

# Modularity-based approach for tracking communities in dynamic social networks

Michele Mazza, Guglielmo Cola\*, Maurizio Tesconi

*Institute of Informatics and Telematics, National Research Council, Via G. Moruzzi, 1, 56124, Pisa, Italy*

## ARTICLE INFO

### Keywords:

Community detection  
Community tracking  
Dynamic communities  
Social network analysis

## ABSTRACT

Community detection is a crucial task to unravel the intricate dynamics of online social networks. The emergence of these networks has dramatically increased the volume and speed of interactions among users, presenting researchers with unprecedented opportunities to explore and analyze the underlying structure of social communities. Despite a growing interest in tracking the evolution of groups of users in real-world social networks, the predominant focus of community detection efforts has been on communities within static networks. In this paper, we introduce a novel framework for tracking communities over time in a dynamic network, where a series of significant events is identified for each community. Our framework adopts a modularity-based strategy and does not require a predefined threshold, leading to a more accurate and robust tracking of dynamic communities. We validated the efficacy of our framework through extensive experiments on synthetic networks featuring embedded events. The results indicate that our framework can outperform the state-of-the-art methods. Furthermore, we utilized the proposed approach on a Twitter network comprising over 60,000 users and 5 million tweets throughout 2020, showcasing its potential in identifying dynamic communities in real-world scenarios. The proposed framework can be applied to different social networks and provides a valuable tool to gain deeper insights into the evolution of communities in dynamic social networks.

## 1. Introduction

With the rapid proliferation and evolution of various social networking systems, such as online social networks, mobile networks [1], and collaboration networks [2], social network analysis has emerged as a critical research topic. In this field, community detection plays a key role in uncovering meaningful group structures within networks, offering insights into the underlying social dynamics. Hence, various community detection algorithms have been proposed, considering factors such as individual cooperative or hostile relationships [3], and the utilization of local information to identify hidden communities [4]. Despite the significant attention that community detection has received in recent literature, most approaches have focused on static networks, overlooking the temporal properties inherent in real-world networks. Traditional methods rely on aggregation to build a static network that represents all the interactions over a specified period [5]. Inevitably, such representation is unable to capture the evolving nature of communities influenced by time-sensitive events. For example, in a Twitter retweet network [6,7] the evolution of community structures can be driven by factors like new users starting to retweet a specific account or existing community members ceasing their interactions. Treating these types of networks as static can lead to invalid associations and

an oversimplified representation of communities. A user may show similarities with a community at a certain time and later move towards a different community: this user would represent the only connection between the two communities, and omitting the temporal properties may lead to the aggregation of these communities as a single large community. In addition, time can also affect content semantically. For example, the meaning of the hashtag #MeToo changed dramatically with the emergence of the #MeToo social movement. Thus, community detection in such networks is crucial to enable a deeper understanding of the underlying dynamics.

A dynamic network can be represented as a time series of static networks called snapshots [8]. Each snapshot corresponds to the interactions aggregated over a defined period, such as a week or an hour. An intuitive method to detect communities in a dynamic network partitioned into snapshots consists in employing well-studied static community detection algorithms on each snapshot. Next, dynamic communities can be tracked by identifying the events that shape the evolution of a community over time [9–11]. This step involves matching the communities found at different snapshots through an algorithm, usually based on the similarity of community members. Therefore, an arbitrary threshold is required to determine if two or

\* Corresponding author.

*E-mail addresses:* [michele.mazza@iit.cnr.it](mailto:michele.mazza@iit.cnr.it) (M. Mazza), [guglielmo.col@iit.cnr.it](mailto:guglielmo.col@iit.cnr.it) (G. Cola), [maurizio.tesconi@iit.cnr.it](mailto:maurizio.tesconi@iit.cnr.it) (M. Tesconi).

more communities found at different snapshots are indeed the result of the evolution of the same community over time. The main drawback is that the final result may change drastically depending on the selected threshold: a high similarity threshold leads to less tolerance for fluctuating members, whereas setting a too-low similarity threshold leads to the merge of dissimilar communities. The scarcity of threshold-free frameworks negatively affects their real-world implementation, which could instead prove useful in countering threats that plague online social platforms such as disinformation [12,13] and information manipulation [14,15].

In this paper, we introduce a threshold-free framework crafted to track the evolution and structure of communities across multiple snapshots of a dynamic network. Given the communities detected across all snapshots, we represent the entire network as a similarity network, where the nodes are the static communities found and the weight of the edges corresponds to their similarity. By applying *Local Modularity Optimization* [16], we reveal the groups of nodes having high modularity and similarity. In this manner, we turn the community matching task into a modularity optimization problem where dynamic communities are found by optimizing modularity locally on all nodes. The node groups found embody the evolution of a community over time. Finally, each group is disentangled to reconstruct the temporal evolution of the represented dynamic community.

Notably, by constructing a single network encompassing all static communities across snapshots, our approach avoids sequential matching, which struggles to track intermittent communities that may not be present in consecutive snapshots. Through local modularity optimization, we derive groups of communities with robust interconnections and marked similarities, at the same time preserving a desirable degree of flexibility between them. Importantly, our threshold-free design guarantees consistent and trustworthy outcomes, promoting the application of our framework in real-world scenarios.

For the evaluation of our framework, we used four synthetic dynamic networks [9] embedded with distinct communities and events. Results indicate that, in most cases, our approach performs better than the state-of-the-art community tracking methods. In the second part of our evaluation, we show how the proposed approach can be applied to a real-world co-hashtag Twitter network [17] to extract relevant information about dynamic communities.

The main contributions of our work can be summarized as follows:

- We propose a novel framework for tracking the evolution and structure of communities in dynamic networks, leveraging local modularity optimization. The ability to track communities in dynamic networks is of paramount importance to fully understand the underlying social dynamics of real-world networks, like online social networks.
- Our framework offers a significant advantage over existing methods by eliminating the need for setting a specific threshold, leading to more accurate and robust tracking of dynamic communities.
- The evaluation of our novel framework on four synthetic datasets demonstrates its superior performance compared to the state of the art.
- We showcase the practical application of our framework on a Twitter dataset generated by real malicious accounts, illustrating its potential to extract relevant and suggestive information about dynamic communities.

The subsequent section offers a brief overview of existing research in dynamic community tracking. In Section 3, we provide foundational concepts pertinent to dynamic networks. Section 4 delves into a comprehensive description of our proposed framework. Section 5 encompasses both the evaluation and comparison of our framework against other methodologies and the presentation of results derived from a real-world network. Finally, Section 6 wraps up the paper with a concise summary and potential avenues for future research.

## 2. Related work

The inherent dynamic nature of most social networks [18] has spurred significant research interest in studying their evolution in recent years. In this section, we briefly review works that, similar to our approach, adapt existing static community detection algorithms to capture the temporal dynamics of social networks [19]. This approach is particularly suitable for social networks with highly dynamic community structures [20]. While here we provide an overview of the key contributions from each work, a detailed examination of selected studies is provided in Section 5.

One pioneering work that utilized static network snapshots to track community evolution over time is presented in [21]. The authors proposed a method based on agglomerative hierarchical clustering to identify and track stable clusters over time. In [22], the clique percolation method was extended to monitor events in the evolution of dynamic networks. They built joint graphs for pairs of consecutive snapshots and then matched the clique-based communities obtained using an autocorrelation function to find overlap between two states of a community. The work presented in [23] introduced a strategy, implemented as bit operations, to identify events between communities found in two consecutive snapshots. Another efficient approach to identify and track dynamic communities across multiple snapshots is described in [9]. The authors employed a community matching strategy based on weighted bipartite matching. A framework proposed by [20] offers an event-based approach for detecting transitions between communities in consecutive snapshots. In a later work from the same authors [10], the event definition formula has been improved to track community transitions throughout the observation time, no longer restricting it to consecutive snapshots. The Group Evolution Discovery (GED) framework, introduced by [11], takes into account not only the similarity of community members but also the positions and importance of nodes within the community. This approach facilitates matching communities and tracking their evolution across consecutive snapshots. More recent studies [24,25] proposed a method to detect and model the evolution of a community using a novel similarity measure termed “mutual transition”. In [19], a method for the Identification of Community Evolution by Mapping (ICEM) has been presented. In their approach, the evolution of a community is tracked by representing community members within a hash map.

We provide a more detailed description and analysis of [9–11,19,20,24,25] in Section 5, where we compare the strengths and weaknesses of these methods with respect to our novel framework.

## 3. Foundational concepts on dynamic networks

Before delving into our framework, let us first introduce foundational concepts pertinent to dynamic social networks.

### 3.1. Dynamic communities

A dynamic social network can be represented as a series of  $n$  graphs  $\mathbb{G} = \{G_1, G_2, \dots, G_n\}$ , where  $G_i$  represents the graph at the  $i$ th snapshot, encompassing only the nodes and edges present at that particular time. Next, we define as  $\mathbb{C}^i = \{C_1^i, C_2^i, \dots, C_{k_i}^i\}$  the set of  $k_i$  static communities found at the  $i$ th snapshot. Our objective is to find a set of  $m$  dynamic communities  $\mathbb{D} = \{D_1, D_2, \dots, D_m\}$  that occur in one or more snapshots. Each dynamic community  $D_j$  is represented by a timeline of its constituent static communities ordered in time. Fig. 1 shows an example of timelines for five dynamic communities. When dynamic communities, such as  $D_2$  and  $D_3$  or  $D_4$  and  $D_5$ , have interconnected static communities within them, we can either treat these dynamic communities independently or merge them into a single entity that encompasses all their constituent static communities. In the latter case, they become  $D_{(2,3)}$  and  $D_{(4,5)}$ .




Fig. 1. Example of five dynamic communities tracked over five snapshots, featuring growth, contraction, merging, splitting, and death events.

### 3.2. Critical events

In order to track communities and their structural changes over time, we need to define structural events to describe the evolutionary behavior of dynamic social networks. In [22,23], some types of events are proposed, but they do not cover all the possible ways in which a community can evolve. This paper uses the same models of critical events as proposed in [9]: *Growth* (a community gains new members), *Contraction* (a community loses some members), *Merging* (two or more communities merge into a new one), *Splitting* (a community is split into two or more new ones), *Birth* (a new community appears), *Death* (a community disappears).

Examples of these events are illustrated in Fig. 1:

- The community in  $D_1$  grows at the 2nd snapshot.
- The community in  $D_4$  contracts at the 3rd snapshot.
- The community in  $D_3$  merges with the community in  $D_2$  at the 4th snapshot.
- The community in  $D_4$  splits into two communities at the 5th snapshot, one still in  $D_4$  and the other in  $D_5$ .
- The communities in  $D_1$ ,  $D_2$ , and  $D_4$  were born at the 1st snapshot, whereas the community in  $D_3$  emerged at the 2nd snapshot and the community in  $D_5$  was formed at the 5th snapshot.
- The community in  $D_1$  experienced a death event at the 4th snapshot.

It is worth noting that a community may not manifest at every snapshot. Considering the example in Fig. 1, the community in  $D_2$  is observed at the 1st snapshot and then again at the 3rd snapshot. This “intermittence” can be caused by the behavior of community members or depend on the duration granularity of each snapshot.

### 3.3. Similarity

When tracking dynamic communities and the evolutionary events of their constituent communities, the key to finding relationships between communities from different snapshots is similarity. Most existing methods require setting a similarity threshold: communities are deemed related only if their similarity exceeds the threshold. In evolving dynamic social networks, unstable communities might undergo member shifts, losing original members and gaining new ones over time. Hence, setting the similarity threshold too high fails to identify these evolving

communities, while a threshold set too low may erroneously link unrelated communities. As detailed in the following section, our framework utilizes similarity without requiring a predefined threshold.

## 4. Proposed framework to track dynamic communities

Fig. 2 provides an outline of our framework. We initiate the process by segmenting the initial network into snapshots. Next, we apply a static community detection method to each snapshot. Subsequently, we construct the community similarity network, where each static community is represented as a node. In this phase, we disregard the origin of the static communities, enabling matches between communities from non-consecutive snapshots. We then identify dynamic communities by optimizing modularity locally across all nodes. Lastly, for each dynamic community, we reorder its constituent communities to reconstruct its temporal evolution and identify the events that occurred across snapshots.

### 4.1. Community similarity network

In our framework, the matching step relies on a weighted undirected community similarity network  $G_c = (V_c, E_c, W_c)$  where the vertex set (nodes)  $V_c$  encompasses the static communities found at each snapshot, the edge set  $E_c$  consists of links connecting these communities, and  $W_c$  assigns weights to these edges based on the similarity values of the connected communities. Notably, an edge is established between two communities only if they originate from distinct snapshots and share a non-zero similarity.

We compute similarity as the *overlap coefficient* between two communities. Let  $\mathcal{C}^i$  and  $\mathcal{C}^j$  denote the sets of communities found at the snapshots  $i$  and  $j$ , where  $i \neq j$ . Further, let  $\alpha \in \mathcal{C}^i$  and  $\beta \in \mathcal{C}^j$ . The overlap coefficient between two communities,  $\alpha$  and  $\beta$ , is calculated as follows:

$$overlap(\alpha, \beta) = \frac{|\alpha \cap \beta|}{\min(|\alpha|, |\beta|)} \quad (1)$$

where  $|\alpha|$  and  $|\beta|$  are the number of members, respectively, in  $\alpha$  and  $\beta$ , while  $|\alpha \cap \beta|$  is the number of shared members. The proposed framework is not tied to any specific similarity metric. For example, one could employ the Jaccard similarity, as shown in [9].

### 4.2. Local modularity optimization

After building the similarity network, we identify groups of similar nodes by optimizing modularity locally on all nodes. Modularity is a leading metric for evaluating community structures [26]. Scaled between  $[-0.5, 1]$ , it quantifies the density of intra-community links compared to inter-community links. By optimizing modularity, we seek a community assignment for each node in the network such that modularity  $M$  is maximized using the function defined by:

$$M = \frac{1}{2m} \sum_{xy} \left[ A_{xy} - \frac{k_x k_y}{2m} \right] \delta(c_x, c_y) \quad (2)$$

where  $A$  is the adjacency matrix with  $A_{xy}$  representing the weight of the edge between the vertices  $x$  and  $y$ ,  $k_x$  and  $k_y$  are the sum of the weights of the edges attached to vertices  $x$  and  $y$ , respectively,  $c_x$  and  $c_y$  are the communities to which vertices  $x$  and  $y$  belong,  $\delta$  is the Kronecker delta function ( $\delta(c_x, c_y)$  returns 1 if  $c_x = c_y$  and 0 otherwise), and  $m$  is the sum of the weights of all edges in the graph. In modularity optimization, edge weights, which represent the connection strength, guide the optimization process. Initially, each vertex in the network is assigned to its own community. Then, each vertex  $x$  is removed from its community and moved to the community of each neighbor vertex  $y$ .

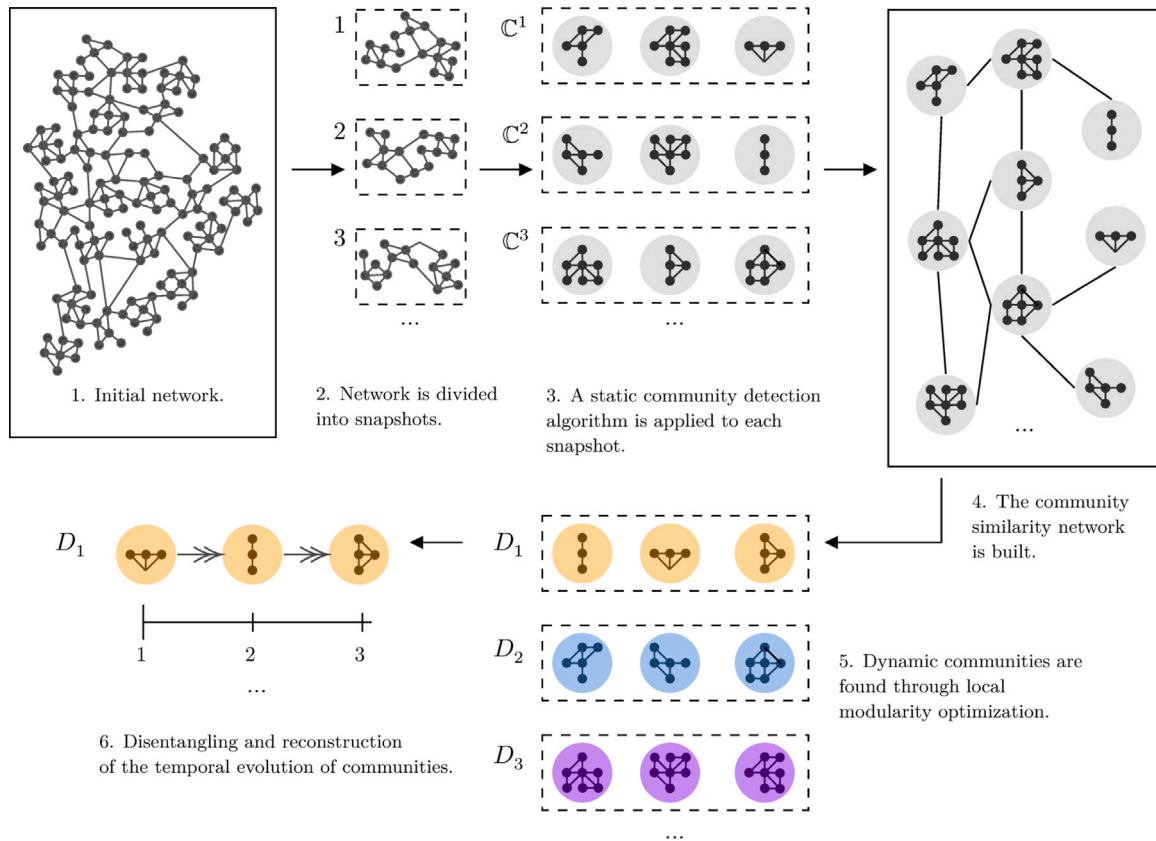


Fig. 2. Workflow of the proposed approach. Starting from a network, we divide it into snapshots. Then, we apply a static community detection algorithm to each snapshot. We construct the community similarity network using the communities found at the previous step. Dynamic communities are subsequently identified using local modularity optimization. Finally, for each dynamic community, we reorder its static communities based on their originating snapshots to highlight critical events.

When a vertex is assigned to a new community, the modularity increase is calculated:

$$\Delta M = \left[ \frac{\Sigma_{in} + 2k_{x,in}}{2m} - \left( \frac{\Sigma_{tot} + k_x}{2m} \right)^2 \right] - \left[ \frac{\Sigma_{in}}{2m} - \left( \frac{\Sigma_{tot}}{2m} \right)^2 - \left( \frac{k_x}{2m} \right)^2 \right] \quad (3)$$

where  $\Sigma_{in}$  is the sum of all the weights of the links within the community to which  $x$  is moving,  $\Sigma_{tot}$  is the sum of all the weights of the links to nodes in the community to which  $x$  is moving,  $k_x$  is the weighted degree of  $x$ ,  $k_{x,in}$  is the sum of the weights of the edges between  $x$  and the other nodes in the community to which  $x$  is moving, and  $m$  is the sum of the weights of all edges in the network. Then, once  $\Delta M$  is calculated for all communities connected to  $x$ , the community that resulted in the most significant increase in modularity encompasses  $x$ . If there is no increase in modularity, vertex  $x$  stays in its initial community. The process described above is applied to all vertices until there is no further increase in modularity. While our approach aligns with the initial phase of the Louvain algorithm [16], the latter then transitions to a second phase, building a new network where the vertices represent the communities found at the previous phase. Next, it applies the first phase to this new network, further optimizing modularity. A complete run of both phases is called a pass. Such passes are carried out repeatedly until a maximum of modularity is achieved. As our goal is to aggregate vertices with high similarity rather than maximizing modularity, we only employ the first step. Furthermore, because of the intrinsic properties of the community similarity network, using Louvain's complete procedure might yield communities with low granularity, potentially leading to clusters where certain members are poorly related to each other. Therefore, by solely applying local modularity optimization to the vertices of the community similarity network,

we obtain clusters with high modularity and thus composed of similar communities. These clusters correspond to the dynamic communities featured in the dynamic network.

#### 4.3. Identify dynamic communities over time

The pseudocode to identify and track dynamic communities over time is presented in Algorithm 1. First, the community similarity network  $G_c$  is initialized as an empty graph (Line 1). For each community within all snapshots, all subsequent snapshots are traversed to find communities with a similarity value greater than 0 (lines 2–7). When two communities exhibit a positive similarity, they are introduced as vertices in  $G_c$ , connected by an edge weighted by this similarity value (lines 8–9). Subsequently, a community set is initialized, where each node in  $G_c$  is assigned to its own distinct community (line 10). Local modularity optimization is then applied, evaluating each vertex to determine if its movement to another community would bolster network modularity (lines 11–20). The final result is a collection of sets of communities, each of these sets being an identified dynamic community. By reordering the communities in each set based on time, we can uncover the events that shaped the evolution of the dynamic community.

#### 4.4. Events reconstruction

Once the dynamic communities have been found from the community similarity network, we proceed with the identification of the events that describe their evolutionary behavior. Given a dynamic community  $D$ , which is a set of static communities found across snapshots, let us define more formally each critical event as follows:

**Algorithm 1:** Tracking dynamic communities

---

**Input :**  $S$  (List of sets of static communities contained at each snapshot)

**Output:**  $O$  (Collection of sets, each set being an identified dynamic community)

```

1  $G_c \leftarrow \{\}$ ; // initialize  $G_c$  with an empty network
  /* compare each community with all communities
  belonging to subsequent sets (snapshots) */
2 for  $C^i \in S$  do
3   for  $\alpha \in C^i$  do
4      $S^j \leftarrow$  all subsequent snapshots of  $C^i$ ;
5     for  $C^j \in S^j$  do
6       for  $\beta \in C^j$  do
7          $sim = overlap(\alpha, \beta)$ ;
8         if  $sim > 0$  then
9           add an edge in  $G_c$  between  $\alpha$  and  $\beta$ , with
             weight  $sim$ ;
/* optimize local modularity of  $G_c$  */
10  $O \leftarrow \{\{v\} | v \in G_c(V)\}$ ; // initialize each node of  $G_c$ 
    with its own community
11 while one or more nodes are moved to a different community
    do
12   for  $x \in G_c(V)$  do
13      $best_q \leftarrow -\infty$ ;
14      $best_c \leftarrow$  community of  $x$ ;
15     for all neighboring nodes  $y$  of  $x$  do
16        $gain_q \leftarrow \Delta M$  moving  $x$  to the community of  $y$ ;
17       if  $best_q \leq gain_q$  then
18          $best_q \leftarrow gain_q$ ;
19          $best_c \leftarrow$  community of  $y$ ;
20      $O \leftarrow$  place  $x$  in  $best_c$ ;
21 return  $O$ ;
```

---

- **Growth** – A community  $C^i \in D$  grows if there is a community  $C^j \in D$  at a prior snapshot ( $i > j$ ) that shares members with  $C^i$  and is smaller in size

$$growth(C^i) = 1 \quad \text{if}$$

$$\exists C^j \in D : \quad i > j, \quad |C^i \cap C^j| > 0, \quad |C^i| > |C^j| .$$

- **Contraction** – A community  $C^i \in D$  contracts if there is a community  $C^j \in D$  at a prior snapshot ( $i > j$ ) that shares members with  $C^i$  and is larger in size.

$$contract(C^i) = 1 \quad \text{if}$$

$$\exists C^j \in D : \quad i > j, \quad |C^i \cap C^j| > 0, \quad |C^i| < |C^j| .$$

- **Merging** – A community  $C^i \in D$  results from a merge if there is a set of communities  $C_{set}^j \in D$  from a prior snapshot ( $i > j$ ) such that each community in  $C_{set}^j$  shares members with  $C^i$

$$merge(C^i) = 1 \quad \text{if}$$

$$\begin{aligned} \exists C_{set}^j = \{C_1^j, C_2^j, \dots, C_n^j\} : \quad i > j, \\ \forall C^j \in C_{set}^j : \quad |C^i \cap C^j| > 0 . \end{aligned}$$

- **Split** – A community  $C^i \in D$  results from a split if it is part of a set of communities  $C_{set}^i \in D$  at snapshot  $i$  and there is a community  $C^j \in D$  from a prior snapshot ( $i > j$ ) such that each community in  $C_{set}^i$  shares members with  $C^j$

$$split(C^i) = 1 \quad \text{if}$$

$$\begin{aligned} \exists C_{set}^i = \{C_1^i, C_2^i, \dots, C_n^i\} : \quad i > j, \\ \forall C^i \in C_{set}^i : \quad |C^i \cap C^j| > 0 . \end{aligned}$$

- **Birth** – A community  $C^i \in D$  is born if there is no community  $C^j \in D$  from any prior snapshot ( $i > j$ ) that shares members with  $C^i$

$$birth(C^i) = 1 \quad \text{if}$$

$$\nexists C^j \in D : \quad i > j, \quad |C^i \cap C^j| > 0 .$$

- **Death** – A community  $C^i \in D$  dies if there is no community  $C^j \in D$  from any subsequent snapshot ( $j > i$ ) that shares members with  $C^i$

$$death(C^i) = 1 \quad \text{if}$$

$$\nexists C^j \in D : \quad j > i, \quad |C^i \cap C^j| > 0 .$$

## 5. Experiments and results

In this section, we detail the experiments conducted to compare our framework with state-of-the-art methods. Additionally, we present an example utilizing real-world data to demonstrate our framework's capability in extracting valuable insights from social interactions.

### 5.1. Synthetic datasets

To evaluate our framework, we employed the synthetic networks derived from the four *benchmark datasets* proposed in [9]. These synthetic networks have been constructed through a dynamic extension of the static LFR benchmark [27] to model different types of evolutionary behaviors of communities over time. Each dataset, consisting of 15,000 vertices, contains five static networks (snapshots), simulating evolving communities. For each snapshot, the ground truth about static communities is available. In each of the four synthetic networks, nodes have an average degree of 20, a maximum degree of 40, and a mixing parameter,  $\mu = 0.2$ , that controls community overlap. Furthermore, at each snapshot, 20% of the nodes change their memberships, mirroring the natural migration of members between communities over time. The synthetic networks were designed to encompass all types of community evolution events:

- **BirthDeath**: 40 new communities are constructed to replace 40 existing communities.
- **ExpandContract**: 40 randomly selected communities expand or contract their size by 25% at each snapshot.
- **MergeSplit**: 40 communities are randomly selected to be split, while another set of 40 undergoes merging, combining two communities into one.
- **Intermittent**: at each snapshot, 10% of the communities are intermittently concealed, making them unobservable.

### 5.2. Evaluation metric and experimental setup

To achieve a fair evaluation, we utilized the Louvain algorithm [16] as the static community detection approach for all the methods. We selected the hierarchical level that best matched the number of ground-truth static communities at the initial snapshot. To evaluate the performance in dynamic community tracking, we adopted an approach partially inspired by [9,28,29]. For each evaluated competitor, dynamic communities were found considering two scenarios: (i) using the ground-truth static communities at each snapshot; (ii) using the static communities derived by Louvain as mentioned above. Then, we employed Normalized Mutual Information (NMI) to compare the memberships to dynamic communities found in both scenarios.

NMI is a common entropy measure in information theory that quantifies the similarity between two clusters. It is defined as:

$$NMI(X, Y) = \frac{H(X) + H(Y) - H(X, Y)}{(H(X) + H(Y))/2} \quad (4)$$

where  $H(X)$  is the entropy of the random variable  $X$  associated with an identified community,  $H(Y)$  is the entropy of the random variable

**Table 1**  
NMI values obtained by the Greene framework.

Dataset	Threshold		
	0.1	0.3	0.5
BirthDeath	0.984	0.948	0.926
ExpandContract	0.986	0.940	0.921
MergeSplit	0.969	0.949	0.936
Intermittent	0.981	0.936	0.913

$Y$  associated with a ground truth, and  $H(X, Y)$  is the joint entropy. A higher NMI value, nearing 1, indicates that the dynamic community detection is more resilient against noise in the static community information at each snapshot. This resilience is crucial for ensuring reliable dynamic community tracking in real-world datasets.

For each dataset, we calculated the NMI in five sequential experiments, one per available snapshot. The first NMI value was found considering only the first snapshot, then for subsequent experiments we added one snapshot at a time. Hence, the last experiment employed the entire “timeline” made of five snapshots. As mentioned above, we compared our framework with other state-of-the-art approaches using the same Louvain-derived static community sets for unbiased comparison. Since some of these approaches require a threshold to be set, we tested them with different threshold values: (0.1; 0.3; 0.5). To select a threshold value, we calculated the average of the five NMI values for each dataset (see Tables 1–5). Then, for each approach, we selected the threshold yielding the highest average NMI.

### 5.3. State-of-the-art approaches

In this section, we describe the state-of-the-art approaches that we compared against our framework: *Greene* [9], *Takaffoli* [10,20], *Brodka* (*GED*) [11], *Tajeuna* [24,25], and *Mohammadmosaferi* (*ICEM*) [19]. Similar to our framework, these methods begin by analyzing the static communities detected at each snapshot using a community detection algorithm. Each of these approaches employs different strategies and similarity measures to track the evolution of communities over time, i.e., across snapshots.

*Greene et al.* [9] proposed a heuristic threshold-based method that enables many-to-many mappings between communities across different snapshots. The strategy proceeds as follows. First, an algorithm to find static communities is applied at each snapshot. Each community from the first snapshot is assigned to a dedicated dynamic community. Subsequently, communities at each snapshot are compared with each dynamic network’s front community. The front community of a dynamic network is the community found at the most recent snapshot belonging to that dynamic community. To perform matching, the authors used the Jaccard coefficient for binary sets [30]:

$$sim(\alpha, \beta) = \frac{|\alpha \cap \beta|}{|\alpha \cup \beta|} \quad (5)$$

A pair is considered a match if its similarity surpasses a threshold  $k$ . Additionally, it is assumed that a community appearing at snapshot  $t$  is deemed dissolved if it lacks a match for  $d$  consecutive snapshots. This condition allows the discovery of evolving communities at non-consecutive snapshots. In our evaluation of this framework, we assumed  $d = inf$  to enable matching with all non-consecutive snapshots. According to our experiments, the best result is obtained with  $k = 0.1$ .

*Takaffoli et al.* [10,20] used the following similarity measure to identify critical events in both consecutive and non-consecutive snapshots:

$$sim(\alpha, \beta) = \begin{cases} \frac{|\alpha \cap \beta|}{\max(|\alpha|, |\beta|)} & \text{if } \frac{|\alpha \cap \beta|}{\max(|\alpha|, |\beta|)} \geq k \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

**Table 2**  
NMI values obtained by the Takaffoli framework.

Dataset	Threshold		
	0.1	0.3	0.5
BirthDeath	0.941	0.950	0.921
ExpandContract	0.943	0.946	0.913
MergeSplit	0.921	0.928	0.909
Intermittent	0.940	0.941	0.903

**Table 3**  
NMI values obtained by the GED framework (Brodka et al.).

Dataset	Threshold		
	0.1	0.3	0.5
BirthDeath	0.933	0.931	0.923
ExpandContract	0.936	0.932	0.925
MergeSplit	0.925	0.923	0.917
Intermittent	0.926	0.923	0.918

The similarity threshold  $k$  was automatically determined using a text-mining approach since it was assessed on networks with content information. Given our use of synthetic networks, we could not determine the threshold automatically. From our experiments, we obtained the best results using  $k = 0.3$  for this method.

*Brodka et al.* [11] introduced the group evolution discovery (*GED*) framework, designed to identify overlapping communities. Unlike previous methods that solely relied on a similarity metric to track community changes, this approach extends the measure by incorporating a topological metric:

$$I(\alpha, \beta) = \frac{|\alpha \cap \beta|}{|\alpha|} \cdot \frac{\sum_{x \in (\alpha \cap \beta)} NI_{\alpha(x)}}{\sum_{x \in (\alpha)} NI_{\alpha(x)}} \quad (7)$$

where  $NI_{\alpha(x)}$  represents the importance of node  $x$  within the community  $\alpha$ . This value can be any centrality metric, e.g., degree centrality, social position, betweenness centrality, or PageRank. Although the comparison is made at consecutive snapshots, the inclusion effect helps to track overlapping and non-overlapping communities. The method requires two threshold values,  $k$  and  $j$ . We used the same value for  $k$  and  $j$  and obtained the best results for  $k, j = 0.1$ .

*Tajeuna et al.* [24,25] characterized each community using a *transition probability vector*. These vectors encapsulate the extent of shared membership between different communities across time. To assess similarity between communities, the method compares their respective transition probability vectors. Given two communities  $\alpha$  and  $\beta$  and their transition probability vectors  $v_{\alpha}$  and  $v_{\beta}$ , the similarity between the communities is calculated as:

$$sim(\alpha, \beta) = \begin{cases} \sum_{x=1}^{N_c} 2 \frac{p_{\alpha,x} \cdot p_{\beta,x}}{p_{\alpha,x} + p_{\beta,x}} & \\ \text{if } \sum_{x=1}^{N_c} 2 \frac{p_{\alpha,x} \cdot p_{\beta,x}}{p_{\alpha,x} + p_{\beta,x}} > k & \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where  $p_{\alpha,x}$  and  $p_{\beta,x}$  are components of the transition probability vectors  $v_{\alpha}$  and  $v_{\beta}$ , respectively, and  $N_c$  is the total number of communities in the dynamic network. The threshold value  $k$  is determined automatically as the approximated point of intersection of two Gamma curves. These curves derive from the non-zero similarity values between pairs of transition probability vectors.

*Mohammadmosaferi et al.* [19] introduced a novel method for Identification of Community Evolution by Mapping (*ICEM*). In this method, the members of each community are mapped using a hash map. This map associates members with a pair consisting of the snapshot and a community index. From the second snapshot, ICEM builds a similarity list for each community and determines the evolution of a community based on that list. Besides identifying common critical events, the

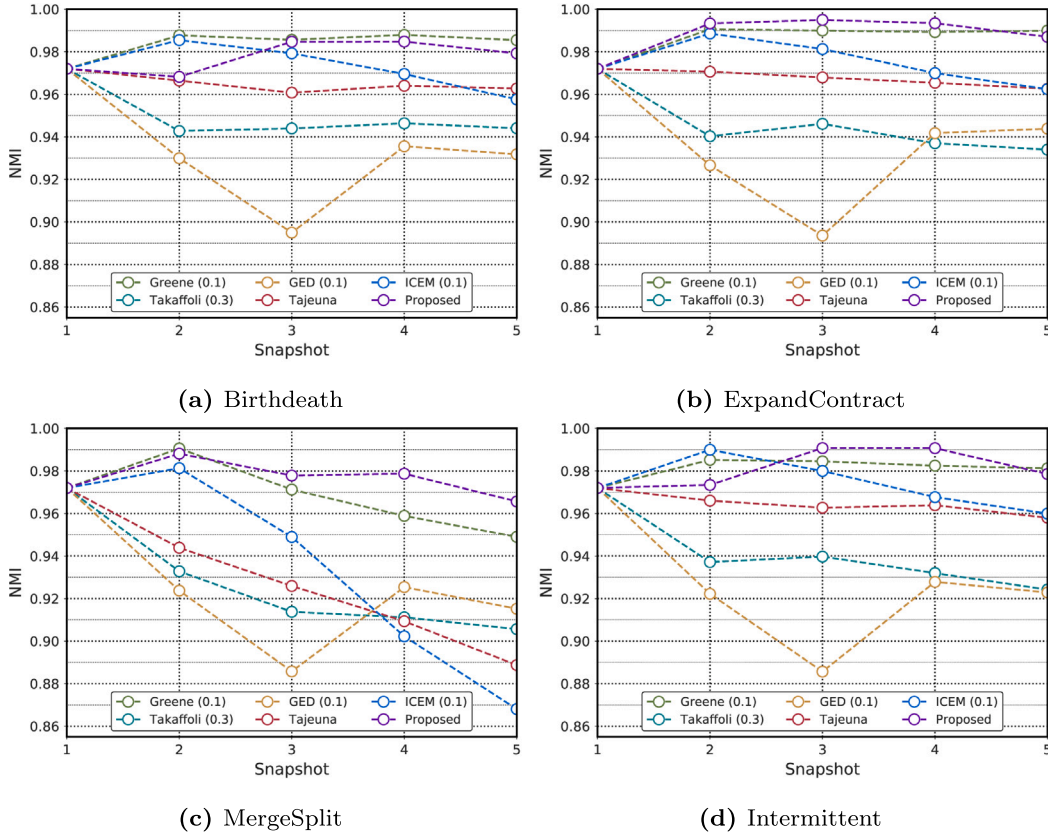


Fig. 3. NMI of the proposed dynamic community tracking framework and state-of-the-art methods on the four synthetic datasets.

Table 4

NMI values obtained by the Tajeuna framework.

Dataset	Threshold
	Automatically set
BirthDeath	0.965
ExpandContract	0.968
MergeSplit	0.929
Intermittent	0.965

Table 5

NMI values obtained by the ICEM framework (Mohammadmosaferi et al.).

Dataset	Threshold		
	0.1	0.3	0.5
BirthDeath	0.973	0.972	0.936
ExpandContract	0.975	0.970	0.930
MergeSplit	0.935	0.956	0.923
Intermittent	0.974	0.967	0.922

authors introduced the concept of *partial* events. These events are identified based on the partial similarity between two communities:

$$\text{sim}(\alpha, \beta) = \frac{|\alpha \cap \beta|}{|\alpha|} \quad (9)$$

$$\text{sim}(\beta, \alpha) = \frac{|\alpha \cap \beta|}{|\beta|} \quad (10)$$

where  $\alpha$  is a community from the  $i$ th snapshot and  $\beta$  a community from the  $j$ th snapshot, where  $i < j$ . The communities  $\alpha$  and  $\beta$  are partially similar if  $\text{sim}(\alpha, \beta) > k$  and  $\text{sim}(\beta, \alpha) > k$ , while they are very similar if  $\text{sim}(\alpha, \beta) > v$ . Therefore,  $k$  and  $v$  are the thresholds to identify partially similar and very similar communities, respectively. Since our concern is not to distinguish different types of critical events in this evaluation, we set  $v = 0.5$  and obtained the best results for  $k = 0.1$ .

Table 6

NMI values obtained by our framework.

Dataset	Snapshot				
	1	2	3	4	5
BirthDeath	0.971	0.968	0.985	0.985	0.979
ExpandContract	0.971	0.993	0.995	0.993	0.987
MergeSplit	0.971	0.988	0.977	0.979	0.966
Intermittent	0.971	0.973	0.991	0.991	0.979

#### 5.4. Results of the evaluation and comparison

Table 6 displays the NMI values achieved by our framework across all datasets and snapshots. Notably, it secured high NMI values (above 0.96) in every experiment.

Fig. 3 contrasts our results with the state-of-the-art methods on the four synthetic datasets. Our approach offers the best performance in most scenarios, although Greene and ICEM occasionally achieve a slightly higher NMI value. Overall, Greene, ICEM, and Tajeuna, exhibit strong results, whereas Takaffoli and particularly GED fall short in most experiments. GED's underwhelming performance may stem from its inability to link non-consecutive communities. Interestingly, all the evaluated methods achieve their worst performance with the *MergeSplit* dataset. Merge and split events significantly distort the dynamic community structures, making them challenging to track. Nevertheless, our method consistently achieves an NMI score above 0.96 even for this dataset. Among others, only Greene managed to retain NMI values above 0.94 in this dataset, while ICEM and Tajeuna show a degrading performance as the number of available snapshots increases. For scenarios with less complex evolutionary patterns, the efficacy of our method becomes more pronounced with the inclusion of more than two snapshots. A similar trend is also observed with Greene. With the expansion of the community similarity network, pinpointing analogous




Fig. 4. Number of dynamic communities found on the *MergeSplit* synthetic dynamic network, by the proposed dynamic community tracking method and state-of-the-art methods, relative to the ground-truth number of dynamic communities.




Fig. 5. Circles represent dynamic communities, plotted with the overall number of members on the  $x$  axis against the days of activity on the  $y$  axis. The color of each circle encodes its average hashtags overlap value.

community clusters with higher precision becomes more feasible. In contrast, some competing methods exhibit increased vulnerability to community matching errors as the snapshot count grows.

Fig. 4 offers a distinct perspective on the performance of the evaluated methods within the *MergeSplit* dataset. It contrasts the count of ground-truth dynamic communities in each snapshot with those identified by the various methods. Our method, together with Tajeuna and Takaffoli, closely mirrors the ground truth in terms of the number of dynamic communities. The discrepancy observed with Greene and ICEM could be attributed to the low threshold value employed, which results in increased community matches and subsequent merging. In contrast, GED, as previously mentioned, suffers from its limited comparison scope, restricted to consecutive snapshot communities.

In summary, the proposed framework excels in handling dynamic community tracking by focusing on local modularity optimization and setting aside the temporal component. This unique design choice allows the framework to adapt to structural changes in dynamic communities, consistently delivering strong results. Unlike most competitors, our approach operates without the need to set a threshold value. Although this entails the drawback that an unsatisfactory result cannot be altered by varying the threshold value, it is crucial to acknowledge that community tracking is a complex process with no universally correct outcomes. As such, the absence of a threshold can be seen as an advantage, preventing misleading adjustments. While methods that permit threshold adjustments might yield seemingly satisfactory results, they may not truly reflect the genuine evolution of dynamic communities. Additionally, by leveraging a denser community similarity network, our framework achieves enhanced performance when using more than two snapshots.

Table 7

Execution time of the proposed framework on networks with different characteristics.

Dataset characteristics			Time
snapshots	nodes	communities	
5	1,000	438	< 1 s
	10,000	4,369	21 s
	100,000	43,704	34 m 45 s
10	1,000	872	1 s
	10,000	8,738	1 m 33 s
	100,000	87,559	2 h 31 m 53 s

### 5.5. Complexity analysis

We excluded the computational cost of network representation and static community detection from our complexity analysis, as our framework builds upon communities derived from any given static community detection method. In our framework, each static community is compared with all others in subsequent snapshots, yielding a time complexity of  $O(c^2)$ , with  $c$  representing the total number of static communities. The local modularity optimization phase takes  $O(c)$ . Thus, the total time and memory complexity of the framework is  $O(c^2)$ . To assess the scalability of our framework, we employed the synthetic graph generator proposed by [9] to create progressively larger dynamic networks. Table 7 reports the properties of the datasets along with the corresponding execution times. We generated synthetic dynamic networks in two sets: one with five snapshots and another with ten. The number of nodes ranges from 1,000 to 100,000, which determines the number of static communities handled by the framework. We used a Python implementation of the proposed framework and executed the experiments using a single core of a machine with an Intel Xeon Processor (2.2 GHz) and 64 GB of RAM. For a dynamic network with 100,000 nodes across five snapshots, the processing time is under 35 min. However, with ten snapshots, it increases to approximately two and a half hours. The execution time is heavily influenced by the number of snapshots in the dynamic network, given its positive correlation with the count of static communities. Since most computations, such as determining the similarity between two communities, can operate independently, there is potential to boost computational performance through parallelization. Further improvements could be achieved by combining Minhashing [31] and Local Sensitive Hashing (LSH) [32] to identify approximate nearest neighbors in the Jaccard space. This would significantly reduce the number of comparisons, edging closer to linear time complexity.

### 5.6. Application to a real-world dataset

For our second evaluation, we assessed the behavior of the proposed framework using a real dataset described by [17]. This dataset comprises the activity of 63,358 fake Twitter accounts, which produced 5,457,758 tweets during 2020. The term “fake account” refers to social media accounts that contain false information or pretend to be a real person or organization [33]. This dataset is ideal for our experiments because the tweets were not confined to specific languages or topics, providing a broad range of content that included, for instance, politics and cryptocurrency [34]. We set snapshot granularity to one day of activity, and produced a weighted undirected graph for each snapshot. Daily graphs were derived from a co-hashtag network, in which two users are linked if they both used the same hashtag on a given day. The weight of each edge is given by the number of hashtags used in common between two users. To find the static communities at each snapshot, we used the Order Statistics Local Optimization Method (OSLOM) [35], which is based on measuring the significance of communities compared to a null model without community structure [36]. After OSLOM identifies statistically significant communities by grouping neighboring nodes, it undergoes multiple iterations of adjustments, like node



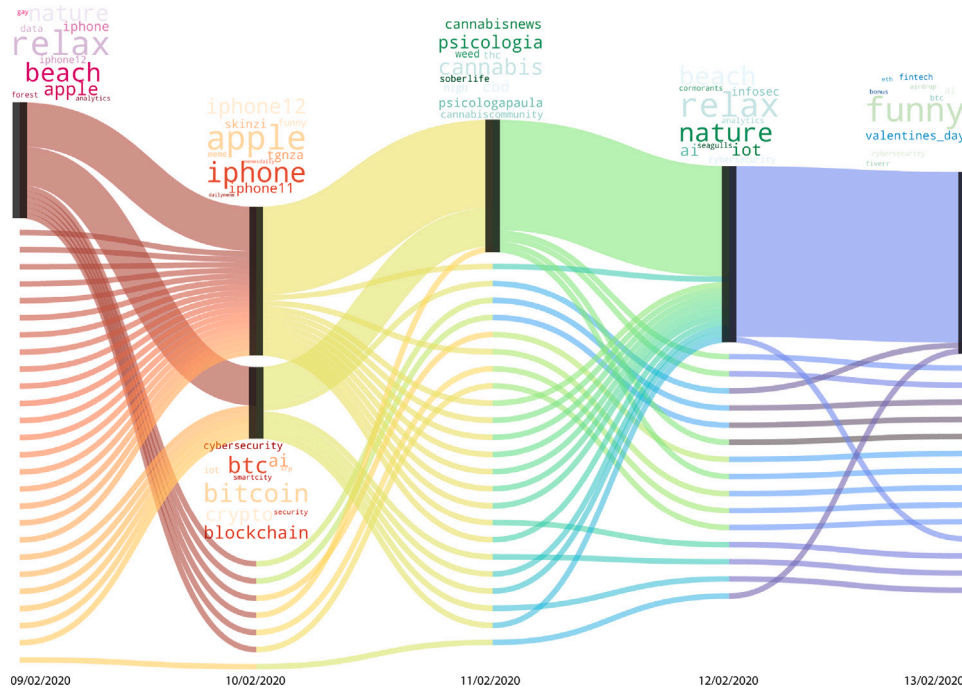


Fig. 6. Alluvial diagram representing the structure of the analyzed dynamic community.

addition or deletion, to enhance the significance of communities. We selected OSLOM due to its noted efficacy in detecting communities within online social networks [37,38]. A possible drawback of OSLOM is its tendency to identify smaller static communities [39], which may not properly capture the community structure of a network [40]. This concern is not relevant to the experiment, as the identified static communities are joined into bigger dynamic communities by the proposed framework.

We applied our framework to the static communities identified by OSLOM, obtaining 103 dynamic communities. To investigate the temporal behavior of the identified dynamic communities, we identified the hashtags used daily within each static community. Then, we introduced a new metric related to a dynamic community: the *average hashtags overlap*. Given two sets of hashtags  $h_1$  and  $h_2$ , the hashtags overlap is found as:

$$\text{Hashtags Overlap} = \frac{|h_1 \cap h_2|}{\min(|h_1|, |h_2|)}. \tag{11}$$

The average hashtags overlap for a dynamic community  $D$  is calculated by averaging the hashtags overlaps across all its constituent static communities. An acceptable alternative metric might have been the one used in [10], where instead of keywords or hashtags, the authors relied on topics. Given that this dataset captures the activity of sold fake accounts over a year, it is plausible that communities employed varied hashtag sets on different days. Despite this, a set of hashtags effectively represents a community on a specific day.

Averaging the hashtags overlap values for all static communities within a dynamic community allows us to gain insights into the characteristics of the identified dynamic communities. Fig. 5 plots the identified dynamic communities, with the overall number of members on the  $x$  axis and the days of activity on the  $y$  axis. Moreover, through color, we reported the average hashtags overlap of each dynamic community. This plot reveals a correlation between the number of members and days of activity; specifically, communities with more users tend to be active for more days, and vice versa. In addition, dynamic communities with high values in terms of members and days of activity also exhibit a low value of average hashtags overlap, meaning that during the one-year period, these dynamic communities discussed disparate topics. A low value of average hashtags overlap does not imply that a

dynamic community is not relevant, as this metric may be affected by other factors, such as the strategy employed.

Fig. 6 provides a glimpse into a dynamic community’s activity from February 9 to February 13, 2020. To better understand temporal behavior, for each constituent community we show the ten most used hashtags. Initially, there is only one community using mostly generic hashtags such as #relax, #beach, and #nature and some more specific ones such as #apple and #iphone. The next day, this community split into two communities, one of which turns out to be larger than the initial one due to the arrival of new users. This community resumes the hashtags used on the previous day regarding Apple products, introducing new ones such as #iphone11 and #iphone12. Instead, the other community focuses on cryptocurrency and blockchain technology through the hashtags #btc, #bitcoin, and #blockchain. On the following day, a portion of users from these two communities merge into a new community centered on hemp-related topics: #cannabis, #cbd, #cannabisnews. However, this community expands on the subsequent day, reverting to more generic hashtags. This behavior continues the next day, where we find #funny as well as #valentines\_day since it is Valentine’s Eve.

This example highlights the variability in topics that users of an online social network can address from one day to the next. Observing this kind of phenomenon and, more generally, studying the behavior of a community over time is only possible by taking into account its temporal properties, that is, by analyzing it as a dynamic community. By extending the framework’s application to similar datasets [41,42], analysts could gain new time-related insights into various phenomena on social media platforms, such as coordinated inauthentic behaviors.

## 6. Conclusions

In this paper, we introduced a framework designed to track the evolution of dynamic communities over time. Unlike most existing frameworks, our community matching phase operates without a threshold value. Instead, it employs modularity optimization applied to a weighted, undirected community similarity network, where the temporal component is omitted. We assessed our framework using synthetic graphs with embedded events and benchmarked our results

against other state-of-the-art frameworks. Our method excelled in handling structural changes in dynamic communities, yielding consistently strong results across diverse scenarios. The absence of a threshold further ensures the consistency of these results. Moreover, our framework is agnostic to the specific community detection algorithm used, offering flexibility in algorithm choice based on network properties, such as whether the network is weighted or unweighted, directed or undirected. We also conducted a preliminary evaluation of the proposed framework on a real-world Twitter network, uncovering 103 dynamic communities with distinct characteristics. This finding suggests that dynamic communities, particularly those derived from online social media, can show considerable variability in the topics they encompass. When applied to online social media analysis, our framework shows promising potential in identifying user groups with aligned interests or behaviors. By tracking dynamic communities, our approach can also enhance the temporal analysis of influencers and key opinion leaders, offering valuable insights for research applications.

### CRedit authorship contribution statement

**Michele Mazza:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Guglielmo Cola:** Writing – review & editing, Writing – original draft, Supervision, Conceptualization. **Maurizio Tesconi:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

Data will be made available on request.

### Acknowledgments

This work was supported by : the European Union – NextGenerationEU – National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) – Project: “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics” – Prot. IR0000013 – Avviso n. 3264 del 28/12/2021; project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU.

### References

- [1] B. Wu, Q. Ye, S. Yang, B. Wang, Group CRM: A New Telecom CRM Framework from Social Network Perspective, CNIKM '09, Association for Computing Machinery, New York, NY, USA, 2009, pp. 3–10, <http://dx.doi.org/10.1145/1651274.1651277>.
- [2] M.E.J. Newman, The structure of scientific collaboration networks, *Proc. Natl. Acad. Sci.* 98 (2) (2001) 404–409, <http://dx.doi.org/10.1073/pnas.98.2.404>.
- [3] C. Xia, Y. Luo, L. Wang, H.-J. Li, A fast community detection algorithm based on reconstructing signed networks, *IEEE Syst. J.* 16 (1) (2022) 614–625, <http://dx.doi.org/10.1109/JSYST.2021.3065378>.
- [4] Y. Luo, L. Wang, S. Sun, C. Xia, Community detection based on local information and dynamic expansion, *IEEE Access* 7 (2019) 142773–142786, <http://dx.doi.org/10.1109/ACCESS.2018.2878674>.
- [5] S. Wasserman, K. Faust, *Social Network Analysis: Methods and Applications*, Cambridge University Press, 1994, <http://dx.doi.org/10.1017/CBO9780511815478>.
- [6] M. Cinelli, S. Cresci, W. Quattrociocchi, M. Tesconi, P. Zola, Coordinated inauthentic behavior and information spreading on Twitter, *Decis. Support Syst.* 160 (2022) 113819, <http://dx.doi.org/10.1016/j.dss.2022.113819>.
- [7] P. Zola, G. Cola, M. Mazza, M. Tesconi, Interaction strength analysis to model retweet cascade graphs, *Appl. Sci.* 10 (23) (2020) <http://dx.doi.org/10.3390/app10238394>.
- [8] T.Y. Berger-Wolf, J. Saia, A framework for analysis of dynamic social networks, in: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '06*, Association for Computing Machinery, New York, NY, USA, 2006, pp. 523–528, <http://dx.doi.org/10.1145/1150402.1150462>.
- [9] D. Greene, D. Doyle, P. Cunningham, Tracking the evolution of communities in dynamic social networks, in: *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining, IEEE*, 2010, pp. 176–183, <http://dx.doi.org/10.1109/ASONAM.2010.17>.
- [10] M. Takaffoli, F. Sangi, J. Fagnan, O.R.Z. ane, Community evolution mining in dynamic social networks, *Procedia - Soc. Behav. Sci.* 22 (2011) 49–58, <http://dx.doi.org/10.1016/j.sbspro.2011.07.055>, *Dynamics of Social Networks*.
- [11] P. Bródka, S. Saganowski, P. Kazienko, GED: The method for group evolution discovery in social networks, *Soc. Netw. Anal. Min.* 3 (1) (2013) 1–14, <http://dx.doi.org/10.1007/s13278-012-0058-8>.
- [12] A. Martín, J. Huertas-Tato, Á. Huertas-García, G. Villar-Rodríguez, D. Camacho, FacTeR-Check: Semi-automated fact-checking through semantic similarity and natural language inference, *Knowl.-Based Syst.* 251 (2022) 109265, <http://dx.doi.org/10.1016/j.knsys.2022.109265>.
- [13] K.A. Qureshi, R.A.S. Malick, M. Sabih, H. Cherifi, Deception detection on social media: A source-based perspective, *Knowl.-Based Syst.* 256 (2022) 109649, <http://dx.doi.org/10.1016/j.knsys.2022.109649>.
- [14] G. Suchacka, J. Iwański, Identifying legitimate web users and bots with different traffic profiles — An information Bottleneck approach, *Knowl.-Based Syst.* 197 (2020) 105875, <http://dx.doi.org/10.1016/j.knsys.2020.105875>.
- [15] S. Li, J. Yang, G. Liang, T. Li, K. Zhao, SybilFlyover: Heterogeneous graph-based fake account detection model on social networks, *Knowl.-Based Syst.* 258 (2022) 110038, <http://dx.doi.org/10.1016/j.knsys.2022.110038>.
- [16] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, *J. Stat. Mech. Theory Exp.* 2008 (10) (2008) P10008, <http://dx.doi.org/10.1088/1742-5468/2008/10/p10008>.
- [17] M. Mazza, G. Cola, M. Tesconi, Ready-to-(ab)use: From fake account trafficking to coordinated inauthentic behavior on Twitter, *Online Soc. Netw. Media* 31 (2022) 100224, <http://dx.doi.org/10.1016/j.osnem.2022.100224>.
- [18] M.E.J. Newman, J. Park, Why social networks are different from other types of networks, *Phys. Rev. E* 68 (2003) 036122, <http://dx.doi.org/10.1103/PhysRevE.68.036122>.
- [19] K. Kadhoda Mohammadmosaferi, H. Naderi, Evolution of communities in dynamic social networks: An efficient map-based approach, *Expert Syst. Appl.* 147 (2020) 113221, <http://dx.doi.org/10.1016/j.eswa.2020.113221>.
- [20] M. Takaffoli, F. Sangi, J. Fagnan, O. Zaiane, A framework for analyzing dynamic social networks, in: *Proceedings of the 7th Conference on Applications of Social Network Analysis, ASNA*, 2010, URL <https://webdocs.cs.ualberta.ca/~zaiane/postscript/ASNA10.pdf>.
- [21] J. Hopcroft, O. Khan, B. Kulis, B. Selman, Tracking evolving communities in large linked networks, *Proc. Natl. Acad. Sci.* 101 (suppl 1) (2004) 5249–5253, <http://dx.doi.org/10.1073/pnas.0307750100>.
- [22] G. Palla, A.-L. Barabási, T. Vicsek, Quantifying social group evolution, *Nature* 446 (7136) (2007) 664–667, <http://dx.doi.org/10.1038/nature05670>.
- [23] S. Asur, S. Parthasarathy, D. Ucar, An event-based framework for characterizing the evolutionary behavior of interaction graphs, *ACM Trans. Knowl. Discov. Data (TKDD)* 3 (4) (2009) 1–36, <http://dx.doi.org/10.1145/1631162.1631164>.
- [24] E.G. Tajeuna, M. Bouguessa, S. Wang, Tracking the evolution of community structures in time-evolving social networks, in: *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics, DSAA, IEEE*, 2015, pp. 1–10, <http://dx.doi.org/10.1109/DSAA.2015.7344876>.
- [25] E.G. Tajeuna, M. Bouguessa, S. Wang, Tracking communities over time in dynamic social network, in: *Proceedings of the International Conference on Machine Learning and Data Mining in Pattern Recognition*, Springer, 2016, pp. 341–345, [http://dx.doi.org/10.1007/978-3-319-41920-6\\_25](http://dx.doi.org/10.1007/978-3-319-41920-6_25).
- [26] M. Chen, T. Nguyen, B.K. Szymanski, On measuring the quality of a network community structure, in: *Proceedings of the International Conference on Social Computing, IEEE*, 2013, pp. 122–127, <http://dx.doi.org/10.1109/SocialCom.2013.25>.
- [27] A. Lancichinetti, S. Fortunato, Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities, *Phys. Rev. E* 80 (2009) 016118, <http://dx.doi.org/10.1103/PhysRevE.80.016118>.
- [28] J. He, D. Chen, C. Sun, Y. Fu, W. Li, Efficient stepwise detection of communities in temporal networks, *Physica A* 469 (2017) 438–446, <http://dx.doi.org/10.1016/j.physa.2016.11.019>.
- [29] F. Folino, C. Pizzuti, An evolutionary multiobjective approach for community discovery in dynamic networks, *IEEE Trans. Knowl. Data Eng.* 26 (8) (2014) 1838–1852, <http://dx.doi.org/10.1109/TKDE.2013.131>.
- [30] P. Jaccard, The distribution of the flora in the alpine zone, *New Phytologist* 11 (2) (1912) 37–50, URL <http://www.jstor.org/stable/2427226>.
- [31] A. Broder, On the resemblance and containment of documents, in: *Proceedings. Compression and Complexity of SEQUENCES 1997 (Cat. No.97TB100171)*, 1997, pp. 21–29, <http://dx.doi.org/10.1109/SEQUEN.1997.666900>.

- [32] P. Indyk, R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality, in: Proceedings of the ACM Symposium on Theory of Computing, STOC '98, Association for Computing Machinery, New York, NY, USA, 1998, pp. 604–613, <http://dx.doi.org/10.1145/276698.276876>.
- [33] M. Mazza, M. Avvenuti, S. Cresci, M. Tesconi, Investigating the difference between trolls, social bots, and humans on Twitter, *Comput. Commun.* (2022) <http://dx.doi.org/10.1016/j.comcom.2022.09.022>.
- [34] G. Cola, M. Mazza, M. Tesconi, From tweet to theft: Tracing the flow of stolen cryptocurrency, in: F. Buccafurri, E. Ferrari, G. Lax (Eds.), Proceeding of the Italian Conference on Cyber Security, ITASEC 2023, in: CEUR Workshop Proceedings, (no. 3488) 2023, URL <https://ceur-ws.org/Vol-3488/paper11.pdf>.
- [35] A. Lancichinetti, F. Radicchi, J.J. Ramasco, S. Fortunato, Finding statistically significant communities in networks, *PLOS One* 6 (4) (2011) 1–18, <http://dx.doi.org/10.1371/journal.pone.0018961>.
- [36] M. Molloy, B. Reed, A critical point for random graphs with a given degree sequence, *Random Struct. Algorithms* 6 (2–3) (1995) 161–180, <http://dx.doi.org/10.1002/rsa.3240060204>.
- [37] D. Darmon, E. Omodei, J. Garland, Followers are not enough: A multifaceted approach to community detection in online social networks, *PLOS One* 10 (8) (2015) 1–20, <http://dx.doi.org/10.1371/journal.pone.0134860>.
- [38] W. Silva, Á. Santana, F. Lobato, M. Pinheiro, A methodology for community detection in Twitter, in: Proceedings of the International Conference on Web Intelligence, WI '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 1006–1009, <http://dx.doi.org/10.1145/3106426.3117760>.
- [39] P. Kumar, S. Gupta, B. Bhasker, An upper approximation based community detection algorithm for complex networks, *Decis. Support Syst.* 96 (2017) 103–118, <http://dx.doi.org/10.1016/j.dss.2017.02.010>.
- [40] C. Shi, Y. Cai, D. Fu, Y. Dong, B. Wu, A link clustering based overlapping community detection algorithm, *Data Knowl. Eng.* 87 (2013) 394–404, <http://dx.doi.org/10.1016/j.datak.2013.05.004>.
- [41] G. Cola, M. Mazza, M. Tesconi, Twitter newcomers: Uncovering the behavior and fate of new accounts through early detection and monitoring, *IEEE Access* 11 (2023) 55223–55232, <http://dx.doi.org/10.1109/ACCESS.2023.3282580>.
- [42] E. Chen, E. Ferrara, Tweets in time of conflict: A public dataset tracking the Twitter discourse on the war between Ukraine and Russia, in: Proceedings of the International AAAI Conference on Web and Social Media, Vol. 17, 2023, pp. 1006–1013, <http://dx.doi.org/10.1609/icwsm.v17i1.22208>.