

Temporal Network Motifs: Models, Limitations, Evaluation

Penghang Liu¹, Valerio Guarrasi², and Ahmet Erdem Sariyüce¹

Abstract—Investigating the frequency and distribution of small subgraphs with a few nodes/edges, i.e., motifs, is an effective analysis method for static networks. Motif-driven analysis is also useful for temporal networks where the spectrum of motifs is significantly larger due to the additional temporal information on edges. This variety makes it challenging to design a temporal motif model that can consider all aspects of temporality. In the literature, previous works have introduced various models that handle different characteristics. In this work, we compare the existing temporal motif models and evaluate the facets of temporal networks that are overlooked in the literature. We first survey four temporal motif models and highlight their differences. Then, we evaluate the advantages and limitations of these models with respect to the temporal inducedness and timing constraints. In addition, we suggest a new lens, event pairs, to investigate temporal correlations. We believe that our comparative survey and extensive evaluation will catalyze the research on temporal network motif models.

Index Terms—Network motifs, temporal networks, temporal motifs

1 INTRODUCTION

AN important local property of networks is the network motifs, which are defined as the limited size, recurrent and statistically significant patterns [1]. Motifs are shown to be more effective when considered as connected, non-isomorphic, and induced subgraphs [2]. Motifs are used to model and examine interactions among small sets of vertices in networks. Finding frequent patterns of interactions can reveal functions of participating entities [3], [4], [5], [6], [7], [8] and help characterize the network. Also known as higher-order structures, motifs are regarded as basic building blocks of complex networks in domains such as social networks, food webs, and neural networks [1]. The triangle, for example, is the most basic motif in simple undirected networks and plays an important role in defining the global network characteristics such as clustering coefficients.

In temporal networks, edges are associated with temporal information (timestamps), which indicate the occurrence time. We call the temporal edges as events. For example, in an email network each node is an email address and an event represents the email sent from one address to another at a specific time, or an event in the call detail records is a phone call from a person to another at a certain time.

Temporality brings new challenges for network analysis [9]. Motif-driven techniques, for instance, should

consider the temporal information on edges which significantly increases the spectrum of motifs with respect to static networks. The event order, inter-event time intervals, and durations are some of the aspects that need to be incorporated [10]. Thus, it is beyond non-trivial to design a model for temporal network motifs that considers all those characteristics while being practical. There are several studies [11], [12], [13], [14] and each proposes a temporal motif model in a different way. A motif can be valid in some models but not in the others, due to the specific constraints required by the different models. Also, the prior studies are introduced in various subfields of computer and network science and mostly unaware of each other. Consequently, there does not exist a unified approach that can address the limitations of those models while leveraging their novelties. A comparative evaluation on these model is essential in that respect. Note that the motif-driven approach is different from the temporal subgraph isomorphism or periodic subgraph mining [15]. The motifs have limited sizes, when compared to the subgraphs, and it enables the exploration of all the combinations of edges among the nodes and all the permutations of edges in temporal order.

In this work, we introduce a comparative survey for the four models: by Kovanen *et al.* [11], [16], Song *et al.* [12], Hulovatyy *et al.* [13], and Paranjape *et al.* [14]. These four models are the first works that introduced new approaches while handling various aspects of temporality. They are also used heavily in various applications (Section 3). We aim to address three questions:

- 1) What are the differences between these models in terms of the aspects of temporality and the constraints incorporated?
- 2) How do these differences affect the frequency and spectrum of the resulting temporal motifs?
- 3) What are the implications of these effects and how do they benefit real-world applications?

- Penghang Liu and Ahmet Erdem Sariyüce are with the Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY 14260 USA. E-mail: {penghang, erdem}@buffalo.edu.
- Valerio Guarrasi is with the Sapienza University of Rome, 00185 Rome, Italy. E-mail: valerio.guarrasi@uniroma1.it.

Manuscript received 11 September 2020; revised 15 March 2021; accepted 26 April 2021. Date of publication 4 May 2021; date of current version 7 December 2022.

(Corresponding author: Penghang Liu.)

Recommended for acceptance by K. Daudjee.

Digital Object Identifier no. 10.1109/TKDE.2021.3077495

For the first, we give a comparison with respect to the different aspects of temporality, and highlight the advantages and limitations of each (Section 4). For the second, we evaluate two key aspects of temporality: temporal inducedness and timing constraints (Section 5). We focus on the temporality features that are overlooked in previous studies, such as the behavior of intermediate events. We also suggest a new lens, event pairs, to analyze the characteristics of the sequences in temporal motifs and investigate the temporal correlations. Lastly, we interpret the results in each experiment and discuss what domains can be affected by the assumptions in those four models (Section 5). We explore all three-event two-/three-nodes (36 in total) and four-event two-/three-/four-nodes (696 in total) motifs. We believe that our work will steer the temporal motif research in a healthy direction.

2 BACKGROUND

$G(V, E)$ is a temporal network where V is the set of nodes and E is the set of events. Each event $e_i \in E$ is a 4-tuple $(u_i, v_i, t_i, \Delta t_i)$ [10]. u_i and v_i are the (source and target) node pair where the i th event occurs, t_i is the starting time of i th event, and Δt_i is its duration. E is a time-ordered list of m events where the starting times are $t_1 \leq t_2 \leq t_3 \leq \dots \leq t_m$, and V is the set of nodes that appear in E . In real-world temporal networks, it is very common that the inter-event time, $t_{i+1} - t_i$, is significantly larger than Δt_i , thus event durations can be ignored. For simplicity, we also follow this convention in our work and consider each event in set E as a 3-tuple (u_i, v_i, t_i) . Here we distinguish edges and events, where the edge (u, v) is the static projection of an event (u, v, t) . We also refer the set of motifs with the number of nodes and events; e.g., 3n3e motifs have three-nodes and three-events.

3 RELATED WORK

In this work, we focus on four temporal motif models [11], [12], [13], [14] and give a detailed overview and comparison in Section 4 (Note that the works by Nicosia *et al.* [17] and [18] also provided short surveys covering a few works available). Here we summarize the previous work on (1) applications of temporal motifs, (2) algorithmic improvements for temporal motif finding, and (3) temporal subgraph isomorphism, a related topic to our subject.

Applications. Motif has been a versatile tool for several application domains that are engaged with the temporal networks. Building on the notion of static motifs [1], [2], there have been several works that considered snapshot-based adaptations to incorporate the temporality, listed below in chronological order. Jin *et al.* [19] investigated the temporal node-weighted networks and devised trend motifs based on the weight changes on the nodes over a specified period of time. They counted the trend motifs in financial and protein-interaction networks. Chechik *et al.* analyzed a yeast metabolic network by activity motifs to understand timings of transcriptions [20]. Activity motif is defined as a partially (or totally) ordered combination of chains, forks, and joins to understand the interactions among gene activations and repressions. Zhao *et al.* proposed communication motifs to

study synchronous and asynchronous human communication networks, such as call detail records (CDR) and Facebook wall post interactions, and understand the information propagation [21]. Communication motif is defined as a static network motif where each connected edge pair satisfies a timing constraint and there is no particular order defined among the edges. Lastly, Sarkar *et al.* [22] investigated the undirected static motifs in the subsequences of microblog network snapshots to study the information diffusion process. They created motif based features to predict the network at the stage of inhibition.

Bajardi *et al.* worked on the cattle trade movements among farms in Italy and used the dynamic motifs, which are defined as the chain motifs ordered in time, to model cause-effect relations [23]. Faisal and Milenkovic introduced static-temporal motifs to study human aging [24]. Kovanen *et al.* introduced the first holistic temporal motif model that is independent of a particular topology and explicitly considers the temporal adjacency [11] (more details in Section 4). Most of the following studies embraced this model. Jurgens and Lu investigated the editor interactions in Wikipedia with temporal motifs [25]. Kovanen *et al.* adapted their model [11] for colored networks, where the colors denote categorical node attributes, and studied a CDR with respect to sex, age group, and subscription type attributes [26]. Zhang *et al.* analyzed two bipartite networks, ship-chartering and ship order-to-build networks, to see how the 2,2-bicliques are formed in time [27]. Their analysis does not rely on Kovanen *et al.*'s model and can be considered as a snapshot based approach. Li *et al.* performed a study on mobile communication networks to analyze the ordering among edges in three-node motifs [28]. Zhang *et al.* introduced an extensive analysis for online and offline human interactions [29] by using the temporal motif model by Kovanen *et al.* [11]. They analyzed phone messages as well as face-to-face interactions (by RFID) and sexual contacts by using 3-event and 4-event temporal motifs. Xuan *et al.* studied the online task-oriented networks, where people collaborate to work on tasks, and used temporal motifs to understand the collaboration patterns [30]. Li *et al.* considered temporal motifs for heterogeneous networks where there are multiple types of nodes and edges [31]. They analyzed the DBLP network (of papers, authors, terms, venues, and years), meme-tracker data, and news articles (of documents, locations, and topics). More recently, Kosyfaki *et al.* proposed flow motifs to analyze the dynamics of evolution in edge-weighted temporal networks, such as bitcoin transactions, passenger flows, and facebook interactions (aggregated over certain time intervals) [32].

Algorithmic Improvements. Algorithmic improvements are studied for faster temporal motif finding and counting. Gurukar *et al.* proposed a fast technique called COMMIT [33] to find the communication motifs proposed by [21]. Building on Paranjape *et al.*'s model [14], Kumar *et al.* introduced an efficient algorithm to enumerate temporal cycles of any length [34] by extending the seminal cycle counting algorithms of Johnson [35]. Sun *et al.* introduced fast algorithms to find temporal motifs (by [11]'s model) by using a time-first search approach [36], [37]. Liu *et al.* presented sampling algorithms to find the approximate frequency of temporal motifs up to two orders of magnitude faster [38].

Temporal Subgraph Isomorphism. A related subject is the problem of temporal subgraph isomorphism. Redmond and Cunningham first introduced the problem to find a query subgraph with temporal constraints in a given temporal network [39], [40]. Unlike temporal motifs, there is no size constraint in the isomorphism problem. Mackey *et al.* proposed another variation where there is a total ordering among the edges [41]. Franzke *et al.* introduced another definition where there is more flexibility in defining the temporal order of edges; e.g., relative time intervals can be specified for each edge in the temporal subgraph [42].

4 OVERVIEW OF TEMPORAL MOTIF MODELS

To the best of our knowledge, there are four models:

- Kovanen *et al.* [11] proposed the first model that has the notion of temporal adjacency to relate the events in a motif.
- Song *et al.* [12] introduced another model for streaming workloads where the motifs are found on-the-fly and the events in a motif can be partially ordered.
- Hulovatyy *et al.* [13] considered new relaxations and restrictions to improve Kovanen *et al.*'s model and also discussed the events with durations.
- Paranjape *et al.* [14] proposed a practical model with a specified time window to bound all events in a motif.

Note that we focus on motif models in this work and refrain from discussing the temporal subgraphs, which are larger in size and often only considered with respect to a template. We first briefly explain the main idea in each paper and then compare them with respect to different aspects.

The first temporal network motif model is introduced by Kovanen *et al.* [11]. The idea is to use edge timestamps to build more expressive motifs than the classical network motifs in static networks [1]. Kovanen *et al.* define the temporal motif as an ordered set of events with two features: (1) time difference between each pair of consecutive events (in the whole set) is less than the threshold Δ_C , an input parameter, (2) for each node in the motif, its adjacent events in the motif are consecutive, i.e., the node does not participate in any other event between its events in the motif. This is aimed to consider causality among the events.

Song *et al.* [12] proposed the event pattern matching problem for the real-time graph streams. By considering the graph structure, event pattern matching improves the traditional complex event processing [43]. They approached the problem from time series and stream processing perspective and used the graph structure as a new feature. Event pattern is simply a temporal motif model that considers node/edge labels, partial orderings among events, and an input parameter, Δ_W , which is an upper bound for the time difference between the first and the last events in the motif.

Hulovatyy *et al.* [13] came up with another temporal motif model based on the notion of graphlets (induced motifs) in static networks [2]. Hulovatyy *et al.* improved the model by Kovanen *et al.* by (1) only considering the induced subgraphs, i.e., all interactions among a given set of nodes are taken into account, and (2) relaxing the constraint that adjacent events of a node should be consecutive. They also discussed the use of events with durations in temporal

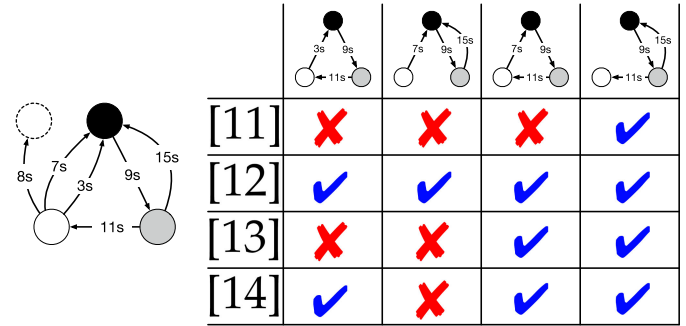


Fig. 1. Comparison of four temporal motif models [11], [12], [13], [14]. A temporal network with six events is shown on the left. The table on the right has four motifs and shows whether they are valid motifs according to the four models, where $\Delta_C=5s$ (inter-event timing constraint) and $\Delta_W=10s$ (entire motif timing constraint). The first motif is not valid according to [11], [13] because it breaks the Δ_C constraint; the second motif is invalid in [14] since it is not an induced subgraph, and it also violates the Δ_C constraint in [11], [13]; the third motif violates the consecutive event restriction, thus not a valid motif in [11]; and the last motif is valid according to all the models.

network motifs, for the first time. Furthermore, they introduced an additional restriction (constrained dynamic graphlets) to reduce the computational complexity while obtaining approximate results. Authors have shown that their new model captures various temporal motifs from each node's perspective and gives more effective results than the prior techniques for predicting aging-related genes in humans.

The last model, by Paranjape *et al.* [14], also considers a relaxation of the first model proposed by [11]. The constraint that the adjacent events of a node should be consecutive (in [11]) is relaxed so that the motifs that occur in a short burst can be caught. Time window, Δ_W , to bound the time difference between the last and the first events in a motif.

A given motif can be valid in some models but not in the others, due to the specific constraints required by the different models. Fig. 1 gives four example motifs in a temporal network and describes the validity of each motif with respect to the four models.

There are several aspects of temporal networks and motifs that are handled differently in each of those four models. Table 1 presents an overview. Those aspects are crucial across the diverse application space of temporal networks. For instance, fraudsters in financial transaction networks camouflage their identities by getting involved in repetitive legal transactions. A strictly induced temporal motif is helpless in this context since it considers all the transactions among a set of entities in which the few fraudulent transactions can be overlooked. In a communication network, on the other hand, it might make more sense to use induced motifs to fully understand the information dissemination. Here we discuss each aspect and highlight the advantages and limitations of the four models accordingly.

4.1 Motif as Induced Subgraph

In static networks, considering all the edges among a given set of nodes (rather than selecting a subset) has been shown to be more effective in motif-based analysis [44]. Because the non-induced motifs become artificially recurrent and shadow the importance of larger induced structures. For

TABLE 1
Aspects of Temporal Motif Models

Article	Kovanen et al. [11]	Song et al. [12]	Hulovatyy et al. [13]	Paranjape et al. [14]
Induced subgraph (Sec. 4.1)	<i>Node-based temporal</i>	✗	<i>Static only</i>	<i>Static only</i>
Event durations (Sec. 4.2)	✗	✗	✓	✗
Partial ordering (Sec. 4.3)	✓	✓	✗	✗
Directed edges (Sec. 4.4)	✓	✓	✗	✓
Node/Edge labels (Sec. 4.4)	✗	✓	✗	✗
Adjacent events in Δ_C (Sec. 4.5)	✓	✗	✓	✗
Entire motif in Δ_W (Sec. 4.5)	✗	✓	✗	✓

instance, an induced square motif ($1 \rightarrow 2$, $2 \rightarrow 3$, $3 \rightarrow 4$, $1 \rightarrow 4$) implies that no diagonal edges exist (i.e., $1 \rightarrow 3$ and $2 \rightarrow 4$ are missing) whereas a non-induced square motif has no such restriction (i.e., every 4-clique is also a square). In temporal networks, “inducedness” is more involved. In order to capture all the interactions among a given set of nodes, one should consider the time interval in which those events occur and also watch for the number of events in that period. Formally, a temporal subgraph induced by the node set $V' \in G$ for a time interval $[t_s, t_f]$ includes all the events (u, v, t) in G such that $u, v \in V'$ and $t_s \leq t \leq t_f$. If there is an additional restriction for the number of events, say k (as in the motif definition), then one can consider only the consecutive k events to form a motif. This means that there should not be any other event $(u, v, t) \in G$ such that $u, v \in V'$ and $t_i \leq t \leq t_{i+1}$ for any consecutive event pair in the motif with timestamps t_i and t_{i+1} .

The first model in [11] does not require a motif to be induced in the static sense, e.g., a diagonal edge in the square example above can be allowed. However, one important condition in Kovanen *et al.*'s model [11] is that a node's adjacent events in a motif should be consecutive in time, i.e., there cannot be any adjacent event outside the motif which occurs between two events in the motif. For example, if there is a temporal motif with events $(u, v, 5)$, $(v, w, 8)$, and $(u, v, 12)$, then there cannot be any other event in the graph that contains u and occurs in $[5, 12]$ interval. Likewise, no event adjacent to v can exist in $[8, 12]$ interval. We call this *consecutive events* restriction. This can be seen as *node-based temporal inducedness* since all the events that are adjacent to a node should be part of the motif for a given time interval. This also avoids the exponentially many motifs in certain scenarios, like when a node has a burst of events in the form of a star. Note that, Hulovatyy *et al.* [13] and Paranjape *et al.* [14] argued that this constraint is too restrictive and shadows important motifs occurring in short bursts. Song *et al.*'s model [12] does not rely on such a restriction either.

Song *et al.* [12] approach the inducedness from a different perspective. They argue that if the temporal data is observed in a streaming setting, it is often desirable to catch certain events in a specific structure, e.g., some temporal and non-induced motifs (like squares) in financial transaction networks are a strong indicator of fraud [34], [45]. Therefore, unlike static networks, non-induced motifs in temporal networks are valuable, especially when the network is being streamed and time-sensitive decision making is demanded. Hence, [12] does not consider induced subgraphs.

Hulovatyy *et al.* [13] is the first to discuss the lack of induced subgraphs in the initial model by Kovanen *et al.* [11]. Authors argue that the temporal motif must be induced by relying on Przulj *et al.*'s work where they proposed to use induced motifs for static networks since the non-induced counts would be artificially amplified [44]. However, the proposed model in [13] only captures the static inducedness, i.e., considering all the edges among the given set of nodes in a static projection (e.g., a diagonal edge in the square example above is not allowed). There is no restriction for the temporal behavior, even in the node-level since the *consecutive events* restriction (of [11]) does not exist. For example, given four events $(a, b, 2)$, $(b, c, 4)$, $(c, a, 5)$, $(c, a, 6)$; the triangle formed by the first, second, and fourth events is a valid motif in [13]'s model (third event can be ignored). There is another approach in [13], named *constrained dynamic graphlet*, that is somehow related to the inducedness. Considering a snapshot-based representation where consecutive snapshots have a similar dense structure for some set of nodes, the events that are being repeated in the new snapshots do not give any new information and also cause redundancy in computation. Hence, watching for the new events that are not observed in the prior snapshots is more interesting and also more efficient. Hulovatyy *et al.* incorporates this observation in the *constrained dynamic graphlet* model such that if two events (u_1, v_1, t_1) and (u_2, v_2, t_2) are consecutive in a temporal motif (where $u_1, v_1 \neq u_2, v_2$), then there is no event (u_2, v_2, t') in the temporal graph for which $t_1 \leq t' \leq t_2$. In a sense, this enables the discovery of temporal motifs where there are causal relationships.

The model by Paranjape *et al.* [14] requires the motifs to be induced, but only in the graphlet projections, as in [13].

4.2 Event Durations

As defined in Section 2, each event in a temporal network can have a duration that denotes how long the event exists, e.g., each phone call in CDR has its duration. Temporal motifs should be modeled to take the durations into account in such cases. This is acknowledged in [11] but omitted from the model for simplicity. Likewise, [14] mentions that their model can be generalized for the events with duration but it is not clear how the timing constraints would be adjusted. In [12], authors consider the event duration as an edge label where the motifs can be defined accordingly, e.g., a motif where the duration of each edge is less than 30 secs. The only work that incorporated the duration into the model definition is [13] where they propose the dynamic

graphlet model; the time difference between two consecutive events are determined with respect to the end time of the first and the start time of the second event.

4.3 Partial Ordering

Temporal networks have timestamps associated with the edges and a temporal motif is often defined with respect to the ordering among its edges. However, the ordering among the events of a motif can be total or partial. In the former, each event pair is ordered and the order of all events are unique. In the partial ordering, there can be some event pairs for which the ordering is undefined. For instance, an acyclic-triangle motif with nodes A, B, C and edges $A \rightarrow B, A \rightarrow C, B \rightarrow C$ can be defined in a way that $B \rightarrow C$ precedes both $A \rightarrow B$ and $A \rightarrow C$. A temporal motif with a partial ordering can always be expressed as a set of multiple motifs where each covers a different total ordering for the nodes who are partially ordered, i.e., the example above is union of $(B \rightarrow C) \prec (A \rightarrow C) \prec (A \rightarrow B)$ and $(B \rightarrow C) \prec (A \rightarrow B) \prec (A \rightarrow C)$. However, finding all possibilities is not practical and also redundant. [11] and [12] consider this and define temporal motif models with a partial ordering. In both models, there is a strict partial ordering among the events, which means that the ordering is irreflexive ($e_i \not\prec e_i$), transitive ($(e_i \prec e_j \wedge e_j \prec e_k) \rightarrow e_i \prec e_k$), and asymmetric ($e_i \prec e_j \rightarrow e_j \not\prec e_i$). Note that, this ordering assumes that each event in the given temporal network has a unique timestamp (timestamps are strictly increasing), which is not realistic: many real-world temporal networks have events (on different edges) with the same timestamp (timestamps are non-strictly increasing), e.g., an email network where a person can send an email to multiple people at the same time. Furthermore, if the temporal network's timespan is too large, it is often desirable to reduce the resolution by creating snapshots. Thus, the models in [11] and [12] do not handle such networks. [14] only mentions that their model can be extended to handle partial ordering but does not provide further details. In [13], partial ordering is not taken into account at all and the motif model is defined to have a total ordering among its events. Note that the models that assume a total ordering also fail to handle the networks where event timestamps are not unique.

4.4 Directed Edges, Node/Edge Labels

All models, except [13], considered directed edges in their definitions. Hulovatty *et al.* only mention that their model is extendible for directed edges [13]. For the node/edge labels, only [12] established their model accordingly. We believe that other models can also handle node/edge labels. Note that none of the models consider edge weights, which is also challenging and mostly overlooked for static networks.

4.5 Timing Constraints

Connectivity in the temporal dimension is a key feature for temporal motif models. There have been different approaches to formalize temporal connectedness in the previous works [11], [12], [13], [14]. Kovanen *et al.* defined the temporal motif as a connected temporal subgraph such that for any pair of consecutive events that share a node, the time difference should be less than Δ_C [11]. The same

approach is also used by Hulovatty *et al.* [13]. Note that both models require graph connectivity to consider consecutive events. On the other hand, [12] and [14] consider a window-based temporal connectivity where all the events in a temporal motif must occur within a given time interval, denoted as Δ_W . Namely, the time difference between the last and the first events is limited by Δ_W .

Those two approaches yield temporal motifs with different semantics. Δ_C is useful to detect temporal correlations since each pair of consecutive events should occur in a certain time period, but it fails to bound all the events in a motif, i.e., can only give a loose limit, $(|E'| - 1) * \Delta_C$, for the entire motif where E' is the set of events. Δ_W , on the other hand, presents a holistic temporal view for the entire motif but cannot consider the temporal correlations between consecutive events. Consider a connected temporal motif with three ordered events where we set $\Delta_C = 5$ for the models in [11], [13] and $\Delta_W = 10$ for the models in [12], [14]. If the timestamps of those three events are 1, 9, and 10, Δ_W based models consider this motif valid, but Δ_C based models do not since the first two events are not close enough for $\Delta_C=5$. One can say that the first two events are too far apart, thus there is no temporal correlations. However, it is also possible that the third event occurred right after the second only because of some important information that is initially conveyed by the first event; thus there is a temporal correlation between the first and third events. Either interpretation can make sense depending on the type of temporal network being considered. For instance, a temporal cycle can be formed this way where person A tells person B that she started a new job. Then B sees C like a week later and tells about A 's new job, which immediately results in C calling A for congratulations. Here the delayed convey of the information can be captured with Δ_W parameter.

One can consider to use both parameters to have a trade-off between the two extremes of Δ_W and Δ_C . Depending on the number of events in the temporal motif, one of those two timing constraints can be useless for certain values of Δ_W and Δ_C . Given a motif with m events and Δ_C/Δ_W ratio, we have

$$\text{Constraints} = \begin{cases} \Delta_C & \text{if } 0 \leq \frac{\Delta_C}{\Delta_W} \leq \frac{1}{m-1} \\ \Delta_C, \Delta_W & \text{if } \frac{1}{m-1} < \frac{\Delta_C}{\Delta_W} < 1. \\ \Delta_W & \text{if } \frac{\Delta_C}{\Delta_W} \geq 1 \end{cases}$$

There are $m - 1$ time intervals among m events. The loose bound defined by Δ_C for the entire motif window is $\Delta_C * (m - 1)$. In the first case, any Δ_W larger than that bound is meaningless; so satisfying Δ_C constraint is sufficient. It is also meaningless to consider a Δ_C value (third case) that is not smaller than Δ_W ; just considering Δ_W would also satisfy the other. The only case where both constraints make sense is when Δ_W is smaller than the loose bound $\Delta_C * (m - 1)$ (second case). Exploring the parameter space in that case may enable to consider temporal motifs with respect to both inter-event timings (Δ_C) and entire motif timing (Δ_W).

5 EXPERIMENTAL EVALUATION

In this section, we present an extensive evaluation for the various aspects considered in the four temporal network

TABLE 2
Important Statistics for the Temporal Network Datasets We Use in Our Experiments

Name	Nodes	Events	Edges	#T	$ E_u / E $	$m(\Delta_t)$
Bitcoin-otc	5.88K	35.6K	35.6K	35.4K	99.2%	707
CollegeMsg	1.90K	59.8K	20.3K	58.9K	97.2%	37
Calls-Copen.	536	3.60K	924	3.59K	99.7%	194
SMS-Copen.	568	24.3K	1.30K	24.0K	97.6%	32
Email	986	332K	24.9K	208K	50.5%	15
FBWall	47.0K	877K	274K	868K	98.0%	42
SMS-A	44.4K	548K	69.0K	470K	73.1%	3
StackOver.	260K	6.35M	4.15M	5.97M	88.2%	6
SuperUser	194K	1.44M	925K	1.44M	99.2%	83

motifs [11], [12], [13], [14]. We are particularly interested in the impacts of temporal inducedness and timing constraints. Our code (including plot scripts) is available for reproducibility: <https://codeocean.com/capsule/dee9010a-48c6-45a1-8609-6ca8fb9ea942/>.

Datasets. We select various directed temporal network datasets from several domains, including phone messages, emails, Facebook wall interactions, posts in Q/A websites, and call detail records (CDR). Table 2 gives several statistics about our datasets. For each network, we provide the number of nodes, events (interactions), edges (unique pairs of nodes in events), timestamps (#T, unique timestamps across the entire timespan), the percentage of events with unique timestamp ($|E_u|/|E|$), and the median interevent-time ($m(\Delta_t)$, in seconds). The time resolution of all the networks is one second. In addition to the number of events, edges, and the percentage of events with unique timestamp, we give the median interevent-time for each dataset, which is the median of time intervals between all pairs of consecutive events in the network. This gives us an idea about how to choose the timing parameters (Δ_C and Δ_W) in order to address the trade-off between discovering more motifs and reducing the computational costs. In the phone message networks, an event (u, v, t) represents a message sent to person v by person u at time t ; we have SMS-A [46], SMS (Copenhagen) [47], and College-messages [48]. Email presents the emails between members of European research institution [48], where an event (u, v, t) denotes an email sent from person u to person v at time t . Among the online social networks; FBWall is the collection of posts between users on Facebook in the New Orleans region [49], where an event (u, v, t) indicates user u posted on the user v 's wall at time t ; StackOverflow and SuperUser are the interaction networks on two stack exchange websites [48], where event (u, v, t) stands for the answer/comment user u posted on user v 's question/answer at time t . Note that we slice out the earliest 10 percents of the events from the original StackOverflow datasets for the efficiency purposes. We also study the Bitcoin-otc, a trust network where users rate each other to reflect their trust regarding the bitcoin transactions [48]. We also use Calls (Copenhagen), which is the collection of phone calls between university students over a period of four weeks [47].

Motif Notation. We consider three-event and four-event motifs in this work and it is not feasible to visualize all. Instead, we introduce a notation to refer to the temporal

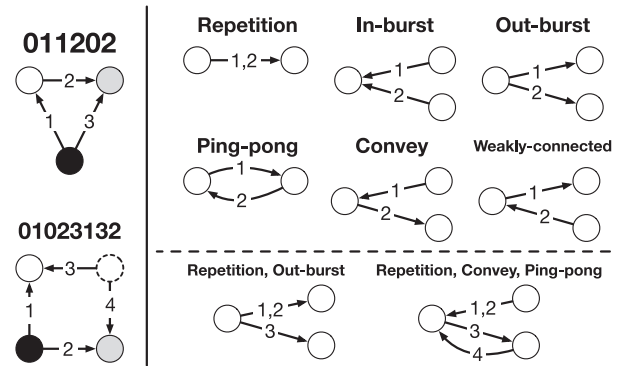


Fig. 2. [Left] We use $2n$ digits to denote a temporal motif with n events. Each event is given by a pair of digits, where the source node is the first and the target node is the second digit. The first two digits are always 01, to denote that the first event occurred from node 0 to node 1. The sequence of pairs and the digit number of each node follow the chronological order of events and nodes. [Right] Event pair representations, all six types are listed. At the bottom, a three-node three-event motif is denoted as a sequence of repetition and out-burst, and a four-event motif is denoted as a sequence of repetition, convey, and ping-pong.

motifs in our experiments. We use $2n$ digits to denote a temporal motif with n events. Each pair of digits is an event from the node represented by the first digit to the node denoted by the second digit. The first two digits are always 01, to denote that the first event occurred from node 0 to node 1. The sequence of pairs and the digit number of each node follow the chronological order of events and nodes. Fig. 2 presents some examples. For instance, 011202 (top-left) corresponds to a triangle temporal motif where the first event is from the black node (0) to the white node (1), the second one is from the white node (1) to the gray node (2), and the last event is from the black node (0) to the gray node (2). Note that we only consider the motifs that grow as a single component, by adding one event at a time.

A New Lens: Event Pairs. In this work, we consider an alternative way to interpret the temporal motifs [29]. We simply look at the sequence of two events that share a node (i.e., event pairs). Event pairs can be seen as the building blocks of the larger motifs (with ≥ 2 events). Given a pair of consecutive events that share a node in the motif, (u_1, v_1, t_1) and (u_2, v_2, t_2) , where $t_1 < t_2$, there are six types of event pairs (also shown in Fig. 2):

- **Repetition.** Two events occur on the same edge ($u_1 = u_2, v_1 = v_2$).
- **Ping-pong.** Second event is reverse of the first one ($u_1 = v_2, v_1 = u_2$).
- **In-burst.** Two events share the same target ($u_1 \neq u_2, v_1 = v_2$).
- **Out-burst.** Two events share the same source ($u_1 = u_2, v_1 \neq v_2$).
- **Convey.** The source of the second event is the target of the first event ($v_1 = u_2, u_1 \neq v_2$).
- **Weakly-connected.** The target of the second event is the source of the first event ($u_1 = v_2, v_1 \neq u_2$).

A motif with m events can be represented as a sequence of $m - 1$ event pairs, as long as each consecutive event pair shares a node (see Fig. 2 for examples). Considering event pairs as the building blocks is useful for three reasons:

- *It is a 6-letter alphabet*, thus a simple and succinct way to deal with the temporal and structural complexity.
- *It is expressive*; can exactly represent all $2n3e$ or $3n3e$ motifs (36 in total, 6^2) and all $3n4e$ or $2n4e$ motifs (216 in total, 6^3). It also gives 216 (6^3) broad descriptions (not exact) for the 480 $4n4e$ motifs that grow as a single component.
- *It is well-suited to capture the temporal correlations*: event pairs within larger motifs can reveal the mesoscale characteristics while providing an easy interpretation.

Questions. In our experimental evaluation, we are looking to answer the following questions about temporal network motifs by using the various real-world temporal networks in our dataset:

Q1: What are the implications of temporal inducedness constraints discussed in Section 4.1? In particular, is there any bias in the spectrum or counts of motifs obtained by consecutive event restriction in [11] and constrained dynamic graphlets in [13]? We answer those in Section 5.1.

Q2: How the timing constraints, Δ_C and Δ_W , impact the spectrum or counts of motifs? What are the biases implied by each constraint? How would the combination of those two constraints compare? We consider those in Section 5.2. Note that we avoid the experiments in prior studies [14] about the choice of parameters.

Q3: What new insights can we capture by using the lens of event pairs? What kind of event pairs observed in each dataset? Is there any commonality for all datasets or for the ones from the same domain? We answer all in Section 5.3.

Comparison Criteria. We are only interested in the motif counts and behavior of the motif spectrum in our evaluation; runtime performance is not in our scope in this work (but is a promising future direction). Note that, in static motif analysis [1] a null model (randomized graph) is often used as reference to measure the statistical significance of motifs. However, it is difficult to choose a proper null model for temporal motifs: the intertwined temporal and structural complexity in temporal networks makes it challenging to consider a null model that can mimic the realistic features. This issue is also highlighted by Kovanen *et al.* [11]. Although there has been several new random reference models for temporal networks [50], our preliminary analysis showed that there is no reliable random graph model which can mimic both the structural and temporal features in real-world networks. We tried several link-shuffling and time-shuffling models from [50], some are too restrictive where the motif counts barely change, and some others are too loose where all the motifs are reported as significant. Therefore, as also done in [11], we consider the motif count as the main indicator of motif significance, which gives a relative comparison. We consider three-event and four-event motifs in our experiments.

5.1 Temporal Inducedness Constraints

As discussed in Section 4.1, temporal network motif models have different approaches to define temporal inducedness. Here we evaluate the impact of two restrictions by using $3n3e$ (three-nodes, three-events) motifs: (1) Consecutive events, proposed in [11] to ensure node-based temporal

TABLE 3
The Impact of Consecutive Events Restriction

Network	Non-cons.	Cons.	010210	011210	012010	012110
Calls-Cop.	1.66K	77	+7	-9	0	0
CollegeMsg	1.59M	2.55K	+18	+23	+10	+16
SMS-Copen.	198K	389	+16	+18	+14	+17
SMS-A	1.02M	35	+11	+11	+2	+6
Email	1.84M	1.82K	+1	+4	+4	+4
FBWall	268K	904	+14	+10	+13	+2
Bitcoin-otc	6.42K	1.93K	0	0	+2	+3
StackOver.	3.78M	1.02K	-11	+8	-5	0
SuperUser	481K	10.9K	+4	+4	-4	+5

inducedness, (2) Constrained dynamic graphlets, proposed in [13] to filter out the stale information.

5.1.1 Consecutive Events Restriction

In this part, we evaluate the advantages and limitations of adding the consecutive events restriction from [11]. As explained in Section 4.1, consecutive events restriction is a node-based constraint. If a node is part of a motif, then the adjacent events of the node in the motif should be consecutive, i.e., the node cannot have an adjacent event outside the motif while it is engaged in the motif. The third column in Fig. 1 also demonstrates this restriction, where white node interacts with the dashed node at $t = 8$ thus the motif is not valid according to [11]. Consecutive events restriction is useful when handling a star node for instance, which is adjacent to many events. The continuity requirement ensures that the node can only be engaged with a linear number of motifs, avoiding quadratically many. However this is a double-edged sword: it permits fast counting and analysis, but can miss important patterns. Here we count the motifs with and without consecutive events restriction on all 32 $3n3e$ motifs and use $\Delta_C = 1500s$ timing configuration (i.e., no Δ_W is specified). Table 3 presents the results (all results are available in the supplementary materials, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TKDE.2021.3077495>). We compare the total counts of $3n3e$ motifs without or with consecutive event restriction for $\Delta_C=1500s$ (2nd and 3rd columns). We also give the ranking changes of the four motifs among all 32 $3n3e$ motifs after the consecutive restriction is added (last four columns). Overall, the given motifs are amplified when the consecutive event restriction is applied.

The motif counts show that in all datasets except Bitcoin-otc, over 95 percent of the motifs are removed if the restriction is applied. We observe that four motifs, 010210, 011210, 012010, and 012110, are often amplified by the consecutive event restriction. These motifs are not frequently observed in the non-consecutive scenario but rise to a significantly higher ranking after the consecutive restriction is applied. The amplification of these motifs is most commonly observed in message networks, such as CollegeMsg and SMS (Copenhagen) All of these motifs follow an ask-reply pattern, where the last event replies the first event. It is not an immediate response and another node is involved in the second event, which appears to be another conversation. Note that if the consecutive events restriction is not

considered, adjacent events of a node are not required to be consecutive. Then, the second event would be allowed to form a different motif, by skipping the third event (reply) and adding another adjacent event from another conversation. This would shadow the original ask-reply pattern between the first and the third events and, in turn, overstates the motifs where the first and second events form a star. When it is important to distinguish ask-reply patterns from the big star structures, consecutive events restriction can be advantageous. For instance, using consecutive events restriction can be useful for temporal networks of text messages and emails. A burst of messages or emails from one individual to others may be spam (or the information value may not be high), while two reciprocal messages or emails in short time are likely to indicate a real conversation. On the other hand, considering this restriction in other temporal networks, like online interactions in social networks, may suggest biased results by artificially amplifying the ask-reply patterns. *All in all, while the consecutive events restriction filters out the discontinuous engagements from the node's perspective, it yields biased motif counts by consistently amplifying ask-reply motifs where the first and last events are reciprocal.* This can be advantageous in the analysis of message and emails (where bursts of messages/emails are too common) but may introduce bias in the analysis of other temporal networks.

5.1.2 Constrained Dynamic Graphlets

As defined in Section 4.1, constrained dynamic graphlets are defined in [13] to exclude stale information in the temporal motifs. Constrained dynamic graphlet only includes the new events that are not observed in the prior timestamps; formally, if two events (u_1, v_1, t_1) and (u_2, v_2, t_2) are consecutive in the graphlet (where $u_1, v_1 \neq u_2, v_2$), then there must be no event (u_2, v_2, t') in the temporal graph for which $t_1 \leq t' \leq t_2$.

Here we investigate how the constrained dynamic graphlets impact the spectrum of observed motifs and give a comparison against the temporal motifs without such restriction. Note that this restriction is mainly motivated by the snapshot-based representation, which implies that there are multiple events occurring at the same timestamp. For most real-world temporal networks, the raw data has a fine resolution (e.g., 1 second) and it is unlikely to observe multiple events occurring at the same timestamp. We also show this in Table 2 in the second last column: more than 97 percent of the events have unique timestamps in most datasets. Hence, nearly all motifs would be able to escape from the constrained dynamic graphlet restriction since almost always only one event occurs in a give timestamp. Thus we degrade the resolution of our datasets to 300s in order to highlight the difference between constrained dynamic graphlet counting and vanilla temporal motif counting (i.e., without constraints). Note that this will have an additional impact on both types of the motif counts: each will have less motifs when compared to 1s resolution (This is because we assume the events in a motif have a total ordering, so events within the same timestamp are not included in the same temporal motif). In both scenarios, we only consider $\Delta_C = 1500s$ as the timing constraint and investigate the counts of 3n3e motifs.

TABLE 4
The Motif Proportion Changes (percentage) When Going From Vanilla Temporal Motifs to the Constrained Dynamic Graphlets

Network	Variance	010102	010202	012020	010201
Calls-Cop.	1.70	+0.22%	-3.45%	+1.39%	-5.60%
CollegeMsg	3.36	+3.31%	+4.36%	+3.76%	-2.12%
SMS-Copen.	3.49	+2.37%	+3.23%	+3.26%	-0.99%
SMS-A	2.34	+4.09%	+4.98%	+1.54%	-1.93%
Email	18.98	-9.63%	-10.05%	+2.88%	-18.00%
FBWall	1.03	+1.06%	+1.09%	+0.75%	-0.78%
Bitcoin-otc	0.00	0.00%	0.00%	0.00%	0.00%
StackOver.	0.04	+0.26%	+0.27%	+0.26%	-0.09%
SuperUser	0.06	+0.63%	+0.65%	+0.24%	-0.14%

We first check the impact of degrading the resolution on vanilla temporal motif counts. Message networks are affected most; we observe 80 percent decrease in counts from 1s to 300s. Degrading the resolution has less impacts on the stack exchange networks, where the inter-event time intervals are larger than the other datasets. Comparing the ratios of each motif at 1s and 300s resolutions shows that the change of proportion is always less than 1 percent. Hence degrading the resolution affects all motifs equally, so does not introduce any bias into our evaluation.

Now we compare the two methods after the resolution is degraded to 300s. Table 4 presents the results (complete results are available in the supplementary materials), available online. We look at the changes in the ratios for all 3n3e motifs (ratio of a particular motif count to the sum). We observe the largest variance in Email, which indicates that some motifs change significantly than others. Message networks also show strong variance while the variances are less than 0.1 in the stack exchange networks. There are some commonalities across all datasets: 010102, 010202, 012020, and 010201 are the motifs that show the most significant proportion changes when going from vanilla counts to restricted counts. The decrease in 010201 is due to the fact that the constrained dynamic graphlet does not favor the *delayed repetitions*: if there exists many 01s after the 02, only the first one is considered and the rest is ignored. Note that it is often likely to see such delayed repetitions in communication networks such as messaging and email. The decrease in 010201 translates to increases in the motifs with *immediate repetitions*, such as 010102, 010202, and 012020. Note that Email shows a different behavior with respect to these motifs where most ratios decrease drastically. The common feature of the motifs with decreases (010102, 010202, 010201) is that there is a repetition and out-burst from the initial sender node (0), either one after the other or interleaved. We believe that the prevalence of out-bursts is due to the the carbon copies (cc) in emails. It is also likely that the carbon copies occur at the both timestamps of the repetition, thus those motifs are removed by the constrained dynamic graphlet restriction. *Using the constrained dynamic graphlet restriction can adversely impact the analysis of communication networks (such as text messages and emails) where two parties are engaged in a conversation.* Delayed repetitions are likely to occur when the sender is engaged in another conversation and ignoring those motifs will amplify the motifs with immediate repetitions.

TABLE 5
Counts of Event Pairs and Their Reduction Rates When Going From $only-\Delta_W$ to Δ_W -and- Δ_C and $only-\Delta_C$ Configurations

Network	Motif Type	$only-\Delta_W$		Δ_W -and- Δ_C		$only-\Delta_C$	
		Count	Ratio	Count	Ratio	Count	Ratio
College.	R, P, I, O	514K	81.9%	421K	81.9%	292K	56.8%
	C, W	68.3K	82.6%	56.4K	82.6%	40.2K	58.9%
FBWall	R, P, I, O	395K	79.7%	315K	79.7%	242K	61.3%
	C, W	45.9K	88.4%	40.6K	88.4%	32.7K	71.2%
Bitcoin.	R, P, I, O	8.91K	80.9%	7.21K	80.9%	5.94K	66.6%
	C, W	338	93.5%	316	93.5%	282	83.4%
SMS-Cop.	R, P, I, O	293K	82.1%	241K	82.1%	177K	60.3%
	C, W	31.1K	86.8%	27.0K	86.8%	21.0K	67.5%
SMS-A	R, P, I, O	894K	83.4%	745K	83.4%	561K	62.8%
	C, W	66.0K	88.8%	58.6K	88.8%	43.6K	66.1%

Summary. Overall, we observe that both the consecutive events restriction and constrained dynamic graphlets exhibit a bias towards certain types of motifs, consistently in most datasets. This can have an adverse affect in the analysis of certain types of temporal networks.

5.2 Impact of Timing Constraints

Both timing constraints (Δ_C and Δ_W) have their specialties in capturing temporal motifs: Δ_C emphasizes the continuity, while Δ_W provides a holistic view and brings a strict bound to the motif timespan. Here we use both parameters and compare the patterns we observed in the two extremes ($only-\Delta_C$ and $only-\Delta_W$). For all experiments, we set the Δ_W as 3000 seconds and change the Δ_C to obtain different Δ_C/Δ_W ratios. Our choice relies on the same principle used in [14]; we use inter-event times as proxy. As we discussed in Section 4.5, the Δ_C/Δ_W ratio must be in $(\frac{1}{m-1}, 1)$ interval to make the both constraints useful, i.e., we are considering $only-\Delta_C$ if the Δ_C/Δ_W ratio is too small and $only-\Delta_W$ if the ratio is too large. For three-event motifs we select three configurations: $\Delta_C/\Delta_W = 0.5$ ($only-\Delta_C$), 0.66 (Δ_W -and- Δ_C), and 1.0 ($only-\Delta_W$); for four-event motifs there are four configurations: $\Delta_C/\Delta_W = 0.33$ ($only-\Delta_C$), 0.5, 0.66, and 1.0 ($only-\Delta_W$). Since we only change the Δ_C here, some motifs in larger Δ_C/Δ_W configurations will not satisfy the Δ_C constraint when the ratio becomes smaller, thus the set of motifs observed under a smaller Δ_C/Δ_W ratio is a subset of a larger Δ_C/Δ_W configuration.

We first discuss the motif counts and spectrum under different timing constraints by using the event pairs (Section 5.2.1). Then, we look at the behavior of intermediate events (non-first, non-last) (Section 5.2.2). Last, we check how the motif timespans are shaped by the different timing parameters (Section 5.2.3).

5.2.1 Motifs Counts and Event Pairs

We compare the motif counts obtained with different timing constraint configurations and analyze the frequency and ratio of the event pairs (R, P, I, O, C, W) in 3n3e motifs.

We Observe That Considering $only-\Delta_W$ Amplifies the Number of R, P, I, O Motifs. Enforcing the Δ_C constraint helps to find less R, P, I, O motifs since consecutive events should be close to each other. Table 5 shows that R, P, I, O motifs are




Fig. 3. Ratios of event pairs in $only-\Delta_W$ and $only-\Delta_C$ configurations. Each figure shows a pair of pie charts for the ratio of event pairs in six categories. Fig. 3a has the distributions for 3e motifs in StackOverflow; Fig. 3b gives the ratios for 4e motifs in Calls-Copenhagen. The proportion of repetitions decreases in almost all datasets when going from $only-\Delta_W$ to $only-\Delta_C$ while the ones with increasing ratios are different.

overrepresented in the situation where Δ_C is ignored. The number of R, P, I, O motifs is 10 times greater than C, W motifs in all datasets. The number of R, P, I, O motifs observed in $only-\Delta_C$ configuration is 40 percent less than the $only-\Delta_W$ configuration. Among all datasets message networks are affected most, where the number of R, P, I, O motifs is reduced to 56.8 percent in CollegeMsg and 60.3 percent in SMS-Copenhagen. The number of C, W motifs is also reduced when switched to $only-\Delta_C$ configuration, but the reduction rates are smaller than the R, P, I, O motifs. Using both constraints can achieve a balance between the two extremes, where the number of R, P, I, O motifs is reduced to nearly 80 percent and the C, W motifs is reduced to 90 percent.

We also compare the ratios of event pair types between $only-\Delta_C$ and $only-\Delta_W$ configurations for all three-events (3e) (two- or three-nodes) and four-events (4e) (two-, three-, or -four nodes) motifs. Fig. 3 presents representative results; for StackOverflow (3e motifs) and Calls-Copenhagen (4e motifs) (complete results are available in the supplementary materials), available online. In most datasets, we observe that the ratio of repetitions decreases when going from $only-\Delta_W$ to $only-\Delta_C$ configuration whereas the increases show variety. The ratio of repetitions for 3e motifs in StackOverflow is reduced from 18.0 percent in $only-\Delta_W$ to 16.4 percent in $only-\Delta_C$ (Fig. 3a). It is similar for 4e motifs; in Calls-Copenhagen, the repetitions decrease from 15.7 percent in $only-\Delta_W$ to 12.5 percent in $only-\Delta_C$ (Fig. 3b). The decrease in repetitions also means an increase in other types of event pairs. For stack exchange interactions, the ratio of in-bursts increases: for 3e motifs in StackOverflow, the in-bursts increases from 22.5 percent in $only-\Delta_W$ to 25.6 percent in $only-\Delta_C$. As interactions in StackOverflow are answers and comments, this indicates that new posts are often answered and commented by many different persons in a short time period, and we can have better grasp of the in-bursts of answers and comments by bringing the Δ_C constraint. For 4e motifs in Calls-Copenhagen, on the other hand, the proportions of ping-pongs and conveys increase when going from $only-\Delta_W$ to $only-\Delta_C$ configuration. This implies that in CDR networks ask-reply and message delivering patterns tend to happen in a short time period, therefore become more prominent in the $only-\Delta_C$ configuration.

5.2.2 Intermediate Event Behaviors

Next, we investigate the intermediate event behaviors under different configurations. Since Δ_W only limits the

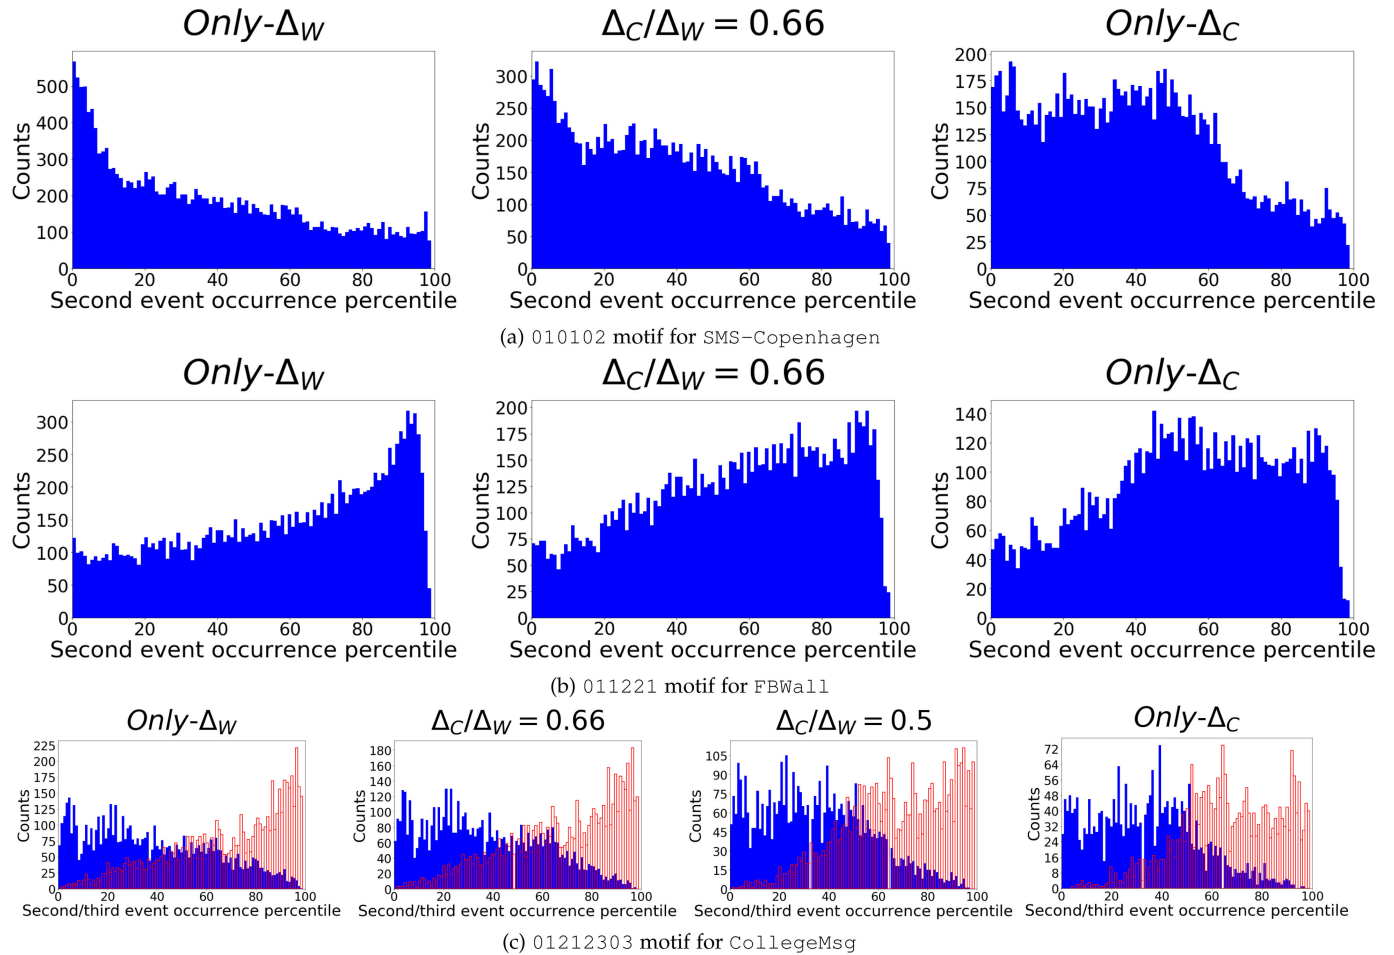


Fig. 4. Behavior of intermediate event occurrences. In each figure, the x -axis denotes the occurrence time of intermediate events with respect to the first and last events. 0 percent denotes the first event occurrence and 100 percent represents the last event occurrence. The y -axis shows the frequency of the intermediate events; second events (blue) in three-event motifs and second & third events (blue & red) in four-event motifs. In all cases, enforcing the Δ_C constraint regularizes the skew in *only- Δ_W* case.

interval between the first and the last events, it has no control on the behavior of intermediate events, i.e., when the second event happens in three-event motifs or when the second/third events happen in four-event motifs. Due to the bursty nature of temporal networks, the intermediate events in the motifs can be heavily skewed to the first event or the last event. Ideally, such bias is greatly reduced by Δ_C constraint, since it limits the time difference between consecutive events. Here we show a few representative motifs in different datasets for 3n3e and 4n4e motifs in Fig. 4 (all results are available in the supplementary materials), available online. Fig. 4a shows the intermediate event occurrence of motif 010102 in SMS (Copenhagen). In *only- Δ_W* case, the occurrence of the second event is significantly skewed to the first event since the second event is a repetition of the first. The lapse between the second and the last event can cover the entire timespan of the motif, and it is hard to guarantee the relevance of the last event in this scenario. *The skewness in the second event occurrence is reduced when $\Delta_C/\Delta_W = 0.66$, and the distribution is further regularized in *only- Δ_C* configuration.* Similar skewed patterns are also observed in the other 3n3e motifs. Fig. 4b displays the intermediate event occurrence of motif 011221 in FBWall. Due to the ping-pong behavior formed by the last two events, the second event occurrence is skewed to the last event for

only- Δ_W , and the bias is reduced in $\Delta_C/\Delta_W = 0.66$ and *only- Δ_C* configurations. We also observed similar patterns in 4n4e motifs. Fig. 4c shows the intermediate event occurrence of motif 01212303 in CollegeMsg. In *only- Δ_W* case, the occurrence of the second event is skewed to the first event and the third event is skewed to the last event, as both first-second and third-fourth pairs are in-bursts. *Tuning Δ_C/Δ_W from 1.0 to 0.33 regularizes the distribution of the occurrence and ensures that the timing of the intermediate event does not depend on the type of the motif.*

5.2.3 The Timespan of Motifs

Lastly, we compare the distributions of motif timespans (the time difference between the last and first events) for *only- Δ_C* and *only- Δ_W* configurations. Note that the Δ_C only gives a loose bound to the timespan of motifs ($\Delta_C * (m - 1)$ where m is the number of events), while the Δ_W gives a hard limit. Fig. 5 shows a representative result; the timespan of all 010102 motifs in CollegeMsg for different timing constraints (other results are in the supplementary materials), available online. In *only- Δ_C* , the distribution of motif timespans follows a normal distribution where the mean is approximately equal to the value of Δ_C . This indicates that, although Δ_C provides a loose bound, using *only- Δ_C* fails to




Fig. 5. Distribution of the motif timespans for 010102 motifs in ColleMsg. The x -axis denotes the timespan of the motif and the y -axis shows the count of such motifs. The distributions are more regularized when going from $only-\Delta_C$ to $only-\Delta_W$.

control the timespan of motifs. We observe that the motif timespans are more uniform in $\Delta_C/\Delta_W = 0.66$ configuration, and the distribution is further regularized in $only-\Delta_W$ configuration. Obtaining a set of motifs with uniform timespan distribution can be important in the analysis of real-world events which shows a variety in the timespan. For example, people have different churn behaviors in subscription services (some leave after a week and some others unsubscribe after a month) and one may be interested to see the temporal motifs that lead to the attrition. In such cases, selecting the motifs with uniform time distribution can enable to see the patterns that are related to customer's timeline rather than the absolute period of subscription.

Summary. We observe that Δ_C and Δ_W have complementary features; the former fails to bound timespans whereas the latter introduces bias for the occurrence of intermediate events. Combining both parameters by choosing a ratio for Δ_C/Δ_W in $(\frac{1}{m-1}, 1)$ interval can yield a trade-off.

5.3 Motifs as Sequence of Event Pairs

In this section, we use our new lens, event pairs, to analyze the characteristics of the sequences. In particular, we are interested in the temporal correlations among different types of event pairs and how the ordered sequences of event pairs show variety in various datasets. Fig. 6 presents the heat maps for all three-event motifs (both two- and three-nodes) in SMS-A, SMS-Copenhagen, Calls-Copenhagen, and Email using both timing constraints ($\Delta_C = 2000s$ and $\Delta_W = 3000s$); densities are color coded with respect to the min. and max. values in each dataset (all results are available in the supplementary materials), available online. Since there are two pairs of events in the three-event motifs, we represent those as an ordered sequence of two pairs in a heat map. Regarding the counts, we often observe the majority of motifs are formed by the sequences involving repetitions, while

only a few motifs are formed by the sequences including weakly-connected event pairs. This is also observed in the previous studies [14], [21]. The weakly-connected event pairs are rare since the two events take place in a non-transitive order, therefore are likely to be irrelevant.

In terms of the ordered sequences, we observe highly similar patterns in all message networks. The sequences involving both the repetitions and ping-pongs are the majority. Note that only two nodes are involved if there are only repetitions and ping-pongs; thus those interactions in message networks tend to be very local, and in most cases one-to-one conversations. Message networks also show a certain preference in the sequences involving in-bursts, out-bursts, and conveys. In-bursts and out-bursts seem to be incompatible as we do not observe many motifs that contain both. Conveys are often followed by out-bursts but not by in-bursts or conveys, and in-bursts are often followed with conveys, while the opposites rarely happen. We believe it is due to nature of information. The information delivered by conveys is likely to be distributed to multiple nodes through out-bursts. Meanwhile, the information gathered from multiple nodes through in-bursts is often further propagated through conveys. In Calls-Copenhagen and Email, there are less motifs formed by the sequences involving ping-pongs; however out-bursts are very dominant. Phone calls are not instantaneous events and happen for a duration; thus the information are mutually exchanged and repetitions are less likely to happen. Similarly, email communications do not have information limit when compared to the messages. Asymmetrical trends observed for message networks are also common for calls and emails.

Summary. Sequences of event pairs suggest interesting findings about the interplays among different types of pairs. In particular, the commonalities observed for message networks and the asymmetrical trends are surprising thanks to event pairs based analysis.

6 DISCUSSION

In this work, we introduced a comparative survey for the existing temporal motif models. We evaluated both advantages and limitations of these models with respect to two key aspects: temporal inducedness and timing constraints. In addition, we use the event pairs to analyze the sequences in the observed motifs. Our experimental evaluation shows that; (1) The temporal inducedness restrictions (consecutive

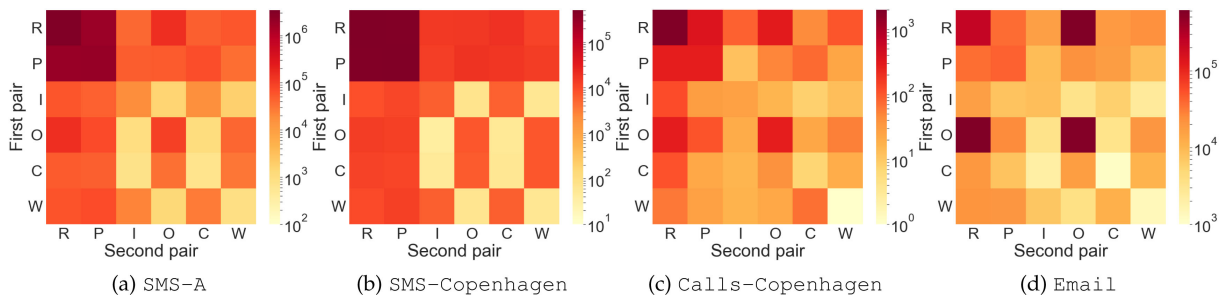


Fig. 6. Ordered sequences of event pairs for three-event motifs. Each block denotes a type of three-event motif, where y -axis shows the first pair of events (first and second event) and the x -axis shows the second pair of events (second and third event). The color indicates the motif counts in log scale and calculated with respect to the minimum and maximum counts in each dataset. Repetitions and ping-pongs are the most dominant in message networks, while repetitions and out-bursts are more common in calls and emails. Also, conveys, in-bursts, and out-bursts have asymmetrical trends in their compatibility.

events restriction and constrained dynamic graphlets) exhibit a bias towards certain types of motifs, consistently in most datasets, (2) Timing constraints, Δ_C and Δ_W , have complementary strengths, where the former fails to bound timespans whereas the latter introduces bias for the occurrence of intermediate events; hence combining both parameters by choosing a ratio for Δ_C/Δ_W in $(\frac{1}{m-1}, 1)$ interval for m event motifs can yield a trade-off; (3) Sequences of event pairs suggest interesting findings about the interplays among different types of pairs, such as asymmetrical trends.

There are several directions worth to explore as a future work. Besides temporal inducedness and timing constraints, for instance, temporal motifs with event durations is a promising avenue. We believe that devising an ultimate unifying model would be too ambitious due to the diverse characteristics of temporal networks and application-driven models (e.g., for phone calls) are likely to yield more effective and practical results. Also, understanding how the observed motif counts are related to the intrinsic network characteristics is important and recently introduced random reference models can be helpful [50]. We also intend to utilize the sequence of event pairs for the event prediction.


ACKNOWLEDGMENT

This research was supported by NSF-1910063 award, JP Morgan Chase and Company Faculty Research Award, and used resources of the Center for Computational Research at the University at Buffalo.

REFERENCES

- [1] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, "Network motifs: Simple building blocks of complex networks," *Science*, vol. 298, no. 5594, pp. 824–827, 2002.
- [2] N. Pržulj, "Biological network comparison using graphlet degree distribution," *Bioinformatics*, vol. 23, no. 2, pp. e177–e183, 2007.
- [3] C. Seshadhri, A. Pinar, and T. Kolda, "Triadic measures on graphs: The power of wedge sampling," in *Proc. 2013 SIAM Int. Conf., Data Mining*, 2013, pp. 10–18.
- [4] M. Jha, C. Seshadhri, and A. Pinar, "Path sampling: A fast and provable method for estimating 4-vertex subgraph counts," in *Proc. 24th Int. Conf. World Wide Web*, 2015, pp. 495–505.
- [5] A. Pinar, C. Seshadhri, and V. Vishal, "Escape: Efficiently counting all 5-vertex subgraphs," in *Proc. 26th Int. Conf. World Wide Web*, 2017, pp. 1431–1440.
- [6] N. K. Ahmed, J. Neville, R. A. Rossi, N. Duffield, and T. L. Willke, "Graphlet decomposition: Framework, algorithms, and applications," *Knowl. Inf. Syst.*, vol. 50, no. 3, pp. 689–722, 2017.
- [7] M. Bressan, F. Chierichetti, R. Kumar, S. Leucci, and A. Panconesi, "Counting graphlets: Space vs time," in *Proc. 10th ACM Int. Conf. Web Search Data Mining*, 2017, pp. 557–566.
- [8] S. Jain and C. Seshadhri, "A fast and provable method for estimating clique counts using Turán's Theorem," in *Proc. 26th Int. Conf. World Wide Web*, 2017, pp. 441–449.
- [9] P. Holme and J. Saramäki, "Temporal networks," *Phys. Rep.*, vol. 519, no. 3, pp. 97–125, 2012.
- [10] N. Masuda and R. Lambiotte, *A Guidance to Temporal Networks*. Singapore: World Scientific, 2016.
- [11] L. Kovananen, M. Karsai, K. Kaski, J. Kertész, and J. Saramäki, "Temporal motifs in time-dependent networks," *J. Stat. Mech.: Theory Exp.*, vol. 2011, no. 11, 2011, Art. no. P11005.
- [12] C. Song, T. Ge, C. Chen, and J. Wang, "Event pattern matching over graph streams," in *Proc. VLDB Endowment*, 2014, pp. 413–424.
- [13] Y. Hulovatyy, H. Chen, and T. Milenković, "Exploring the structure and function of temporal networks with dynamic graphlets," *Bioinformatics*, vol. 31, no. 12, pp. i171–i180, 2015.
- [14] A. Paranjape, A. R. Benson, and J. Leskovec, "Motifs in temporal networks," in *Proc. 10th ACM Int. Conf. Web Search Data Mining*, 2017, pp. 601–610.
- [15] M. Lahiri and T. Y. Berger-Wolf, "Mining periodic behavior in dynamic social networks," in *Proc. IEEE Int. Conf. Data Mining*, 2008, pp. 373–382.
- [16] L. Kovananen, M. Karsai, K. Kaski, J. Kertész, and J. Saramäki, "Temporal motifs," in *Temporal Networks*. Berlin, Germany: Springer, 2013, pp. 119–135.
- [17] V. Nicosia, J. Tang, C. Mascolo, M. Musolesi, G. Russo, and V. Latora, "Graph metrics for temporal networks," in *Temporal Networks*. Berlin, Germany: Springer, 2013, pp. 15–40.
- [18] P. Holme, "Modern temporal network theory: A colloquium," *Eur. Phys. J. B*, vol. 88, no. 9, 2015, Art. no. 234.
- [19] R. Jin, S. McCallen, and E. Almaas, "Trend motif: A graph mining approach for analysis of dynamic complex networks," in *Proc. Seventh IEEE Int. Conf. Data Mining*, 2007, pp. 541–546.
- [20] G. Chechik, E. Oh, O. Rando, J. Weissman, A. Regev, and D. Koller, "Activity motifs reveal principles of timing in transcriptional control of the yeast metabolic network," *Nat. Biotechnol.*, vol. 26, no. 11, pp. 1251–1259, 2008.
- [21] Q. Zhao, Y. Tian, Q. He, N. Oliver, R. Jin, and W.-C. Lee, "Communication motifs: A tool to characterize social communications," in *Proc. 19th ACM Int. Conf. Inf. Knowl. Manage.*, 2010, pp. 1645–1648.
- [22] S. Sarkar, R. Guo, and P. Shakarian, "Using network motifs to characterize temporal network evolution leading to diffusion inhibition," *Social Netw. Anal. Mining*, vol. 9, no. 1, pp. 1–24, 2019.
- [23] P. Bajardi, A. Barrat, F. Natale, L. Savini, and V. Colizza, "Dynamical patterns of cattle trade movements," *PLoS One*, vol. 6, no. 5, 2011, Art. no. e19869.
- [24] F. E. Faisal and T. Milenković, "Dynamic networks reveal key players in aging," *Bioinformatics*, vol. 30, no. 12, pp. 1721–1729, 2014.
- [25] D. Jurgens and T.-C. Lu, "Temporal motifs reveal the dynamics of editor interactions in wikipedia," in *Proc. 6th Int. AAAI Conf. Weblogs Social Media*, 2012, pp. 162–169.
- [26] L. Kovananen, K. Kaski, J. Kertész, and J. Saramäki, "Temporal motifs reveal homophily, gender-specific patterns, and group talk in call sequences," *PNAS*, vol. 110, no. 45, pp. 18070–18075, 2013.
- [27] X. Zhang, S. Shao, H. E. Stanley, and S. Havlin, "Dynamic motifs in socio-economic networks," *Europhysics Lett.*, vol. 108, no. 5, 2014, Art. no. 58001.
- [28] M.-X. Li *et al.*, "Statistically validated mobile communication networks: The evolution of motifs in european and chinese data," *New J. Phys.*, vol. 16, no. 8, 2014, Art. no. 083038.
- [29] Y.-Q. Zhang, X. Li, J. Xu, and A. V. Vasilakos, "Human interactive patterns in temporal networks," *IEEE Trans. Syst. Man Cybern., Syst.*, vol. 45, no. 2, pp. 214–222, Feb. 2015.
- [30] Q. Xuan, H. Fang, C. Fu, and V. Filkov, "Temporal motifs reveal collaboration patterns in online task-oriented networks," *Phys. Rev. E*, vol. 91, no. 5, 2015, Art. no. 052813.
- [31] Y. Li, Z. Lou, Y. Shi, and J. Han, "Temporal motifs in heterogeneous information networks," in *Proc. MLG Workshop@ KDD*, 2018.
- [32] C. Kosyfaki, N. Mamoulis, E. Pitoura, and P. Tsaparas, "Flow motifs in interaction networks," in *Proc. Int. Conf. Extending Database Technol.*, 2019, pp. 241–252.
- [33] S. Gurukur, S. Ranu, and B. Ravindran, "Commit: A scalable approach to mining communication motifs from dynamic networks," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2015, pp. 475–489.
- [34] R. Kumar and T. Calders, "2scent: An efficient algorithm for enumerating all simple temporal cycles," in *Proc. VLDB Endowment*, vol. 11, no. 11, pp. 1441–1453, 2018.
- [35] D. B. Johnson, "Finding all the elementary circuits of a directed graph," *SIAM J. Comput.*, vol. 4, no. 1, pp. 77–84, 1975.
- [36] X. Sun, Y. Tan, Q. Wu, B. Chen, and C. Shen, "Tm-miner: Tfs-based algorithm for mining temporal motifs in large temporal network," *IEEE Access*, vol. 7, pp. 49 778–49 789, 2019, doi: [10.1109/ACCESS.2019.2911181](https://doi.org/10.1109/ACCESS.2019.2911181).
- [37] X. Sun, Y. Tan, Q. Wu, J. Wang, and C. Shen, "New algorithms for counting temporal graph pattern," *Symmetry*, vol. 11, no. 10, 2019, Art. no. 1188.
- [38] P. Liu, A. R. Benson, and M. Charikar, "Sampling methods for counting temporal motifs," in *Proc. 12th ACM Int. Conf. Web Search Data Mining*, 2019, pp. 294–302.
- [39] U. Redmond and P. Cunningham, "Temporal subgraph isomorphism," in *Proc. IEEE/ACM Int. Conf. Adv. Social Netw. Anal. Mining*, 2013, pp. 1451–1452.

- [40] U. Redmond and P. Cunningham, "Subgraph isomorphism in temporal networks," 2016, *arXiv:1605.02174*.
- [41] P. Mackey, K. Porterfield, E. Fitzhenry, S. Choudhury, and G. Chin, "A chronological edge-driven approach to temporal subgraph isomorphism," in *Proc. IEEE Int. Conf. Big Data*, 2018, pp. 3972–3979.
- [42] M. Franzke, T. Emrich, A. Züfle, and M. Renz, "Pattern search in temporal social networks," in *Proc. 21st Int. Conf. Extending Database Technol.*, 2018, pp. 289–300.
- [43] CEP, 2020. [Online]. Available: wikipedia.org/wiki/Complex_event_processing
- [44] N. Pržulj, D. G. Corneil, and I. Jurisica, "Modeling interactome: Scale-free or geometric?," *Bioinformatics*, vol. 20, no. 18, pp. 3508–3515, 2004.
- [45] F. Hoffmann and D. Krasle, "Fraud detection using network analysis," *European Patent Office EP2884418A1*, Jun. 17, 2015.
- [46] Y. Wu, C. Zhou, J. Xiao, J. Kurths, and H. J. Schellnhuber, "Evidence for a bimodal distribution in human communication," *PNAS*, vol. 107, no. 44, pp. 18803–18808, 2010.
- [47] P. Sapiezynski, A. Stopczynski, D. D. Lassen, and S. Lehmann, "Interaction data from the copenhagen networks study," *Sci. Data*, vol. 6, no. 1, pp. 1–10, 2019.
- [48] J. Leskovec and A. Krevl, Jun. 2014, "SNAP Datasets," Stanford University. [Online]. Available: <http://snap.stanford.edu/data>
- [49] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi, "On the evolution of user interaction in facebook," in *Proc. 2nd ACM Workshop Online Social Netw.*, Aug. 2009, pp. 37–42.
- [50] L. Gauvin *et al.*, "Randomized reference models for temporal networks," 2018, *arXiv:1806.04032*.



Penghang Liu received the BS degree in geographic information science from Wuhan University, Wuhan, China, and the MS degree from University at Buffalo, where he is currently working toward the PhD degree in computer science. His research interests include temporal network motifs and dense subgraph discovery.



Valerio Guarrasi received the BC degree in management engineering and the MS degree in data science from the Sapienza University of Rome, Italy, where he is currently working toward the PhD degree in data science. His research interests include machine learning, with a focus on bio-medical applications.



Ahmet Erdem Sariyuce received the BS degree in computer engineering from Middle East Technical University and the PhD degree in computer science and engineering from Ohio State University. He is currently an assistant professor with the University at Buffalo. His research interests include graph mining and management.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.