Raw). In table 4 we report the average test AUC \pm std, along with the average precision (AP) \pm std for the **Link prediction** task. Results were again computed after retraining the best model obtained through hyperparameter optimization, utilizing 10 distinct random seeds. We can observe that the graph-sprints model is the best for link prediction on the Mooc dataset, but performs poorly on the Wikipedia dataset. On the Reddit dataset, the graph-sprints model is second in the transductive setting, but is considerably worse for the inductive setting.

Table 3: Results on Node classification task, we report the average test AUC \pm std achieved by retraining the best model after hyperparameter optimization using 10 random seeds. Our models, Raw, GS, and GS+Raw, use the same ML classifier but differ in the features employed for training. Raw uses raw edge features, GS uses graph-sprints histograms, and GS+Raw combines both. We identify the **best model and highlight the **second best** model.**

Method	AUC \pm std		
	Wikipedia	Mooc	Reddit
Raw	58.5 \pm 2.2	62.8 \pm 0.9	55.3 \pm 0.8
TGN-ID	88.9 \pm 0.2	63.0 \pm 17	61.3 \pm 2.0
Jodie	87.2 \pm 0.9	63.7 \pm 16.7	61.9 \pm 2.0
TGN-attn	86.6 \pm 2.8	75.8 \pm 0.4	67.9 \pm 1.6
GS	90.7 \pm 0.3	75.0 \pm 0.2	68.5 \pm 1.0
GS+Raw	89.2 \pm 0.4	76.5 \pm 0.3	63.7 \pm 0.4




Figure 4: Speedup vs. number of edges. The speedups increase almost linearly with the number of edges in the graph.

3.2.2 Inference runtime. We compare the latency of our framework to baseline GNN architectures. For this purpose, we run 200 batches of 200 events on the external datasets, Wikipedia, Mooc, and Reddit using the node classification task. We compute the average time over 10 runs. Both models were running on Linux PC with 24 Intel Xeon CPU cores (3.70GHz) and a NVIDIA GeForce RTX 2080 Ti GPU (11GB). As depicted in Figure 2, our graph-sprints consistently outperforms other baselines (TGN-attn, TGN-ID, Jodie) in the node classification task while also demonstrating a significantly lower inference latency. Compared to TGN-attn, the GS achieves better classification results but is close to an order of magnitude faster (Figure 4).

To investigate the impact of graph size on runtime, Figure 4 shows our observations. Notably, in the utilized datasets, TGN’s runtime increases as the number of edges in the dataset grows, requiring more time to score 200 batches. Conversely, since GS does not require neighborhood sampling, it exhibits constant inference time regardless of the graph size. Furthermore, the speedups achieved by graph-sprints are expected to be significantly higher in a big-data context, where the data is stored in a distributed manner rather than in memory as in our current experiments. In such scenarios, graph operations used in graph-neural networks like TGN-attn would incur even higher computational costs.

Recently, APAN [22] has attempted to build a low-latency framework for CTDGs. Their approach consisted of performing the expensive graph operations asynchronously, out of the inference loop. In that way, they achieved inference speeds of 4.3ms per batch on the Wikipedia dataset, but we cannot directly compare those results with ours (GS: 1.4ms, TGN-attn: 6ms) due to the different setup and hardware. Importantly, their approach achieves low-latency by sacrificing up-to-date information at inference time. Indeed, the inference step is performed without access to the most recent embeddings, because the expensive graph operations to compute the embeddings are performed asynchronously.

3.2.3 Memory reduction. Both the Wikipedia and Reddit datasets consist of 172 edge features. By calculating graph-sprints with 10 quantiles per feature, along with incorporating in/out degrees histograms and time-difference histograms, we obtain a node embedding of 1742 features (one feature per histogram bin). In our

Table 4: Results on Link prediction task, we report test average AUC and average precision (AP) \pm std resulting from retraining the best model after hyperparameter optimization using 10 random seeds. We report results on both Transductive (T) or Inductive (I) settings. Our models, GS, and GS+Raw, use the same ML classifier but differ in the features employed for training. GS uses GS histograms, and GS+Raw combines GS histograms with raw edge features. We identify the **best model and highlight the **second best** model.**

	Method	Wikipedia		Mooc		Reddit	
		AUC	AP	AUC	AP	AUC	AP
T	TGN-ID	<u>95.6 \pm 0.2</u>	<u>95.8 \pm 0.1</u>	80.4 \pm 5.8	75.0 \pm 6.1	94.7 \pm 0.7	93.2 \pm 1.0
	Jodie	94.3 \pm 0.3	94.5 \pm 0.3	<u>85.1 \pm 1.8</u>	<u>80.0 \pm 3.5</u>	94.9 \pm 1.2	93.4 \pm 1.7
	TGN-attn	97.0 \pm 0.3	97.3 \pm 0.3	80.3 \pm 8.1	75.6 \pm 8.4	96.1 \pm 0.4	95.1 \pm 0.7
	GS	92.5 \pm 0.6	92.9 \pm 0.7	82.7 \pm 0.8	81.1 \pm 0.7	<u>96.1 \pm 0.2</u>	<u>95.3 \pm 0.3</u>
	GS+Raw	92.1 \pm 0.4	92.6 \pm 0.4	85.4 \pm 0.3	83.7 \pm 0.3	96.8 \pm 0.1	96.1 \pm 0.2
I	TGN-ID	<u>92.2 \pm 0.2</u>	<u>92.8 \pm 0.2</u>	68.5 \pm 8.6	63.5 \pm 6.7	93.4 \pm 0.6	92.2 \pm 0.8
	Jodie	87.0 \pm 0.6	89.1 \pm 0.7	71.1 \pm 2.2	66.1 \pm 2.9	92.3 \pm 1.3	90.8 \pm 1.9
	TGN-attn	94.5 \pm 0.2	95.0 \pm 0.2	71.4 \pm 4.1	66.9 \pm 3.9	95.0 \pm 0.4	94.3 \pm 0.5
	GS	92.0 \pm 0.3	91.7 \pm 0.4	<u>78.2 \pm 0.6</u>	<u>76.5 \pm 0.6</u>	92.7 \pm 0.5	92.7 \pm 0.6
	GS+Raw	91.4 \pm 0.2	91.1 \pm 0.3	83.0 \pm 0.5	80.3 \pm 0.5	<u>93.5 \pm 0.4</u>	<u>92.2 \pm 0.5</u>

Table 5: Hyperparameters used in node classification (NC) and link prediction (LP). Learning rate and Weight decay are approximated due to space constraints

Task	Hyperparameter	Wikipedia	Mooc	Reddit
NC	α	0.7	0.65	0.7
	β	0.15	0.1	0.1
	Learning rate	0.0003	0.0001	0.0003
	Dropout perc	0.1	0.3	0.1
	Weight decay	1.1e-07	0.0005	1.1e-07
	Num of dense layers	12	4	4
	Size of dense layer	48	32	48
LP	Batch size	512	2048	512
	α	0.65	0.65	0.15
	β	0.55	0.45	0.15
	Learning rate	4.9e-05	0.002	5.8e-05
	Dropout perc	0.2	0.3	0.2
	Weight decay	8.3e-06	3.5e-05	7.4e-07
	Num of dense layers	14	20	18
	Size of dense layer	64	64	48
	Batch size	1024	512	512

experimental setup, similar to state-of-the-art approaches, we concatenate the source and destination node embeddings for source label prediction, resulting in a 3484-feature vector. To reduce the size of the node embeddings, we propose a similarity hashing-based memory reduction technique (Section 2.3)). Our experiments, as presented in Table 6, demonstrate that our technique significantly reduces storage requirements sacrificing the AUC in the node classification task. In the Reddit dataset, storage can be reduced to 50% with a 0.6% AUC sacrifice or to 10% with a 2% AUC sacrifice. The reduction percentage can be fine-tuned as a hyperparameter, considering the use case and dataset, to strike a balance between precision and memory trade-off.

Table 6: Effect of memory reduction on the node classification task. We report average test AUC \pm std resulting from retraining the best model, after optimization, using 10 random seeds. We use only the GS embeddings to train the classifier. Since Mooc has less features we add '-' where number of features < 1.

Space used	Wikipedia	Mooc	Reddit
100%	90.7 \pm 0.3	75.0 \pm 0.2	68.5 \pm 1.0
50%	90.8 \pm 0.1	75.0 \pm 0.1	67.9 \pm 1.1
25%	91.1 \pm 0.1	74.9 \pm 0.3	65.1 \pm 1.9
10%	90.9 \pm 0.2	74.0 \pm 0.3	66.5 \pm 0.9
0.5%	89.7 \pm 0.3	-	58.0 \pm 2.6
0.12%	86.4 \pm 0.3	-	55.2 \pm 1.1

3.3 Anti-money laundering (AML) experiments

In money laundering, the criminals' objective is to hide the illegal source of their money by moving funds between various accounts and financial institutions. In these experiments, our objective is to enrich a classifier with graph-based features generated by our *graph-sprints* framework.

3.3.1 Datasets. We evaluate the graph-sprints framework in the AML domain using two real-world banking datasets. Due to privacy concerns, we can not disclose the identity of the financial institutions (FIs) nor provide exact details regarding the node features. We refer to the datasets as *FI-A* and *FI-B*. The graphs in this use-case are constructed by considering the accounts as nodes and the money transfers between accounts as edges. Table 7 shows approximate details of these datasets.

3.3.2 Models and Experimental setup. As before, we train the neural network classifier that uses raw node features only, i.e., no graph information is present (Raw). We compare that baseline performance against models that include only *graph-sprint* features

(GS), and models that use both *graph-sprints* features and raw features (GS+Raw). Finally, we train the same GNN architectures as in the public datasets (TGN-ID, Jodie, and TGN-attn).

Table 7: Information about AML datasets.

	FI-A	FI-B
#Nodes	≈ 400000	≈ 10000
#Edges	≈ 500000	≈ 2000000
Positive labels	2-5%	20-40%
Duration	≈ 300 days	≈ 600 days
Edges/day (mean \pm std)	1500 ± 750	3000 ± 5000
Used split	60%-10%-30%	60%-10%-30%

3.3.3 Results. Due to privacy considerations, we are unable to disclose the actual obtained AUC values. Instead, we present the relative improvements in AUC (Δ AUC) when compared to a baseline model that does not utilize graph features. In this context, the baseline model corresponds to a Δ AUC value of 0, and any increase in recall compared to the baselines is represented by positive values of Δ AUC.

Table 8 displays the Δ AUC values for our GS variations and other state-of-the-art baselines. Our GS variations exhibit the most favorable outcomes in both datasets, with an approximate 3.3% improvement in AUC for the *FI-A* dataset and a 27.8% improvement in AUC for the *FI-B* dataset.

Table 8: Node classification results in AML data. We report relative gain in AUC (Δ AUC) compared to a baseline model that does not utilize graph features. We report the average test Δ AUC \pm std achieved by retraining the best model after hyperparameter optimization using 10 random seeds. We identify the **best model and highlight the **second best** model.**

Method	Δ AUC \pm std	
	FI-A	FI-B
TGN-ID	+0.1 \pm 0.1	+24.4 \pm 0.2
Jodie	+0.0 \pm 0.1	+24.5 \pm 0.2
TGN-attn	+0.3 \pm 0.7	+25.1 \pm 0.3
GS	+1.8 \pm 0.5	+27.8 \pm 0.4
GS+Raw	+3.3 \pm 0.3	+20.1 \pm 3.9

4 RELATED WORK

4.1 Random-walk based methods

DeepWalk [15] and node2vec [6] are two high-latency random-walk based methods to extract node embeddings on static graphs. The main limiting factors of these methods are that they disregard node and edge features as well as temporal information, since they are designed for static graphs. Sajjad et al. [17] extend these random-walk based methods to discrete-time dynamic graphs. While some efficiency is gained with the proposed method, it is far from applicable to CTDGs and in low-latency settings. Node2bits [8] considers the temporal information by defining several time windows over the sampled random-walks, and aggregates node attributes in

these time windows into histograms. Node2bits does not include edge-features and performs costly computations which cannot be performed with low latency.

Continuous-time Dynamic Node Embeddings (CTDNE) [13, 14] were proposed to generate time-aware embeddings, generalizing the node2vec framework to CTDGs. The authors consider the graph as a stream of edges, and propose to perform temporal walks starting from seed nodes chosen from a temporally biased distribution. Similarly, temporal random-walks have been used to extract embeddings into hyperbolic spaces [21]. Causal anonymous walks [23] propose to use anonymized walks in order to encode motif information. Similarly, NeurTWs [11] explicitly model time in the anonymous walks using Neural Ordinary Differential Equations (NeuralODEs). Unlike our graph-sprints framework, full random-walks need to be performed.

4.2 K-hop neighborhood based methods

Most GNN-based methods require a K-hop neighborhood on which message-passing operations lead to node embeddings. To deal with CTDGs, a simple approach is to consider a series of discrete snapshots of the graph over time, on which static methods are applied. Such approaches however do not take time properly into account and several works propose techniques to alleviate this issue [5, 10, 18, 28]. To better deal with CTDGs, other works focus on including time-aware features or inductive biases into the architecture. DeepCoevolve [2] and Jodie [12] train two RNNs for bipartite graphs, one for each node type. Importantly, the previous hidden state of one RNN is also added as an input to the other RNN. In this way, the two RNNs interact, in essence performing single-hop graph aggregations. TGAT [26] proposes to include temporal information in the form of time encodings, while TGN [16] extends this framework and also includes a memory module taking the form of a recurrent neural network. In [9], the authors replace the discrete-time recurrent network of TGN with a NeuralODE modeling the continuous dynamics of node embeddings.

APAN [22] proposes to reduce the latency at inference time by decoupling the more costly graph operations from the inference module. The authors propose a more light-weight inference module that computes the predictions based on a node’s embedding as well as a node’s mailbox, which contains messages of recently interacting nodes. The mailbox is updated asynchronously, i.e. separated from the inference module, and involves the more expensive k-hop message passing. While APAN improves the latency at inference time, it sacrifices some memory since each node’s state is now expanded with a mailbox and it potentially uses outdated information at inference time due to asynchronous update of this mailbox. Also towards reducing computational costs of GNNs, HashGNN [24] leverages MinHash to generate node embeddings suitable for the link prediction task, where nodes that results in the same hashed embedding are considered similar. SGSketch [27] is a streaming node embedding framework uses a mechanism to gradually forget outdated edges, achieving significant speedups. Differently than our approach SGSketch uses the gradual forgetting strategy to update the adjacency matrix and therefore only considers the graph structure.

5 CONCLUSIONS

This paper introduces the graph-sprints framework, which enables the computation of time-aware embeddings for CTDGs with minimal latency. The study demonstrates that the graph-sprints features, when combined with a neural network classifier, achieve competitive predictive performance compared to state-of-the-art methods while having a significantly faster inference time, up to approximately an order of magnitude improvement. Interestingly, our method performs more strongly on the node classification tasks compared to the link prediction tasks. It would be interesting to investigate what features make it powerful on one task, and what is still missing in the other. In future work, it would be interesting to extend the graph-sprints framework to heterogeneous graphs, and explore how GNNs could inherit some of the strengths of graph-sprints.

REFERENCES

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2623–2631.
- [2] Hanjun Dai, Yichen Wang, Rakshit Trivedi, and Le Song. 2016. Deep coevolutionary network: Embedding user and item features for recommendation. *arXiv preprint arXiv:1609.03675* (2016).
- [3] Wengqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *The world wide web conference*. 417–426.
- [4] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [5] Palash Goyal, Nitin Kamra, Xinran He, and Yan Liu. 2018. Dyngem: Deep embedding method for dynamic graphs. *arXiv preprint arXiv:1805.11273* (2018).
- [6] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [7] Jin Guo, Zhen Han, Zhou Su, Jiliang Li, Volker Tresp, and Yuyi Wang. 2022. Continuous Temporal Graph Networks for Event-Based Graph Data. *arXiv preprint arXiv:2205.15924* (2022).
- [8] Di Jin, Mark Heimann, Ryan A Rossi, and Danai Koutra. 2019. Node2bits: Compact time- and attribute-aware node representations for user stitching. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 483–506.
- [9] Di Jin, Sungchul Kim, Ryan A Rossi, and Danai Koutra. 2020. From static to dynamic node embeddings. *arXiv preprint arXiv:2009.10017* (2020).
- [10] Di Jin, Sungchul Kim, Ryan A Rossi, and Danai Koutra. 2022. On Generalizing Static Node Embedding to Dynamic Settings. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 410–420.
- [11] Ming Jin, Yuan-Fang Li, and Shirui Pan. 2022. Neural Temporal Walks: Motif-Aware Representation Learning on Continuous-Time Dynamic Graphs. In *Advances in Neural Information Processing Systems*.
- [12] Srikanth Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks. In *Proceedings of the 25th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM.
- [13] John Boaz Lee, Giang Nguyen, Ryan A Rossi, Nesreen K Ahmed, Eunye Koh, and Sungchul Kim. 2020. Dynamic node embeddings from edge streams. *IEEE Transactions on Emerging Topics in Computational Intelligence* 5, 6 (2020), 931–946.
- [14] Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunye Koh, and Sungchul Kim. 2018. Continuous-time dynamic network embeddings. In *Companion proceedings of the the web conference 2018*. 969–976.
- [15] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [16] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. 2020. Temporal Graph Networks for Deep Learning on Dynamic Graphs. In *ICML 2020 Workshop on Graph Representation Learning*.
- [17] Hooman Peiro Sajjad, Andrew Docherty, and Yuriy Tyshetskiy. 2019. Efficient representation learning using random walks for dynamic graphs. *arXiv preprint arXiv:1901.01346* (2019).
- [18] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. 2020. Dysat: Deep neural representation learning on dynamic graphs via self-attention networks. In *Proceedings of the 13th international conference on web search and data mining*. 519–527.
- [19] Amauri Souza, Diego Mesquita, Samuel Kaski, and Vikas Garg. 2022. Provably expressive temporal graph networks. *Advances in Neural Information Processing Systems* 35 (2022), 32257–32269.
- [20] Jianan Wang, Sheng Zhang, Yanghua Xiao, and Rui Song. 2021. A review on graph neural network methods in financial applications. *arXiv preprint arXiv:2111.15367* (2021).
- [21] Lili Wang, Chenghan Huang, Weicheng Ma, Ruibo Liu, and Soroush Vosoughi. 2021. Hyperbolic node embedding for temporal networks. *Data Mining and Knowledge Discovery* 35, 5 (2021), 1906–1940.
- [22] Xuhong Wang, Ding Lyu, Mengjian Li, Yang Xia, Qi Yang, Xinwen Wang, Xinguang Wang, Ping Cui, Yupu Yang, Bowen Sun, et al. 2021. Apan: Asynchronous propagation attention network for real-time temporal graph embedding. In *Proceedings of the 2021 international conference on management of data*. 2628–2638.
- [23] Yanbang Wang, Yen-Yu Chang, Yunyu Liu, Jure Leskovec, and Pan Li. 2021. Inductive representation learning in temporal networks via causal anonymous walks. *arXiv preprint arXiv:2101.05974* (2021).
- [24] Wei Wu, Bin Li, Chuan Luo, and Wolfgang Nejdl. 2021. Hashing-accelerated graph neural networks for link prediction. In *Proceedings of the Web Conference 2021*. 2910–2920.
- [25] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [26] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2020. Inductive representation learning on temporal graphs. *arXiv preprint arXiv:2002.07962* (2020).
- [27] Dingqi Yang, Bingqing Qu, Jie Yang, Liang Wang, and Philippe Cudre-Mauroux. 2022. Streaming graph embeddings via incremental neighborhood sketching. *IEEE Transactions on Knowledge and Data Engineering* 35, 5 (2022), 5296–5310.
- [28] Jiaxuan You, Tianyu Du, and Jure Leskovec. 2022. ROLAND: graph learning framework for dynamic graphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2358–2366.
- [29] Xiao-Meng Zhang, Li Liang, Lin Liu, and Ming-Jing Tang. 2021. Graph neural networks and their current applications in bioinformatics. *Frontiers in genetics* 12 (2021), 690049.