

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

10-2020

Efficient sampling algorithms for approximate temporal motif counting

Jingjing WANG
Hunan University

Yanhao WANG
National University of Singapore

Wenjun JIANG
Hunan University

Yuchen LI
Singapore Management University, yuchenli@smu.edu.sg

Kian-Lee TAN
National University of Singapore

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Numerical Analysis and Scientific Computing Commons](#), and the [Theory and Algorithms Commons](#)

Citation

WANG, Jingjing; WANG, Yanhao; JIANG, Wenjun; LI, Yuchen; and TAN, Kian-Lee. Efficient sampling algorithms for approximate temporal motif counting. (2020). *CIKM '20: Proceedings of the 29th ACM International Conference on Information and Knowledge Management, Virtual, October 19-23*. 1505-1514. Available at: https://ink.library.smu.edu.sg/sis_research/5934

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Efficient Sampling Algorithms for Approximate Temporal Motif Counting

Jingjing Wang*
Hunan University
wangjingjing2019@hnu.edu.cn

Yanhao Wang*
National University of Singapore
yanhao90@comp.nus.edu.sg

Wenjun Jiang†
Hunan University
jiangwenjun@hnu.edu.cn

Yuchen Li
Singapore Management University
Alibaba-Zhejiang University Joint
Research Institute of Frontier
Technologies
yuchenli@smu.edu.sg

Kian-Lee Tan
National University of Singapore
tankl@comp.nus.edu.sg

ABSTRACT

A great variety of complex systems ranging from user interactions in communication networks to transactions in financial markets can be modeled as *temporal graphs*, which consist of a set of vertices and a series of timestamped and directed edges. *Temporal motifs* in temporal graphs are generalized from subgraph patterns in static graphs which take into account edge orderings and durations in addition to structures. Counting the number of occurrences of temporal motifs is a fundamental problem for temporal network analysis. However, existing methods either cannot support temporal motifs or suffer from performance issues. In this paper, we focus on approximate temporal motif counting via random sampling. We first propose a generic edge sampling (ES) algorithm for estimating the number of instances of any temporal motif. Furthermore, we devise an improved EWS algorithm that hybridizes edge sampling with wedge sampling for counting temporal motifs with 3 vertices and 3 edges. We provide comprehensive analyses of the theoretical bounds and complexities of our proposed algorithms. Finally, we conduct extensive experiments on several real-world datasets, and the results show that our ES and EWS algorithms have higher efficiency, better accuracy, and greater scalability than the state-of-the-art sampling method for temporal motif counting.

CCS CONCEPTS

• **Mathematics of computing** → **Graph algorithms.**

KEYWORDS

temporal networks; motif counting; random sampling

*Both authors contributed equally to this work.

†Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6859-9/20/10...\$15.00

<https://doi.org/10.1145/3340531.3411862>

ACM Reference Format:

Jingjing Wang, Yanhao Wang, Wenjun Jiang, Yuchen Li, and Kian-Lee Tan. 2020. Efficient Sampling Algorithms for Approximate Temporal Motif Counting. In *The 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, October 19–23, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3340531.3411862>

1 INTRODUCTION

Graphs are one of the most fundamental data structures that are widely used for modeling complex systems across diverse domains from bioinformatics [30], to neuroscience [38], to social sciences [5]. Modern graph datasets increasingly incorporate temporal information to describe the dynamics of relations over time. Such graphs are referred to as *temporal graphs* [11] and typically represented by a set of vertices and a sequence of timestamped and directed edges between vertices called *temporal edges*. For example, a communication network [9, 46–48, 50] can be denoted by a temporal graph where each person is a vertex and each message sent from one person to another is a temporal edge. Similarly, computer networks and financial transactions can also be modeled as temporal graphs. Due to the ubiquitousness of temporal graphs, they have attracted much attention [6, 8, 9, 20, 25, 27, 32, 50] recently.

One fundamental problem in temporal graphs with wide real-world applications such as network characterization [27], structure prediction [22], and fraud detection [18], is to count the number of occurrences of small (connected) subgraph patterns (i.e., *motifs* [24]). To capture the temporal dynamics in network analysis, the notion of *motif* [16, 17, 22, 27] in temporal graphs is more general than its counterpart in static graphs. It takes into account not only the subgraph structure (i.e., *subgraph isomorphism* [7, 36]) but also the temporal information including edge ordering and motif duration. As an illustrative example, M and M' in Figure 1 are different temporal motifs. Though M and M' have exactly the same structure, they are different in the ordering of edges. Consequently, although there has been a considerable amount of work on subgraph counting in static graphs [2, 4, 14, 15, 29, 34, 35, 42, 43, 45], they cannot be used for temporal motif counting directly.

Generally, it is a challenging task to count temporal motifs. Firstly, the problem is at least as hard as subgraph counting in static graphs, whose time complexity increases exponentially with




Figure 1: Examples for temporal motifs

the number of edges in the query subgraph. Secondly, it becomes even more computationally difficult because the temporal information is considered. For example, counting the number of instances of k -stars is simple in static graphs; however, counting temporal k -stars is proven to be NP-hard [22] due to the combinatorial nature of edge ordering. Thirdly, temporal graphs are a kind of *multigraph* that is permitted to have multiple edges between the same two vertices at different timestamps. As a result, there may exist many different instances of a temporal motif within the same set of vertices, which leads to more challenges for counting problems. There have been a few methods for exact temporal motif counting [27] or enumeration [18, 23]. However, they suffer from efficiency issues and often cannot scale well in massive temporal graphs with hundreds of millions of edges [22].

In many scenarios, it is not necessary to count motifs exactly, and finding an approximate number is sufficient for practical use. A recent work [22] has proposed a sampling method for approximate temporal motif counting. It partitions a temporal graph into equal-time intervals, utilizes an exact algorithm [23] to count the number of motif instances in a subset of intervals, and computes an estimate from the per-interval counts. However, this method still cannot achieve satisfactory performance in massive datasets. On the one hand, it fails to provide an accurate estimate when the sampling rate and length of intervals are small. On the other hand, its efficiency does not significantly improve upon that of exact methods when the sampling rate and length of intervals are too large.

Our Contributions: In this paper, we propose more efficient and accurate sampling algorithms for approximate temporal motif counting. First of all, we propose a generic Edge Sampling (ES) algorithm to estimate the number of instances of any k -vertex l -edge temporal motif in a temporal graph. The basic idea of our ES algorithm is to first uniformly draw a set of random edges from the temporal graph, then exactly count the number of *local* motif instances that contain each sampled edge by enumerating them, and finally compute the global motif count from local counts. The ES algorithm exploits the BACKTRACKING (BT) algorithm [23, 36] for subgraph isomorphism to enumerate local motif instances. We devise simple heuristics to determine the matching order of a motif for the BT algorithm to reduce the search space.

Furthermore, temporal motifs with 3 vertices and 3 edges (i.e., triadic patterns) are one of the most important classes of motifs, whose distribution is an indicator to characterize temporal networks [5, 15, 27, 37]. Therefore, we propose an improved Edge-Wedge Sampling (EWS) algorithm that combines *edge sampling* with *wedge sampling* [15, 35] specialized for counting any 3-vertex 3-edge temporal motif. Instead of enumerating all instances containing a sampled edge, the EWS algorithm generates a sample of *temporal wedges* (i.e., 3-vertex 2-edge motifs) and estimates the number of local instances by counting how many edges can match the query motif together with each sampled temporal wedge. In

this way, EWS avoids the computationally intensive enumeration and greatly improves the efficiency upon ES. Moreover, we analyze the theoretical bounds and complexities of both ES and EWS.

Finally, we test our algorithms on several real-world datasets. The experimental results confirm the efficiency and effectiveness of our algorithms: ES and EWS can provide estimates with relative errors less than 1% and 2% in 37.5 and 2.3 seconds on a temporal graph with over 100M edges, respectively. In addition, they run up to 10.3 and 48.5 times faster than the state-of-the-art sampling method while having lower estimation errors.

Organization: The remainder of this paper is organized as follows. Section 2 reviews the related work. Section 3 introduces the background and formulation of *temporal motif counting*. Section 4 presents the ES and EWS algorithms for temporal motif counting and analyzes them theoretically. Section 5 describes the setup and results of the experiments. Finally, Section 6 provides some concluding remarks.

2 RELATED WORK

Random Sampling for Motif Counting: In recent years, there have been great efforts to (approximately) count the number of occurrences of a motif in a large graph via random sampling. First of all, many sampling methods such as subgraph sampling [33], edge sampling [1, 21, 44], color sampling [26], neighborhood sampling [28], wedge sampling [13, 15, 34, 35], and reservoir sampling [3], were proposed for approximate triangle counting (see [49] for an experimental analysis). Moreover, sampling methods were also used for estimating more complex motifs, e.g., 4-vertex motifs [14, 31], 5-vertex motifs [29, 42, 43, 45], motifs with 6 or more vertices [2], and k -cliques [12]. However, all above methods were proposed for static graphs and did not consider the temporal information and ordering of edges. Thus, they could not be applied to temporal motif counting directly.

Motifs in Temporal Networks: Prior studies have considered different types of *temporal network motifs*. Viard et al. [39, 40] and Himmel et al. [10] extended the notion of *maximal clique* to temporal networks and proposed efficient algorithms for maximal clique enumeration. Li et al. [20] proposed the notion of (θ, τ) -*persistent k -core* to capture the persistence of a community in temporal networks. However, these notions of temporal motifs were different from ours since they did not take *edge ordering* into account. Zhao et al. [50] and Gurukar et al. [9] studied the *communication motifs*, which are frequent subgraphs to characterize the patterns of information propagation in social networks. Kovanen et al. [17] and Kosyfaki et al. [16] defined the *flow motifs* to model flow transfer among a set of vertices within a time window in temporal networks. Although both definitions accounted for edge ordering, they were more restrictive than ours because the former assumed any two adjacent edges must occur within a fixed time span while the latter assumed edges in a motif must be consecutive events for a vertex [27].

Temporal Motif Counting & Enumeration: There have been several existing studies on counting and enumerating temporal motifs. Paranjape et al. [27] first formally defined the notion of *temporal motifs* we use in this paper. They proposed exact algorithms for counting temporal motifs based on subgraph enumeration in

static graphs and timestamp-based pruning. Kumar and Calders [18] proposed an efficient algorithm called 2SCENT to enumerate all simple temporal cycles in a directed interaction network. Although 2SCENT was shown to be effective for cycles, it could not be used for enumerating temporal motifs of any other type. Mackey et al. [23] proposed an efficient BACKTRACKING algorithm for temporal subgraph isomorphism. The algorithm could count temporal motifs exactly by enumerating all of them. Liu et al. [22] proposed an interval-based sampling framework for counting temporal motifs. To the best of our knowledge, this is the only existing work on approximate temporal motif counting via sampling. In this paper, we present two improved sampling algorithms for temporal motif counting and compare them with the algorithms in [18, 22, 23, 27] for evaluation.

3 PRELIMINARIES

In this section, we formally define *temporal graphs*, *temporal motifs*, and the problem of *temporal motif counting* on a temporal graph. Here, we follow the definition of *temporal motifs* in [22, 23, 27] for its simplicity and generality. Other types of temporal motifs have been discussed in Section 2.

Temporal Graph: A *temporal graph* $T = (V_T, E_T)$ is defined by a set V_T of n vertices and a sequence E_T of m temporal edges among vertices in V_T . Each temporal edge $e = (u, v, t)$ where $u, v \in V_T$ and $t \in \mathbb{R}^+$ is a timestamped directed edge from u to v at time t . There may be many temporal edges from u to v at different timestamps (e.g., a user can comment on the posts of another user many times on *Reddit*). For ease of presentation, we assume the timestamp t of each temporal edge e is unique so that the temporal edges in E_T are strictly ordered. Note that our algorithms can also handle the case when timestamps are non-unique by using any consistent rule to break ties.

Temporal Motif: We formalize the notion of *temporal motifs* [22, 27] in the following definition.

Definition 3.1 (Temporal Motif). A temporal motif $M = (V_M, E_M, \sigma)$ consists of a (connected) graph with a set of k vertices V_M and a set of l edges E_M , and an ordering σ on the edges in E_M .

Intuitively, a temporal motif M can be represented as an ordered sequence of edges $\langle e'_1 = (u'_1, v'_1), \dots, e'_l = (u'_l, v'_l) \rangle$. Given a temporal motif M as a template pattern, we aim to count how many times this pattern appears in a temporal graph T . Furthermore, we only consider the instances where the pattern is formed within a short time span. For example, an instance formed in an hour is more interesting than one formed accidentally in one year on a communication network [9, 27, 50]. Therefore, given a temporal graph T and a temporal motif M , our goal is to find a sequence of edges $S \subseteq E_T$ such that (1) S exactly matches (i.e., is *isomorphic to*) M , (2) S is in the same order as specified by σ , and (3) all edges in S occur within a time span of at most δ . We call such an edge sequence S as a δ -instance [22, 27] of M and the difference between t_l and t_1 as the *duration* $\Delta(S)$ of instance S . The formal definition is given in the following.

Definition 3.2 (Motif δ -instance). A sequence of l edges $S = \langle (w_1, x_1, t_1), \dots, (w_l, x_l, t_l) \rangle$ ($t_1 < \dots < t_l$) from a temporal graph T is a δ -instance of a temporal motif $M = \langle (u'_1, v'_1), \dots, (u'_l, v'_l) \rangle$ if

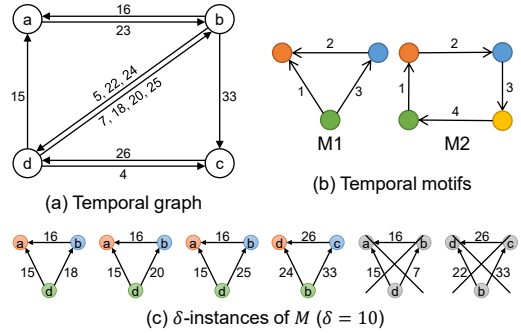


Figure 2: Example for temporal graph and motifs

(1) there exists a bijection f between the vertex sets of S and M such that $f(w_i) = u'_i$ and $f(x_i) = v'_i$ for $i = 1, \dots, l$; and (2) the duration $\Delta(S)$ is at most δ , i.e., $t_l - t_1 \leq \delta$.

Example 3.3. In Figure 2(a), we illustrate a temporal graph with 4 vertices and 13 temporal edges. Let us consider the problem of finding all δ -instances ($\delta = 10$) of temporal motif $M1$ in Figure 2(b). As shown in Figure 2(c), there are 4 valid 10-instances of $M1$ found. These instances can match $M1$ in terms of both structure and edge ordering and their durations are within 10. In addition, we also give 2 invalid instances of $M1$, which are isomorphic to $M1$ but violate either the edge ordering or duration constraint.

Temporal Motif Counting: According to the above notions, we present the *temporal motif counting* problem studied in this paper.

Definition 3.4 (Temporal Motif Counting). For a temporal graph T , a temporal motif M , and a time span δ , the temporal motif counting problem returns the number C_M of δ -instances of M appeared in T .

The temporal motif counting problem has proven to be NP-hard for very simple motifs, e.g. k -stars [22], because the edge ordering is taken into account. According to previous results [22], although there is a simple polynomial algorithm to count the number of k -stars on a static graph, it is NP-hard to exactly count the number of temporal k -stars. Typically, counting temporal motifs exactly on massive graphs with millions or even billions of edges is a computationally intensive task [22, 27]. Therefore, we focus on designing efficient and scalable sampling algorithms for estimating the number of temporal motifs approximately in Section 4. The frequently used notations are summarized in Table 1.

4 OUR ALGORITHMS

In this section, we present our proposed algorithms for approximate temporal motif counting in detail. We first describe our generic Edge Sampling (ES) algorithm in Section 4.1. Then, we introduce our improved EWS algorithm specific for counting 3-vertex 3-edge temporal motifs in Section 4.2. In addition, we theoretically analyze the expected values and variances of the estimates returned by both algorithms. Finally, we discuss the streaming implementation of our algorithms in Section 4.3.

4.1 The Generic Edge Sampling Algorithm

The Edge Sampling (ES) algorithm is motivated by an exact subgraph counting algorithm called *edge iterator* [49]. Given a temporal

Table 1: Frequently used notations

Symbol	Description
T	Temporal graph
V_T, E_T	Set of vertices and edges in T
n, m	Number of vertices and edges in T
M	Temporal motif
V_M, E_M	Set of vertices and edges in M
k, l	Number of vertices and edges in M
δ	Maximum time span of a motif instance
S	Motif δ -instance
C_M	Number of δ -instances of M in T
\widehat{C}_M	Unbiased estimator of C_M
p	Probability of edge sampling
\widehat{E}_T	Set of sampled edges from E_T
$\eta(e)$	Number of δ -instances of M containing edge e
$\eta_j(e)$	Number of δ -instances of M when e is mapped to e'_j
q	Probability of wedge sampling
W	Temporal wedge
$\eta(W)$	Number of δ -instances of M containing W
W'_j	Temporal wedge pattern for M when e is mapped to e'_j
$\widehat{W}'_j(e)$	Set of sampled δ -instances of W'_j
$\widehat{\eta}_j(e)$	Unbiased estimator of $\eta_j(e)$

graph T , a temporal motif M , and a time span δ , we use $\eta(e)$ to denote the number of local δ -instances of M containing an edge e . To count all δ -instances of M in T exactly, we can simply count $\eta(e)$ for each $e \in E_T$ and then sum them up. In this way, each instance is counted l times and the total number of instances is equal to the sum divided by l , i.e., $C_M = \frac{1}{l} \sum_{e \in E_T} \eta(e)$.

Based on the above idea, we propose the ES algorithm for estimating C_M : For each edge $e \in E_T$, we randomly sample it and compute $\eta(e)$ with fixed probability p . Then, we obtain an unbiased estimator \widehat{C}_M of C_M by adding up $\eta(e)$ for each sampled edge e and scaling the sum by a factor of $\frac{1}{pl}$, i.e., $\widehat{C}_M = \frac{1}{pl} \sum_{e \in \widehat{E}_T} \eta(e)$ where \widehat{E}_T is the set of sampled edges.

Now the remaining problem becomes how to compute $\eta(e)$ for an edge e . The ES algorithm adopts the well-known BACKTRACKING algorithm [23, 36] to enumerate all δ -instances that contain an edge e for computing $\eta(e)$. Specifically, the BACKTRACKING algorithm runs l times for each edge e ; in the j^{th} run, it first maps edge e to the j^{th} edge e'_j of M and then uses a tree search to find all different combinations of the remaining $l-1$ edges that can form δ -instances of M with edge e . Let $\eta_j(e)$ be the number of δ -instances of M where e is mapped to e'_j . It is obvious that $\eta(e)$ is equal to the sum of $\eta_j(e)$ for $j = 1, \dots, l$, i.e., $\eta(e) = \sum_{j=1}^l \eta_j(e)$.

We depict the procedure of our ES algorithm in Algorithm 1. The first step of ES is to generate a random sample \widehat{E}_T of edges from the edge set E_T where the probability of adding any edge is p (Lines 1–5). Then, in the second step (Lines 6–11), it counts the number $\eta(e)$ of local δ -instances of M for each sampled edge e by running the BACKTRACKING algorithm to enumerate each instance $S_j(e)$ that is a δ -instance of M and maps e to e'_j for $j = 1, \dots, l$. Note that BACKTRACKING (BT) runs on a subset $\widehat{E}_T[t - \delta, t + \delta]$ of E_T which consists of all edges with timestamps from $t - \delta$ to $t + \delta$ for edge $e = (u, v, t)$ since it is safe to ignore any other edge due to the duration constraint. Here, we omit the detailed procedure of the

Algorithm 1: Edge Sampling

Input: Temporal graph T , temporal motif M , time span δ , edge sampling probability p .
Output: Estimator \widehat{C}_M of the number of δ -instances of M in T

- 1 Initialize $\widehat{E}_T \leftarrow \emptyset$;
- 2 **foreach** $e \in E_T$ **do**
- 3 Toss a biased coin with success probability p ;
- 4 **if** *success* **then**
- 5 $\widehat{E}_T \leftarrow \widehat{E}_T \cup \{e\}$;
- 6 **foreach** $e = (u, v, t) \in \widehat{E}_T$ **do**
- 7 Set $\eta(e) \leftarrow 0$;
- 8 **for** $j \in 1, \dots, l$ **do**
- 9 Generate an initial instance $S_j^{(1)}$ by mapping e to e'_j ;
- 10 Run BACKTRACKING on $E_T[t - \delta, t + \delta]$ starting from $S_j^{(1)}$ to find the set $S_j(e) = \{S_j(e) : S_j(e) \text{ is a } \delta\text{-instance of } M \text{ where } e \text{ is mapped to } e'_j\}$;
- 11 Set $\eta_j(e) \leftarrow |S_j(e)|$ and $\eta(e) \leftarrow \eta(e) + \eta_j(e)$;
- 12 **return** $\widehat{C}_M \leftarrow \frac{1}{pl} \sum_{e \in \widehat{E}_T} \eta(e)$;

BT algorithm because it generally follows an existing algorithm for subgraph isomorphism in temporal graphs [23]. The main difference between our algorithm and the one in [23] lies in the matching order, which will be discussed later. After counting $\eta(e)$ for each sampled edge e , it finally returns an estimate \widehat{C}_M of C_M (Line 12).

Matching Order for BACKTRACKING: Now we discuss how to determine the matching order of a temporal motif. The BT algorithm in [23] adopts a time-first matching order: it always matches the edges of M in order of $\langle e'_1, \dots, e'_l \rangle$. The advantage of this matching order is that it best exploits the temporal information for search space pruning. For a partial instance $S^{(j)} = \langle (w_1, x_1, t_1), \dots, (w_j, x_j, t_j) \rangle$ after e'_j is mapped, the search space for mapping e'_{j+1} is restricted to $E_T[t_j, t_j + \delta]$. However, the time-first matching order may not work well in the ES algorithm. First, it does not consider the *connectivity* of the matching order: If e'_{j+1} is not connected with any prior edge, it has to be mapped to all edges in $E_T[t_j, t_j + \delta]$, which may lead to a large number of redundant partial matchings. Second, the time-first order is violated by Line 9 of Algorithm 1 when $j > 1$ since it first maps e to e'_j .

In order to overcome the above two drawbacks, we propose two heuristics to determine the matching order of a given motif M for reducing the search space, and generate l matching orders for M , in each of which e'_j ($j = 1, \dots, l$) is placed first: (1) *enforcing connectivity*: For each $i = 2, \dots, l$, the i^{th} edge in the matching order must be adjacent to at least one prior edge that has been matched; (2) *boundary edge first*: If there are multiple unmatched edges that satisfy the *connectivity* constraint, the boundary edge (i.e., the first or last unmatched edge in the ordering σ of M) will be matched first. The first rule can avoid redundant partial matchings and the second rule can restrict the temporal range of tree search, both of which are effective for search space pruning.

Example 4.1. We consider how to decide the matching orders of M_2 (i.e., 4-simple temporal cycle) in Figure 2(b). When e'_1 is placed first, we can select e'_2 or e'_4 as the second edge according to




Figure 3: Example of enumerating δ -instances ($\delta = 10$) of M_1 for edge $(b, a, 16)$ in Figure 2 using BACKTRACKING

the *enforcing connectivity* rule; and e'_4 is selected according to the *boundary edge first* rule. Then, either e'_2 or e'_3 can be selected as the next edge since they both satisfy two rules. Therefore, either $\langle e'_1, e'_4, e'_2, e'_3 \rangle$ or $\langle e'_1, e'_4, e'_3, e'_2 \rangle$ is a valid matching order. Accordingly, $\langle e'_2, e'_1, e'_4, e'_3 \rangle$, $\langle e'_3, e'_4, e'_1, e'_2 \rangle$, and $\langle e'_4, e'_1, e'_2, e'_3 \rangle$ are valid matching orders when e'_2, e'_3 , and e'_4 in M_2 are placed first, respectively.

Example 4.2. In Figure 3, we show how to use BACKTRACKING to enumerate δ -instances of M_1 ($\delta = 10$) for $e = (b, a, 16)$ in Figure 2. There are three tree search procedures in each of which e is mapped to e'_1, e'_2 , and e'_3 , respectively. The condition of each mapping step is given in form of $(v_s, v_t, [t_s, t_t])$ where v_s and v_t are the starting and ending vertices and $[t_s, t_t]$ is the range of timestamps. Here, “*” means that it can be mapped to an arbitrary unmapped vertex. Moreover, we use ‘✓’ and ‘X’ to denote a successful matching and a failed partial matching, respectively. We find three 10-instances of M_1 and thus $\eta(e) = 3$. When we run ES with $p = 0.25$ and $\widehat{E}_T = \{(d, c, 4), (b, a, 16), (b, d, 24)\}$, since the numbers of 10-instances containing $(d, c, 4)$ and $(b, d, 24)$ are respectively 0 and 1, we can compute $\widehat{C}_M = \frac{3+0+1}{0.25 \times 3} \approx 5.33$.

Theoretical Analysis: Next, we analyze the estimate \widehat{C}_M returned by Algorithm 1 theoretically.

We first prove that \widehat{C}_M is an unbiased estimator of C_M in Theorem 4.3. The variance of \widehat{C}_M is given in Theorem 4.4. The proofs of Theorems 4.3–4.5 are provided in the extended version [41].

THEOREM 4.3. *The expected value $\mathbb{E}[\widehat{C}_M]$ of \widehat{C}_M returned by Algorithm 1 is C_M .*

THEOREM 4.4. *The variance $\text{Val}[\widehat{C}_M]$ of \widehat{C}_M returned by Algorithm 1 is at most $\frac{1-p}{p} \cdot C_M^2$.*

Finally, we can derive Theorem 4.5 by applying Chebyshev’s inequality to the result of Theorem 4.4.

THEOREM 4.5. $\Pr[|\widehat{C}_M - C_M| \geq \varepsilon \cdot C_M] \leq \frac{1-p}{p\varepsilon^2}$

According to Theorem 4.5, we can say \widehat{C}_M is an (ε, γ) -estimator of C_M for parameters $\varepsilon, \gamma \in (0, 1)$, i.e., $\Pr[|\widehat{C}_M - C_M| < \varepsilon \cdot C_M] > 1 - \gamma$, when $p = \frac{1}{1+\gamma\varepsilon^2}$.

Time Complexity: We first analyze the time complexity of computing $\eta(e)$ for an edge e . For BACKTRACKING, the search space of each matching step is at most the number of (in-/out-)edges within range $[t - \delta, t]$ or $[t, t + \delta]$ connected with a vertex v . Here, we use d_δ to denote the maximum number of (in-/out-)edges connected with one vertex within any δ -length time interval. The time complexity




Figure 4: Temporal wedges, (3,3)-stars, and triangles

of BACKTRACKING is $O(d_\delta^{l-1})$ and thus the time complexity of computing $\eta(e)$ is $O(ld_\delta^{l-1})$. Therefore, ES provides an (ε, γ) -estimator of C_M in $O(\frac{mld_\delta^{l-1}}{1+\gamma\varepsilon^2})$ time.

4.2 The Improved EWS Algorithm

The ES algorithm in Section 4.1 is generic and able to count any connected temporal motif. Nevertheless, there are still opportunities to further reduce the computational overhead of ES when the query motif is limited to 3-vertex 3-edge temporal motifs (i.e., triadic patterns), which are one of the most important classes of motifs to characterize temporal networks [5, 15, 27, 37]. In this section, we propose an improved Edge-Wedge Sampling (EWS) algorithm that combines *edge sampling* with *wedge sampling* for counting 3-vertex 3-edge temporal motifs.

Wedge sampling [15, 34, 35, 49] is a widely used method for triangle counting. Its basic idea is to draw a sample of wedges (i.e., 3-vertex 2-edge subgraph patterns) uniformly from a graph and check the ratio of “closed wedges” (i.e., form a triangle in the graph) to estimate the number of triangles. However, traditional wedge-sampling methods are proposed for undirected static graphs and cannot be directly used on temporal graphs. First, they consider that all wedges are isomorphic and treat them equally. But there are four temporal wedge patterns with different edge directions and orderings as illustrated in Figure 4(a). Second, they are designed for simple graphs where one wedge can form at most one triangle. However, since temporal graphs are multigraphs and there may exist multiple edges between the same two vertices, one temporal wedge can participate in more than one instance of a temporal motif. Therefore, in the EWS algorithm, we extend *wedge sampling* for temporal motif counting by addressing both issues.

The detailed procedure of EWS is presented in Algorithm 2. First of all, it uses the same method as ES to sample a set \widehat{E}_T of edges (Line 1). For each sampled edge $e \in \widehat{E}_T$ and $j = 1, 2, 3$, it also maps e to e'_j for computing $\eta_j(e)$ (Line 4), i.e., the number of δ -instances of M where e is mapped to e'_j . But, instead of running BACKTRACKING to compute $\eta_j(e)$ exactly, it utilizes *temporal wedge sampling* to estimate $\eta_j(e)$ approximately without full enumeration (Lines 5–15), which is divided into two subroutines as discussed later. At last, it obtains an estimate \widehat{C}_M of C_M from each estimate $\widehat{\eta}_j(e)$ of $\eta_j(e)$ using a similar method to ES (Line 16).

Sample Temporal Wedges (Lines 5–12): The first step of *temporal wedge sampling* is to determine which temporal wedge pattern is to be matched according to the query motif M and the mapping

Algorithm 2: Edge-Wedge Sampling

Input: Temporal graph T , temporal motif M , time span δ , edge sampling probability p , wedge sampling probability q .
Output: Estimator \widehat{C}_M of the number of δ -instances of M in T

- 1 Generate \widehat{E}_T using Line 1–5 of Algorithm 1;
- 2 **foreach** $e = (u, v, t) \in \widehat{E}_T$ **do**
- 3 **for** $j \leftarrow 1, 2, 3$ **do**
- 4 Map edge e to e'_j ;
- 5 Initialize $\widehat{\eta}_j(e) \leftarrow 0$ and $\widehat{\mathcal{W}}_j(e) \leftarrow \emptyset$;
- 6 **if** M is a temporal $(3, 3)$ -star **then**
- 7 Select W'_j including e'_j centered at the center of M ;
- 8 **else if** M is a temporal triangle **then**
- 9 Select W'_j including e'_j centered at the vertex mapped to the one with a lower degree in u and v ;
- 10 $E_j(e) \leftarrow$ all edges that form δ -instances of W'_j with e ;
- 11 **foreach** $g \in E_j(e)$ **do**
- 12 Add a δ -instance W of W'_j comprising e and g to $\widehat{\mathcal{W}}_j(e)$ with probability q ;
- 13 **foreach** $W \in \widehat{\mathcal{W}}_j(e)$ **do**
- 14 Let $\eta(W)$ be the number of edges that form δ -instances of M together with W ;
- 15 $\widehat{\eta}_j(e) \leftarrow \widehat{\eta}_j(e) + \frac{\eta(W)}{q}$;
- 16 **return** $\widehat{C}_M \leftarrow \frac{1}{3p} \sum_{e \in \widehat{E}_T} \sum_{j=1}^3 \widehat{\eta}_j(e)$;

from e to e'_j . Specifically, we categorize 3-vertex 3-edge temporal motifs into two types, i.e., *temporal (3, 3)-stars* and *temporal triangles* as shown in Figure 4, based on whether they are closed. Interested readers may refer to [27] for a full list of all 3-vertex 3-edge temporal motifs. For a star or wedge pattern, the vertex connected with all edges is its *center*. Given that $e = (u, v, t)$ has been mapped to e'_j , EWS should find a temporal wedge pattern W'_j containing e'_j from M for sampling. Here, different strategies are adopted to determine W'_j for star and triangle motifs (Lines 6–9): If M is a temporal $(3, 3)$ -star, it must select W'_j that contains e'_j and has the same center as M ; If M is a temporal triangle, it may use the vertex mapped to either u or v as the center to generate a wedge pattern. In this case, the center of W'_j will be mapped to the vertex with a lower degree between u and v for search space reduction. After deciding W'_j , it enumerates all edges that form a δ -instance of W'_j together with e as $E_j(e)$ from the adjacency list of the central vertex (Line 10). By selecting each edge $g \in E_j(e)$ with probability q , it generates a sample $\widehat{\mathcal{W}}_j(e)$ of δ -instances of W'_j (Lines 11–12).

Estimate $\eta_j(e)$ (Lines 13–15): Now, it estimates $\eta_j(e)$ from the set $\widehat{\mathcal{W}}_j(e)$ of sampled temporal wedges. For each $W \in \widehat{\mathcal{W}}_j(e)$, it counts the number $\eta(W)$ of δ -instances of M that contain W (Line 14). Specifically, after matching W with W'_j , it can determine the starting and ending vertices as well as the temporal range for the mapping of the third edge of M . For the fast computation of $\eta(W)$, EWS maintains a hash table that uses an ordered combination $\langle u, v \rangle$ ($u, v \in V_T$) as the key and a sorted list of the timestamps of all edges from u to v as the value on the edge set E_T of T . In this way, $\eta(W)$ can be computed by a hash search followed by at




Figure 5: Examples for the EWS algorithm

most two binary searches on the sorted list. Finally, $\eta_j(e)$ can be estimated by summing up $\eta(W)$ for each $W \in \widehat{\mathcal{W}}_j(e)$ (Line 15), i.e., $\widehat{\eta}_j(e) = \frac{1}{q} \sum_{W \in \widehat{\mathcal{W}}_j(e)} \eta(W)$.

Example 4.6. In Figure 5, we show how to compute $\widehat{\eta}_j(e)$ using EWS on the temporal graph in Figure 2. In this example, q is set to 1, i.e., all temporal wedges found are sampled. When $(a, b, 23)$ is mapped to e'_3 of $M11$ in Figure 4, we have $W'_3 = W2$ and find 2 instances $W21$ and $W22$ of $W2$. Then, we get $\eta(W21) = 1$ and $\eta(W22) = 0$ and thus $\widehat{\eta}_3((a, b, 23)) = 1$. For an edge $(c, d, 26)$ mapped to e'_2 of $M21$ in Figure 4, c is used as the central vertex since $\deg(c) = 3 < \deg(d) = 5$. Then, we have $W'_2 = W4$ and there is only one instance $W41$ of $W4$ found. As $\eta(W41) = 1$, we get $\widehat{\eta}_2((c, d, 26)) = 1$ accordingly.

Theoretical Analysis: Next, we analyze the estimate \widehat{C}_M returned by Algorithm 2 theoretically. We prove the unbiasedness and variances of \widehat{C}_M in Theorem 4.7 and Theorem 4.8, respectively. Detailed proofs are also provided in the extended version [41].

THEOREM 4.7. *The expected value $\mathbb{E}[\widehat{C}_M]$ of \widehat{C}_M returned by Algorithm 2 is C_M .*

THEOREM 4.8. *The variance $\text{Val}[\widehat{C}_M]$ of \widehat{C}_M returned by Algorithm 2 is at most $\frac{1-pq}{pq} \cdot C_M^2$.*

According to the result of Theorem 4.8 and Chebyshev's inequality, we have $\Pr[|\widehat{C}_M - C_M| \geq \varepsilon \cdot C_M] \leq \frac{1-pq}{pq\varepsilon^2}$ and \widehat{C}_M is an (ε, γ) -estimator of C_M for parameters $\varepsilon, \gamma \in (0, 1)$ when $pq = \frac{1}{1+\gamma\varepsilon^2}$.

Time Complexity: We first analyze the time to compute $\widehat{\eta}_j(e)$. First, $|E_j(e)|$ is bounded by the maximum number of (in-/out-)edges connected with one vertex within any δ -length time interval, i.e., d_δ . Second, the time to compute $\eta(W)$ using a hash table is $O(\log h)$ where h is the maximum number of edges between any two vertices. Therefore, the time complexity per edge in EWS is $O(d_\delta \log h)$. This is lower than $O(d_\delta^2)$ time per edge in ES ($k = l = 3$). Finally, EWS provides an (ε, γ) -estimator of C_M in $O(\frac{md_\delta \log h}{1+\gamma\varepsilon^2})$ time.

4.3 Streaming Implementation

To deal with a dataset that is too large to fit in memory or generated in a streaming manner, it is possible to adapt our algorithms to a streaming setting. Assuming that all edges are sorted in chronological order, our algorithms can determine whether to sample an edge or not when it arrives. Then, for each sampled edge $e = (u, v, t)$, we only need the edges with timestamps in $[t - \delta, t + \delta]$ to compute its local count $\eta(e)$ or $\widehat{\eta}(e)$. After a one-pass scan over the temporal graph stream, we can obtain an estimate of the number of a temporal motif in the stream. Generally, our algorithms can process

Table 2: Statistics of datasets

Dataset	$ V_T $	$ E $	$ E_T $	Time span
AU	157, 222	544, 621	726, 639	7.16 years
SU	192, 409	854, 377	1, 108, 716	7.60 years
SO	2, 584, 164	34, 875, 684	47, 902, 865	7.60 years
BC	48, 098, 591	86, 798, 226	113, 100, 979	7.08 years
RC	5, 688, 164	329, 485, 956	399, 523, 749	7.44 years

any temporal graph stream in one pass by always maintaining the edges in the most recent time interval of length 2δ while having the same theoretical bounds as in the batch setting.

5 EXPERIMENTAL EVALUATION

In this section, we evaluate the empirical performance of our proposed algorithms on real-world datasets. We first introduce the experimental setup in Section 5.1. The experimental results are presented in Section 5.2.

5.1 Experimental Setup

Experimental Environment: All experiments were conducted on a server running Ubuntu 18.04.1 LTS with an Intel[®] Xeon[®] Gold 6140 2.30GHz processor and 250GB RAM. All datasets and our code are publicly available¹. We downloaded the code^{2,3,4} of baselines published by the authors and followed the compilation and usage instructions. All algorithms were implemented in C++11 compiled by GCC v7.4 with -O3 optimizations, and ran on a single thread.

Datasets: We used five different real-world datasets in our experiments including AskUbuntu (AU), SuperUser (SU), StackOverflow (SO), BitCoin (BC), and RedditComments (RC). All datasets were downloaded from publicly available sources like the SNAP repository [19]. Each dataset is a sequence of temporal edges in chronological order. We report the statistics of these datasets in Table 2, where $|V_T|$ is the number of vertices, $|E|$ is the number of (static) edges, $|E_T|$ is the number of temporal edges, and *time span* is the overall time span of the entire dataset.

Algorithms: The algorithms compared are listed as follows.


- **EX:** An exact algorithm for temporal motif counting in [27]. The available implementation is applicable only to 3-edge motifs and cannot support motifs with 4 or more edges (e.g., Q5 in Figure 6).
- **2SCENT:** An algorithm for simple temporal cycle (e.g., Q4 and Q5 in Figure 6) enumeration in [18].
- **BT:** A BACKTRACKING algorithm for temporal subgraph isomorphism in [23]. It provides the exact count of any temporal motif by enumerating all of them.
- **IS-BT:** An interval-based sampling algorithm for temporal motif counting in [22]. BT [23] is used as a subroutine for any motif with more than 2 vertices.
- **ES:** Our generic edge sampling algorithm for temporal motif counting in Section 4.1.

¹<https://github.com/jingjing-hnu/Temporal-Motif-Counting>

²<http://snap.stanford.edu/temporal-motifs/>

³<https://github.com/rohit13k/CycleDetection>

⁴<https://gitlab.com/paul.liu.ubc/sampling-temporal-motifs>

**Figure 6: Query motifs**

- **EWS:** Our improved algorithm that combines edge sampling with wedge sampling for counting temporal motifs with 3 vertices and 3 edges (e.g. Q1–Q4 in Figure 6) in Section 4.2.

Queries: The five query motifs we use in the experiments are listed in Figure 6. Since different algorithms specialize in different types of motifs, we select a few motifs that can best represent the specializations of all algorithms. As discussed above, an algorithm may not be applicable to some of the motifs. In this case, the algorithm is ignored in the experiments on these motifs.

Performance Measures: The efficiency is measured by the CPU time (in seconds) of an algorithm to count a query motif in a temporal graph. The accuracy of a sampling algorithm is measured by the relative error $\frac{|\hat{x}-x|}{x}$ where x is the exact number of instances of a query motif in a temporal graph and \hat{x} is an estimate of x returned by an algorithm. In each experiment, we run all algorithms 10 times and use the average CPU time and relative errors for comparison.

5.2 Experimental Results

The overall performance of each algorithm is reported in Table 3. Here, the time span δ is set to 86400 seconds (i.e., one day) on AU and SU, and 3600 seconds (i.e., one hour) on SO, BC, and RC (Note that we use the same values of δ across all experiments, unless specified). For IS-BT, we report the results in the default setting as indicated in [22], i.e., we fix the interval length to 30δ and present the result for the smallest interval sampling probability that can guarantee the relative error is at most 5%. For ES and EWS, we report the results when $p = 0.01$ by default; in a few cases when the numbers of motif instances are too small or their distribution is highly skewed among edges, we report the results when $p = 0.1$ (marked with “**” in Table 3) because ES and EWS cannot provide accurate estimates when $p = 0.01$. In addition, we set q to 1 on AU and SU, and 0.1 on SO, BC, and RC for EWS.

First of all, the efficiencies of EX and 2SCENT are lower than the other algorithms. This is because they use an algorithm for subgraph isomorphism or cycle detection in static graphs for candidate generation without considering temporal information. As a result, a large number of redundant candidates are generated and lead to the degradation in performance. Second, on medium-sized datasets (i.e., AU and SU), ES runs faster than IS-BT in most cases; and meanwhile, their relative errors are close to each other. On large datasets (i.e., SO, BC, and RC), ES demonstrates both much higher efficiency (up to 10.3x speedup) and lower estimation errors (2.42% vs. 4.61%) than IS-BT. Third, EWS runs 1.7x–19.6x faster than ES due to its lower computational cost per edge. The relative errors of ES and EWS are the same on AU and SU because $q = 1$. When $q = 0.1$, EWS achieves further speedups at the expense of higher relative errors. A more detailed analysis of the effect of q is provided in the following paragraph.

Table 3: Running time (in seconds) and average errors (%) of all algorithms on each dataset. We use “—” and “×” to denote “motif not supported” and “running out of memory”, respectively. For IS-BT, ES, and EWS, we show their speedup ratios over BT for comparison. We use “*” to mark the results of ES and EWS for $p = 0.1$ instead of $p = 0.01$.

Dataset	Motif	EX	2SCENT	BT	IS-BT		ES		EWS	
		time (s)	time (s)	time (s)	error	time (s)	error	time (s)	error	time (s)
AU	Q1	1.8	—	0.758	4.84%	0.402/1.9x	4.32%	0.059/12.8x	4.32%	0.027/28.1x
	Q2			1.104	4.16%	0.434/2.5x	4.57%	0.048/23.0x	4.57%	0.029/38.1x
	Q3	2.3		0.884	3.97%	0.50/1.8x	*3.73%	*0.605/1.5x	*3.73%	*0.183/4.8x
	Q4			1.038	4.67%	0.492/2.1x	*4.63%	*0.628/1.7x	*4.63%	*0.173/6x
	Q5	—		23.68	1.262	3.98%	0.536/2.4x	*4.62%	*0.322/3.9x	—
SU	Q1	3.26	—	1.499	3.99%	0.620/2.4x	3.06%	0.102/14.7x	3.06%	0.052/28.8x
	Q2			1.650	3.23%	0.671/2.5x	2.47%	0.083/19.9x	2.47%	0.046/35.9x
	Q3	4.6		1.506	4.85%	0.723/2.1x	4.66%	0.113/13.3x	4.66%	0.030/50.2x
	Q4			1.434	3.79%	0.725/2.0x	4.63%	0.128/11.2x	4.63%	0.042/34.1x
	Q5	—		46.0	1.521	4.55%	0.759/2.0x	*4.52%	*0.453/3.4x	—
SO	Q1	169	—	105.8	4.82%	8.626/12.3x	0.97%	4.419/23.9x	1.22%	1.528/69.2x
	Q2			110.7	4.82%	27.48/4.0x	0.20%	3.985/27.8x	0.89%	1.514/73.1x
	Q3	466		107.4	4.30%	25.70/4.2x	1.36%	4.031/26.6x	3.6%	1.235/87x
	Q4			105.5	4.90%	6.775/15.6x	1.78%	3.936/26.8x	3.31%	1.153/91.5x
	Q5	—		243.7	91.83	4.91%	9.451/9.7x	3.48%	1.505/61.0x	—
BC	Q1	8143	—	220.0	4.75%	50.02/4.4x	0.64%	59.12/3.7x	0.67%	9.463/23.2x
	Q2			399.8	4.90%	125.1/3.2x	1.11%	34.74/11.5x	1.16%	8.126/49.2x
	Q3	8116		396.8	3.89%	90.19/4.4x	1.49%	41.49/9.6x	3.02%	2.121/187x
	Q4			473.4	4.93%	95.47/5.0x	0.83%	37.43/12.6x	1.91%	2.262/209x
	Q5	—		473.7	596.4	4.83%	319.7/1.9x	2.92%	20.47/29.1x	—
RC	Q1	2799	—	1966	4.76%	840.5/2.3x	3.27%	257.4/7.6x	3.36%	31.49/62.4x
	Q2			2113	4.67%	428/4.9x	0.63%	120.6/17.5x	0.6%	30.57/69.1x
	Q3	×		2069	4.61%	784.4/2.6x	2.42%	76.09/27.2x	2.27%	16.17/128x
	Q4			1897	4.86%	683/2.8x	3.47%	68.60/27.7x	4.57%	15.91/119x
	Q5	—		2245	1613	4.41%	706.6/2.3x	*4.32%	*120.3/13.4x	—




Figure 7: Comparison of the performance of EWS when $q = 0.1$ and 1

Effect of q for EWS: In Figure 7, we compare the relative errors and running time of EWS for $q = 1$ and 0.1 when p is fixed to 0.01 . We observe different effects of q on medium-sized (e.g., SU) and large (e.g., BC) datasets. On the SU dataset, the benefit of smaller q is marginal: the running time decreases slightly but the errors become obviously higher. But on the BC dataset, by setting $q = 0.1$, EWS achieves 2x–3x speedups without affecting the accuracy seriously. These results imply that *temporal wedge sampling* is more effective on larger datasets. Therefore, we set $q = 1$ on AU and SU, and $q = 0.1$ on SO, BC, and RC for EWS in the remaining experiments.

Accuracy vs. Efficiency: Figure 8 demonstrates the trade-offs between *relative error* and *running time* of three sampling algorithms, namely IS-BT, ES, and EWS. For IS-BT, we fix the interval length to 30δ and vary the interval sampling probability from 0.01 to 1 . For ES and EWS, we vary the edge sampling probability p from 0.0001 to 0.25 . First of all, ES and EWS consistently achieve better trade-offs between accuracy and efficiency than IS-BT in almost

all experiments. Specifically, ES and EWS can run up to 60x and 330x faster than IS-BT when the relative errors are at the same level. Meanwhile, in the same elapsed time, ES and EWS are up to 10.4x and 16.5x more accurate than IS-BT, respectively. Furthermore, EWS can outperform ES in all datasets except SO because of lower computational overhead. But on the SO dataset, since the distribution of motif instances is highly skewed among edges and thus the *temporal wedge sampling* leads to large errors in estimation, the performance of EWS degrades significantly and is close to or even worse than that of ES. Nevertheless, the effectiveness of *temporal wedge sampling* for EWS can still be confirmed by the results on the BC and RC datasets.

Scalability: We evaluate the scalability of different algorithms with varying the time span δ and dataset size m . In both experiments, we use the same parameter settings as used for the same motif on the same dataset in Table 3. We test the effect of δ for Q3 on the BC dataset by varying δ from 1h to 24h. As shown in Figure 9(a), the

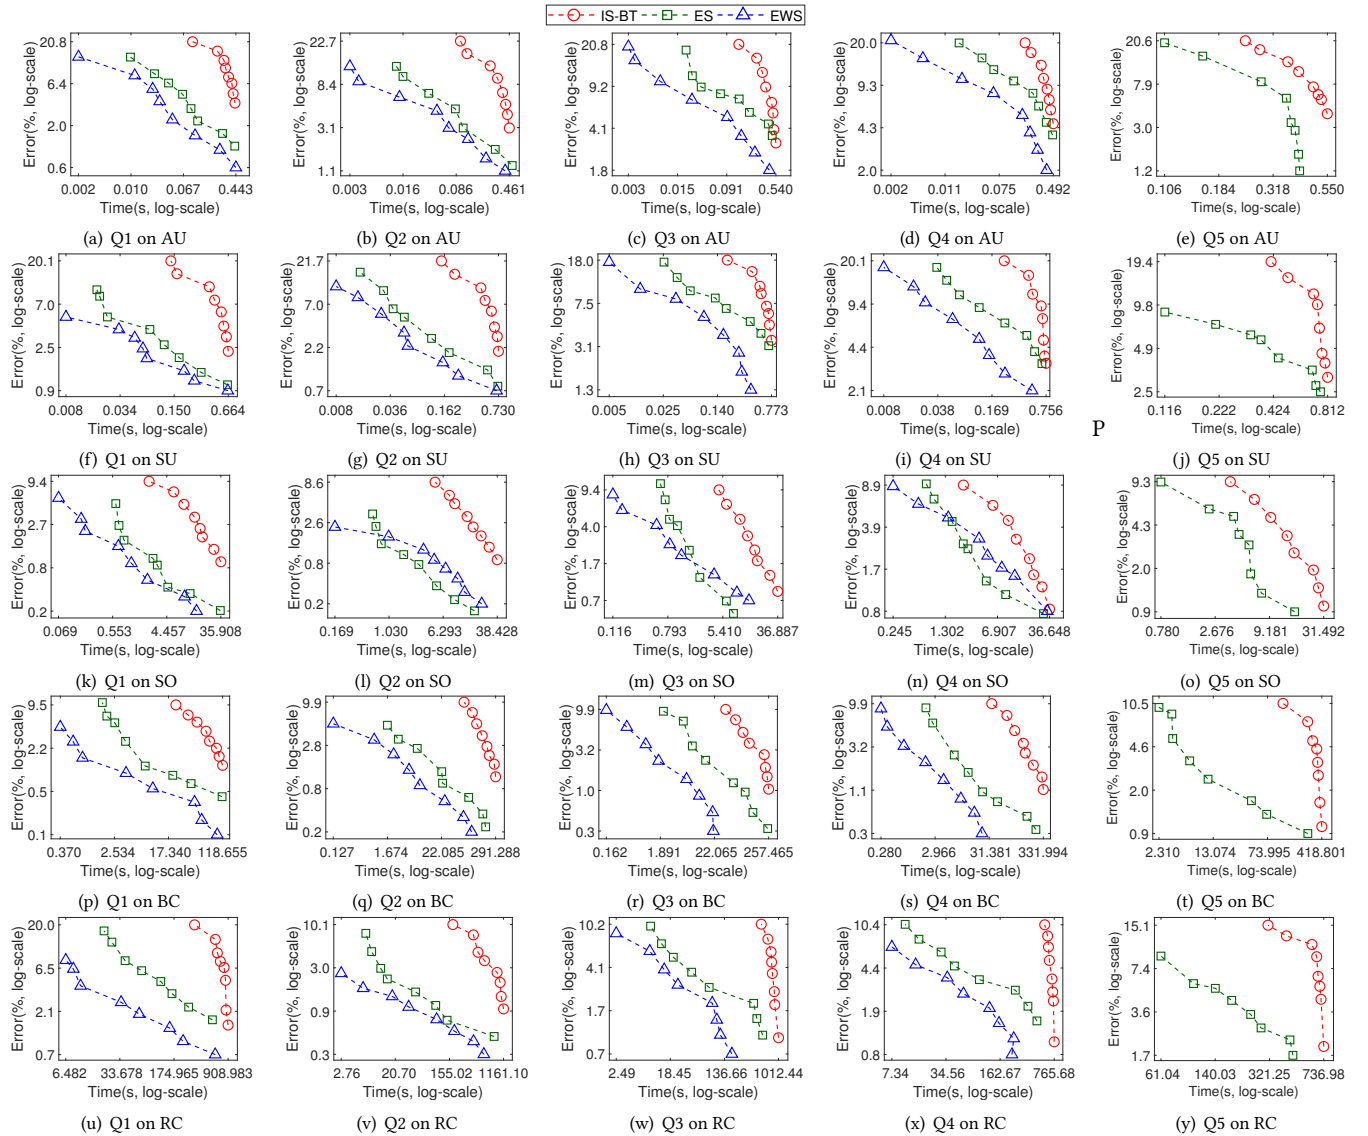


Figure 8: Relative error (%) vs. running time (seconds) with varying sampling probability

running time of all algorithms increases near-linearly w.r.t. δ . BT runs out of memory when $\delta > 10h$. The relative errors of ES and EWS keep steady for different δ but the accuracy of IS-BT degrades seriously when δ increases. This is owing to the increase in cross-interval instances and the skewness of instances among intervals. Meanwhile, ES and EWS run up to 2.2x and 180x faster than IS-BT, respectively, while always having smaller errors. The results for Q2 on the RC dataset with varying m are presented in Figure 9(b). Here, we vary m from 50M to near 400M by extracting the first m temporal edges of the RC dataset. The running time of all algorithms grows near-linearly w.r.t. m . The fluctuations of relative errors of IS-BT explicate that it is sensitive to the skewness of instances among intervals. ES and EWS always significantly outperform IS-BT for different m : they run much faster, have smaller relative errors, and provide more stable estimates than IS-BT.

6 CONCLUSION

In this paper, we studied the problem of approximately counting a temporal motif in a temporal graph via random sampling. We first proposed a generic Edge Sampling (ES) algorithm to estimate the number of any k -vertex l -edge temporal motif in a temporal graph. Furthermore, we improved the ES algorithm by combining edge sampling with wedge sampling and devised the EWS algorithm for counting 3-vertex 3-edge temporal motifs. We provided comprehensive theoretical analyses on the unbiasedness, variances, and complexities of our algorithms. Extensive experiments on several real-world temporal graphs demonstrated the accuracy, efficiency, and scalability of our algorithms. Specifically, ES and EWS ran up to 10.3x and 48.5x faster than the state-of-the-art sampling method while having lower estimation errors.




Figure 9: Scalability tests with varying time span δ and number of temporal edges m

ACKNOWLEDGMENTS

This research was partially supported by the China Scholarships Council grant 201906130131, the NSFC grant 61632009, and the grant MOE2017-T2-1-141 from the Singapore Ministry of Education.

REFERENCES

- [1] Nesreen K. Ahmed, Nick G. Duffield, Jennifer Neville, and Ramana Rao Kompella. 2014. Graph sample and hold: a framework for big-graph analytics. In *KDD*. 1446–1455.
- [2] Marco Bressan, Flavio Chierichetti, Ravi Kumar, Stefano Leucci, and Alessandro Panconesi. 2017. Counting Graphlets: Space vs Time. In *WSDM*. 557–566.
- [3] Lorenzo De Stefani, Alessandro Epasto, Matteo Riondato, and Eli Upfal. 2016. TRIEST: Counting Local and Global Triangles in Fully-Dynamic Streams with Fixed Memory Size. In *KDD*. 825–834.
- [4] Roohollah Etemadi, Jianguo Lu, and Yung H. Tsin. 2016. Efficient Estimation of Triangles in Very Large Graphs. In *CIKM*. 1251–1260.
- [5] Katherine Faust. 2010. A puzzle concerning triads in social networks: Graph constraints and the triad census. *Soc. Netw.* 32, 3 (2010), 221–233.
- [6] Edoardo Galimberti, Alain Barrat, Francesco Bonchi, Ciro Cattuto, and Francesco Gullo. 2018. Mining (maximal) Span-cores from Temporal Networks. In *CIKM*. 107–116.
- [7] Wentian Guo, Yuchen Li, Mo Sha, Bingsheng He, Xiaokui Xiao, and Kian-Lee Tan. 2020. GPU-Accelerated Subgraph Enumeration on Partitioned Graphs. In *SIGMOD*. 1067–1082.
- [8] Wentian Guo, Yuchen Li, Mo Sha, and Kian-Lee Tan. 2017. Parallel Personalized PageRank on Dynamic Graphs. *PVLDB* 11, 1 (2017), 93–106.
- [9] Saket Gururkar, Sayan Ranu, and Balaraman Ravindran. 2015. COMMIT: A Scalable Approach to Mining Communication Motifs from Dynamic Networks. In *SIGMOD*. 475–489.
- [10] Anne-Sophie Himmel, Hendrik Molter, Rolf Niedermeier, and Manuel Sorge. 2016. Enumerating maximal cliques in temporal graphs. In *ASONAM*. 337–344.
- [11] Petter Holme and Jari Saramäki. 2012. Temporal networks. *Phys. Rep.* 519, 3 (2012), 97–125.
- [12] Shweta Jain and C. Seshadhri. 2017. A Fast and Provable Method for Estimating Clique Counts Using Turán’s Theorem. In *WWW*. 441–449.
- [13] Madhav Jha, C. Seshadhri, and Ali Pinar. 2013. A space efficient streaming algorithm for triangle counting using the birthday paradox. In *KDD*. 589–597.
- [14] Madhav Jha, C. Seshadhri, and Ali Pinar. 2015. Path Sampling: A Fast and Provable Method for Estimating 4-Vertex Subgraph Counts. In *WWW*. 495–505.
- [15] Tamara G. Kolda, Ali Pinar, and C. Seshadhri. 2013. Triadic Measures on Graphs: The Power of Wedge Sampling. In *SDM*. 10–18.
- [16] Chrysanthis Kosyfaki, Nikos Mamoulis, Evangelia Pitoura, and Panayiotis Tsaparas. 2019. Flow Motifs in Interaction Networks. In *EDBT*. 241–252.
- [17] Lauri Kovonen, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. 2011. Temporal motifs in time-dependent networks. *J. Stat. Mech.: Theory Exp.* 2011, 11 (2011), P11005.
- [18] Rohit Kumar and Toon Calders. 2018. 2SCENT: An Efficient Algorithm to Enumerate All Simple Temporal Cycles. *PVLDB* 11, 11 (2018), 1441–1453.
- [19] Jure Leskovec and Rok Sosič. 2016. SNAP: A General-Purpose Network Analysis and Graph-Mining Library. *TIST* 8, 1 (2016), 1:1–1:20.
- [20] Rong-Hua Li, Jiao Su, Lu Qin, Jeffrey Xu Yu, and Qiangqiang Dai. 2018. Persistent Community Search in Temporal Networks. In *ICDE*. 797–808.
- [21] Yongsul Lim and U Kang. 2015. MASCO: Memory-efficient and Accurate Sampling for Counting Local Triangles in Graph Streams. In *KDD*. 685–694.
- [22] Paul Liu, Austin R. Benson, and Moses Charikar. 2019. Sampling Methods for Counting Temporal Motifs. In *WSDM*. 294–302.
- [23] Patrick Mackey, Katherine Porterfield, Erin Fitzhenry, Sutanay Choudhury, and George Chin Jr. 2018. A Chronological Edge-Driven Approach to Temporal Subgraph Isomorphism. In *BigData*. 3972–3979.
- [24] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network Motifs: Simple Building Blocks of Complex Networks. *Science* 298, 5594 (2002), 824–827.
- [25] Mohammad Hossein Namaki, Yinghui Wu, Qi Song, Peng Lin, and Tingjian Ge. 2017. Discovering Graph Temporal Association Rules. In *CIKM*. 1697–1706.
- [26] Rasmus Pagh and Charalambos E. Tsourakakis. 2012. Colorful triangle counting and a MapReduce implementation. *Inf. Process. Lett.* 112, 7 (2012), 277–281.
- [27] Ashwin Paranjape, Austin R. Benson, and Jure Leskovec. 2017. Motifs in Temporal Networks. In *WSDM*. 601–610.
- [28] A. Pavan, Kanat Tangwongsan, Srikanta Tirthapura, and Kun-Lung Wu. 2013. Counting and Sampling Triangles from a Graph Stream. *PVLDB* 6, 14 (2013), 1870–1881.
- [29] Ali Pinar, C. Seshadhri, and Vaidyanathan Vishal. 2017. ESCAPE: Efficiently Counting All 5-Vertex Subgraphs. In *WWW*. 1431–1440.
- [30] Natasa Przulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), 177–183.
- [31] Seyed-Vahid Sanei-Mehri, Ahmet Erdem Sariyüce, and Srikanta Tirthapura. 2018. Butterfly Counting in Bipartite Networks. In *KDD*. 2150–2159.
- [32] Mo Sha, Yuchen Li, Bingsheng He, and Kian-Lee Tan. 2017. Accelerating Dynamic Graph Analytics on GPUs. *PVLDB* 11, 1 (2017), 107–120.
- [33] Charalambos E. Tsourakakis, U Kang, Gary L. Miller, and Christos Faloutsos. 2009. DOULION: counting triangles in massive graphs with a coin. In *KDD*. 837–846.
- [34] Ata Turk and Duru Türkoglu. 2019. Revisiting Wedge Sampling for Triangle Counting. In *WWW*. 1875–1885.
- [35] Duru Türkoglu and Ata Turk. 2017. Edge-Based Wedge Sampling to Estimate Triangle Counts in Very Large Graphs. In *ICDM*. 455–464.
- [36] Julian R. Ullmann. 1976. An Algorithm for Subgraph Isomorphism. *J. ACM* 23, 1 (1976), 31–42.
- [37] Ruta Uzupyte and Ernst C. Wit. 2020. Test for triadic closure and triadic protection in temporal relational event data. *Social Netw. Anal. Mining* 10, Article 21 (2020).
- [38] Lav R. Varshney, Beth L. Chen, Eric Paniagua, David H. Hall, and Dmitri B. Chklovskii. 2011. Structural Properties of the *Caenorhabditis elegans* Neuronal Network. *PLoS Comput. Biol.* 7, 2 (2011), 1–21.
- [39] Jordan Viard, Matthieu Latapy, and Clémence Magnien. 2015. Revealing contact patterns among high-school students using maximal cliques in link streams. In *ASONAM*. 1517–1522.
- [40] Tiphaine Viard, Matthieu Latapy, and Clémence Magnien. 2016. Computing maximal cliques in link streams. *Theor. Comput. Sci.* 609 (2016), 245–252.
- [41] Jingjing Wang, Yanhao Wang, Wenjun Jiang, Yuchen Li, and Kian-Lee Tan. 2020. Efficient Sampling Algorithms for Approximate Temporal Motif Counting (Extended Version). arXiv:2007.14028 [cs.LG].
- [42] Pinghui Wang, John C. S. Lui, Bruno F. Ribeiro, Don Towsley, Junzhou Zhao, and Xiaohong Guan. 2014. Efficiently Estimating Motif Statistics of Large Networks. *TKDD* 9, 2 (2014), 8:1–8:27.
- [43] Pinghui Wang, John C. S. Lui, Donald F. Towsley, and Junzhou Zhao. 2016. Minfer: A method of inferring motif statistics from sampled edges. In *ICDE*. 1050–1061.
- [44] Pinghui Wang, Yiyun Qi, Yu Sun, Xiangliang Zhang, Jing Tao, and Xiaohong Guan. 2017. Approximately Counting Triangles in Large Graph Streams Including Edge Duplicates with a Fixed Memory Usage. *PVLDB* 11, 2 (2017), 162–175.
- [45] Pinghui Wang, Junzhou Zhao, Xiangliang Zhang, Zhenguo Li, Jiefeng Cheng, John C. S. Lui, Don Towsley, Jing Tao, and Xiaohong Guan. 2018. MOSS-5: A Fast Method of Approximating Counts of 5-Node Graphlets in Large Graphs. *IEEE Trans. Knowl. Data Eng.* 30, 1 (2018), 73–86.
- [46] Yanhao Wang, Qi Fan, Yuchen Li, and Kian-Lee Tan. 2017. Real-Time Influence Maximization on Dynamic Social Streams. *PVLDB* 10, 7 (2017), 805–816.
- [47] Yanhao Wang, Yuchen Li, Ju Fan, and Kian-Lee Tan. 2018. Location-aware Influence Maximization over Dynamic Social Streams. *ACM Trans. Inf. Syst.* 36, 4 (2018), 43:1–43:35.
- [48] Yanhao Wang, Yuchen Li, and Kian-Lee Tan. 2019. Semantic and Influence aware k-Representative Queries over Social Streams. In *EDBT*. 181–192.
- [49] Bin Wu, Ke Yi, and Zhenguo Li. 2016. Counting Triangles in Large Graphs by Random Sampling. *IEEE Trans. Knowl. Data Eng.* 28, 8 (2016), 2013–2026.
- [50] Qiankun Zhao, Yuan Tian, Qi He, Nuria Oliver, Ruoming Jin, and Wang-Chien Lee. 2010. Communication motifs: a tool to characterize social communications. In *CIKM*. 1645–1648.