



On analyzing user preference dynamics with temporal social networks

Fabiola S. F. Pereira¹ · João Gama² · Sandra de Amo¹ · Gina M. B. Oliveira¹

Received: 31 March 2017 / Accepted: 25 June 2018 / Published online: 23 July 2018
© The Author(s) 2018

Abstract

The preferences adopted by individuals are constantly modified as these are driven by new experiences, natural life evolution and, mainly, influence from friends. Studying these temporal dynamics of user preferences has become increasingly important for personalization tasks in information retrieval and recommendation systems domains. However, existing models are too constrained for capturing the complexity of the underlying phenomenon. Online social networks contain rich information about social interactions and relations. Thus, these become an essential source of knowledge for the understanding of user preferences evolution. In this work, we investigate the interplay between user preferences and social networks over time. First, we propose a temporal preference model able to detect preference change events of a given user. Following this, we use temporal networks concepts to analyze the evolution of social relationships and propose strategies to detect changes in the network structure based on node centrality. Finally, we look for a correlation between preference change events and node centrality change events over Twitter and Jam social music datasets. Our findings show that there is a strong correlation between both change events, specially when modeling social interactions by means of a temporal network.

Keywords Temporal networks · User preferences · Temporal node centrality · Change detection

Editors: Toon Calders and Michelangelo Ceci.

✉ Fabiola S. F. Pereira
fabiola.pereira@ufu.br

João Gama
jgama@fep.up.pt

Sandra de Amo
deamo@ufu.br

Gina M. B. Oliveira
gina@ufu.br

¹ Federal University of Uberlândia, Uberlândia, Brazil

² INESC TEC, University of Porto, Porto, Portugal

1 Introduction

Online social networks, such as Facebook, Twitter, and music social networks, facilitate the building of social relations among people who share similar interests. Users can stay connected with each others and be informed of new trends, consumption preferences and opinions of social friends. A natural process is that people tend to change their interest over time, specially in scenarios where they interact customarily with a wide range of items. At the same time, these social networks grow and also change quickly over time with the addition of new nodes and edges representing new interactions/relations in the underlying social structure.

The development of formalisms for preference specification and reasoning are essential tasks in literature since they can be used for sorting and selecting the objects that most fulfill user wishes. There are mining techniques for the automatic discovery of preferences and user profile building (de Amo et al. 2015), and there also exists research in the development of powerful mechanisms for preference reasoning (Wilson 2004). We are interested in *user preference dynamics*, i.e., the observation of how a user forms and evolves her preferences over time. A user preference is a specific type of opinion derived from comparative perception between two objects (Hansson 1995). For instance, when a user expresses “I don’t like to read about politics. Sports news are much better”, we clearly identify her preference to sports news over politics.

Just as preferences change over time, new links and nodes are continuously created on a wide variety of social networks as new users join the network, and new friendships are created. This leads to a number of important analysis such as event and anomaly detection (Aggarwal and Subbian 2014) in evolutionary networks analysis field. Indeed, key changes in the network structure often reflect individuals reaction to external events and trends (Arias et al. 2014). One can imagine that as the network evolves users evolve their social influence as well, which can directly result in changes to individual preferences. Recent research has made considerable advances towards the understanding of fundamental structural properties (Boccaletti et al. 2006), community structure (Oliveira et al. 2014), information diffusion (Guille et al. 2013), and social influence (Sun and Tang 2011) on online social networks. However, the impact of the online social networks on user preferences remains elusive. For example, little is known about whether and to what degree node centralities on social networks are related to user tastes and behavior changes (Althoff et al. 2017).

Motivated by the scenario described, this work aims at investigating the interplay between user preferences and social networks over time for systems personalization. We hypothesize that the evolution of user preferences is related to the evolution of her social network structure, specially when it comes to the detection of changes. The main research question we seek to answer is: is there a correlation between *preference change* events of a given user and *node centrality change* events in her evolving social network?

Motivating Example Let us consider a context concerning news that users like to read in everyday life. Suppose that analyzing the preferences of a given user A , we detect that on Aug 21st, A prefers to read about *politics* and *economy* than other news categories such as *sports* or *health*. Then, in a second moment, A ’s preferences remain stable, just appearing a preference of *politics* over *economy* news. However, in a third moment, on Aug 30th, we observe that A ’s preferences have changed and now, *economy* is preferred over *politics*. This situation is illustrated in the upper part of Fig. 1, where preferences are represented by better-than graphs (a directed edge (u, v) indicates that u is preferred over v). In the lower part of Fig. 1, snapshots of A ’s social network are represented. We notice that the network is




Fig. 1 Evolving perspective of A's temporal preferences (top) and A's temporal social network (bottom). Preferences are represented by better-than graphs where an edge (u, v) means that u is preferred over v . In the network, nodes are Twitter users and an edge (a, b) means that b retweeted a . On 08/30 there was a significant preference change. At the same time, A significantly changed her structural position (node centrality) on the network

also evolving with nodes appearing, disappearing, associating and disassociating with each other as time flies. In the network, nodes are Twitter users and a directed edge (x, y) means that x retweeted¹ y , i.e., the information flow. We conjecture that many aspects of A's social network can influence on A's preferences evolution. For instance:

- Around Aug 27th, A was being influenced by users who also like *politics*.
- From 27th to Aug 30th, a new connection with an influential personality in *economy* may have appeared and influenced A.
- A is always in contact with people who like *sports*.

It is an essential point to detect and predict A's preferences evolution and changes over time. We show in this paper that the temporal-topological social network structure of a given user is strongly correlated with her preference dynamics. According to our findings, by just observing A's social network evolution, we could increase the assertiveness of a news recommendation system for example, when recommending economy instead of politics news to A from Aug 30th.

In order to investigate the interplay between temporal dynamics of user preferences and her social network evolving over time we analyze the social network as a *temporal network*, where the times when edges are active are an explicit element of the representation (Holme and Saramaki 2012). Considering the order that social interactions occurred lead us to a more realistic model than just analyzing an aggregate static network that ignores when these contacts occurred. We carry experiments based on the following workflow: given a set of users, their preference traces over a given domain and their interactions with each other (social interactions), we (1) infer users preferences over the domain, (2) model an evolving social network based on users interactions and (3) for a given user, look for a correlation between her preference changes and her structural position changes on the network. Finally, we (4) check if the significant correlation extends to all users. Our correlation findings open doors to prediction/recommendation tasks over users tastes based only on the observation of the interactions between them in their social network.

Main contributions The main contributions of this paper can be summarized as follows: (1) proposal of a temporal preference model for representing and reasoning with preferences

¹ Retweet is to share some content originally posted by someone else in Twitter.

over time; (2) a preference change detection algorithm; (3) formalization of the node event detection problem based on node centrality changes; (4) a node event detection algorithm; (5) a set of experiments validating our proposals and finding that there is a correlation between preferences and centrality measures in temporal networks, specially against static networks counterpart.

Organization of the paper The paper is organized as follows. In Sect. 2 we discuss the state of the art in preference dynamics, social networks analysis and temporal networks fields. Section 3 describes user preference dynamics defining a temporal preference model and proposing a preference change detection strategy. In Sect. 4 we propose the use of temporal social networks and define centrality-based metrics to detect changes in nodes position on the network. Section 5 describes our methodology focusing on the preference mining strategy used to extract preferences from the social network content. Section 6 presents a rich experimental evaluation conducted over two datasets, Twitter and This Is My Jam, aiming at investigating the correlation between changes in preferences and changes in node centralities. Finally, Sect. 7 concludes the paper.

2 Related work

Our work is related to a number of research topics, including user preferences, social networks and temporal analysis in general. In literature, there are a lot of contributions combining these topics in pairs, specifically (i) temporal dynamics of user preferences, (ii) user preferences in social networks and (iii) temporal social networks. In this section we identify, organize and discuss the state of art. The originality of our proposal lies at the junction of these topics.

2.1 Temporal dynamics of user preferences

According to Liu (2015) modeling dynamics of preferences requires addressing two challenges: (i) precise preference representation and user profile building and (ii) accurate preference evolution inference.

Time-aware Personalized Recommendation Time-aware personalized recommendation systems generally represent preferences as feature vectors and consider the history of past profiles to predict preferences. Rafailidis and Nanopoulos (2014) proposed a measure of user-preference dynamics (UPD) that captures the rate with which the current preferences of each user has been shifted when providing recommendations. More recently, Liu (2015) also proposed to capture user's dynamic preference to provide timely personalized recommendation. The work of Wu et al. (2016) also deals with temporal behavior of preferences in the recommendation field. A network structure is used to model interactions among users and items rather than a utility matrix. In general, recommendation models assume that user transitions are driven by a static transition matrix. At present, the recommendation community lacks models that predict changes in user preferences (Kapoor et al. 2013).

Modeling Evolving Preferences In this line of research, the nature of preference evolution is studied seeking to describe them qualitatively (Thimm 2013), quantitatively (Sun et al. 2008), in a visual way through trajectories (Moore et al. 2013) and communities (Schlitter and Falkowski 2009) or predicting changes (Kapoor et al. 2013; Kapoor 2014). The latter is closer to our proposal.

Kapoor's works (Kapoor et al. 2013; Kapoor 2014) are the most expressive on predicting changes in user preferences. The idea is that predicting temporal choices is not a trivial task

just based on past behavior. Their approach is founded on psychology theories that state the presence of both stickiness and devaluation effects in user preferences. These studies, however, are orthogonal to ours. While in Kapoor et al. (2013) and Kapoor (2014) stickiness and boredom guided the preference change model, we seek to understand user preference dynamics founded on social influence.

Similar studies analyze the evolution of social action (Tan et al. 2010), sentiment change (Macropol et al. 2013) and user behavior (Zhang et al. 2014a) using social data, but not the evolution of user preferences. In Tan et al. (2010) the authors discuss how to simultaneously model the social network structure, user attributes and user actions over time. Examples of actions are whether a user discusses the topic “Haiti Earthquake” on Twitter or whether a user adds a photo to her favorite list on Flickr. Macropol et al. (2013) hypothesized that there is a strong relationship between users’ activity acceleration and topic sentiment change. Finally, in Zhang et al. (2014a, b) a generative dynamic behavior model is proposed. The model considers the temporal item-adoption behavior as a joint effect of dynamic social influence and varying personal preference over continuous time. Our approach also uses social influence in an environment of continuous preferences but the focus is on changes in preference.

2.2 User preferences on social networks

Works that join the topics of user preferences and social networks can be analyzed from three main perspectives: social recommendation, preference propagation and mining preferences from social networks. We describe related work from these perspectives, highlighting preference data. After all, what are these preferences that come from online social network (OSN)?

When considering time dimensions, research goes in the direction of opinion propagation and diffusion of preferences (Zhang et al. 2011; Lou et al. 2013). Generally, preferences are modeled from information outside the network and then the network structure is used in cascading models and influence detection of these preferences. They are not mined directly from the network.

Regarding preference mining, researchers seek to answer *how can we extract and model preferences from social networks?* In particular, given personal preferences about some of the social media users, how can we infer the preferences of unobserved individuals on the same network? Abbasi et al. (2014) is an approach that infers users’ missing attributes and preferences from networked data. We have proposed mining preferences from social media text using comparative sentences (Pereira 2015; Pereira and de Amo 2015). A comparative opinion is a statement like *car X is much better than car Y*, from which we can clearly extract a preference order: car X is preferred over car Y. Using a genetic algorithm we mined comparative sentences from tweets about PlayStation, Wii and XBox video games.

The preference mining task is generally formulated as user profiling, specially in recommendation models. There is a lack of specific techniques for mining preferences (preference order relation) purely from social networks.

2.3 Temporal social networks

The literature based on temporal networks (Holme and Saramaki 2012), focused on social media (Holme 2014), is essentially concentrated on understanding patterns of information diffusion by identifying key mediators and how temporal and topological structure of inter-

action affects spreading processes. The research in Pereira et al. (2016a); Wu et al. (2014) discusses the various concepts of shortest path for temporal graphs and proposes efficient algorithms to compute them. In Nicosia et al. (2013) graph metrics are revisited for temporal networks in order to take into account the effects of time ordering on causality. There are works addressing community detection in evolving networks (Rossetti et al. 2016; Cordeiro et al. 2016). Instead of focusing on local nodes, the literature is concentrated on the evolving patterns of groups of users in the network. Community detection is orthogonal to our proposal since we focus on nodes not on groups of nodes.

We highlight two main directions from this topic: Online Social Networks Event Detection and Event Detection over Dynamic Graphs. The former is related to social streams processing, where the message content being published on the network is analyzed (Aggarwal and Subbian 2012; Cordeiro and Gama 2016; Imran et al. 2016). For example, a set of posts sharing the same topic and words within a short time. The latter focus on events on the network structure evolution, for instance an increasing number of new connections on the social graph (Eberle and Holder 2016; Ranshous et al. 2015). Our proposal concentrates on the latter category: discovering events on evolving networks.

The most representative work in anomaly detection for dynamic graphs is Ide and Kashima (2004). It addresses the problem considering a time sequence of graphs (graph sequences). The focus is on faults occurring in the application layer of Web-based systems. First, they extract activity vectors from the principal eigenvector of a dependency matrix. Next, via singular value decomposition, it is possible to find a typical activity pattern (in $t - 1$) and the current activity vector (t). In the end, the angular variable between the vectors defines the anomaly metric. Akoglu and Faloutsos (2010) used this Eigen Behavior based Event Detection (EBED) method to detect events in SMS interactions – a *who-texts-whom network*. The main difference in comparison to ours is that it detects events in a global perspective of the network, while ours is node-centric.

3 User preference dynamics

In this section, we come to the task of defining our problem. What kind of user preferences we address and how do we handle temporal dynamics? As in Cadilhac et al. (2015), we distinguish preferences from opinions. Opinions represent a point of view that a person may have about an item; preferences involve an order relation that establishes a comparison between two items, often referred to as pairwise preferences. There are different ways to perceive how preferences of a given user vary over time. One can consider *novelty*, for instance the emergence of a new accessory in the fashion domain or new car models. Another way is considering *selectivity*, where the user becomes more or less restrictive in her preferences. And finally, the *changes*, i.e., whether in the past the user preferred some item and nowadays not anymore. Our approach focuses on the latter: *preference changes*.

3.1 Temporal preference model

There is no consensus concerning the definition of preference dynamics (Liu 2011). We adopt the following definition:

Definition 1 (*User Preference Dynamics (UPD)*) UPD refer to the observation of how a user evolves her preferences over time.




Fig. 2 Better-than graphs representing temporal preferences at days 1, 4 and 9. Edges inferred by transitivity are not depicted for better visualization

A preference is an order relation between two objects. For example, when a user says: “I prefer sports over politics”, if we order sports and politics in a ranking, we can clearly identify that sports will be at the top position.

Definition 2 (Temporal Preference Relation \succ_t) A temporal preference relation (or temporal preference, for short) on a finite set of objects $A = \{a_1, a_2, \dots, a_n\}$ is a strict partial order over A inferred at time t , i.e., a binary relation $R \subseteq A \times A$ satisfying the irreflexivity and transitivity properties at t . Typically, a strict partial order is represented by the symbol \succ . Considering \succ_t as a temporal preference relation, we denote by $a_1 \succ_t a_2$ the fact that a_1 is preferred to a_2 at t .

Definition 3 (Temporal Profile Γ_t^u) A temporal profile Γ_t^u is the transitive closure (TC) of all temporal preferences of user u at time t .

Example 1 Let $A = \{sports, tv, religion, music\}$ be the set of objects in our running domain representing themes of interest of user A. Figure 2 illustrates the temporal preferences of A at days 1, 4 and 9 through better-than graphs. Remark that an edge (a_1, a_2) indicates that a_1 is preferred to a_2 and edges inferred by transitivity are not represented. We have: $\Gamma_1^A = \{sports \succ_1 tv, tv \succ_1 religion, sports \succ_1 religion, sports \succ_1 music\}$, $\Gamma_4^A = \{sports \succ_4 tv, tv \succ_4 religion, sports \succ_4 religion, tv \succ_4 music, sports \succ_4 music\}$ and $\Gamma_9^A = \{sports \succ_9 tv, tv \succ_9 religion, sports \succ_9 religion, music \succ_9 tv, music \succ_9 religion\}$.

3.2 Detecting changes on temporal preferences

A key property of temporal preferences is irreflexivity. We say that a temporal profile Γ_t^u is *inconsistent* when there is a preference $a_1 \succ_t a_1 \in \Gamma_t^u$. It would mean that “I prefer X better than X!”, which does not hold for a strict partial order.

Our proposal for detecting preference change is based on the *consistency* of user temporal profiles. The idea is to compute the union of user profiles collected over time, infer temporal preferences by transitivity considering all timestamps and verify if there is any inconsistency in the resulting set of preferences. If yes, we detect an event of preference change. These concepts are formalized in the following.

Definition 4 (Temporal Profile Union Ω_t^u) Two temporal preferences of type $a_1 \succ_{t-1} a_2$ and $a_2 \succ_t a_3$, can unite to infer a third temporal preference $a_1 \succ_t a_3$, once considering transitivity of both, temporal preference relation and time order. A temporal profile union Ω_t^u is the transitive closure (TC) of all irreflexive relations given by

$$\Omega_t^u = \begin{cases} \Gamma_t^u & t = 1 \\ \Gamma_t^u \cup \Omega_{t-1}^u & t > 1 \end{cases} \tag{1}$$

Definition 5 (*Preference Change δ_t^u*) If there is a temporal preference inconsistency in Ω_t^u , a preference change has been detected at time t for user u . In other words, a preference change δ_t^u is defined as:

$$\delta_t^u = \begin{cases} 1 & \text{if there is a temporal preference inconsistency in } \Omega_t^u \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Remarking on Example 1, the temporal profile union $\Omega_9^A = \{..., tv \succ_4 music, music \succ_9 tv, tv \succ_9 tv, \dots\}$ contains the inconsistency $tv \succ_9 tv$. So, a preference change has been detected at time 9 ($\delta_9^A = 1$). Intuitively, we have that on day 1, for example, A prefers to read/post/share on her social network news about *sports*, but between tv and *religion* she is in the mood for tv . On the following days, A 's preferences practically do not change, just appearing a preference of tv over *music*. However, on day 9, A 's presented a preference change, as *music* became preferred over tv . Figure 10 illustrates a preference change event for a real Twitter user during 2016 Olympic Games.

3.3 PrefChangeDetection algorithm

In order to formalize the detection of changes in our temporal preference model, we propose the *PrefChangeDetection* algorithm. The intuition of this algorithm is to analyze *better-than graphs* (BTG) of a user during some observation period T . For each $t \in T$ we compute $BTG_t^u \cup BTG_{union}^u$, where BTG_t^u is the current BTG derived from the temporal profile Γ_t^u , and BTG_{union}^u refers to the temporal preferences in Ω_{t-1}^u , accumulated during the period $[t - |W|, t - 1]$, for W being a window over temporal profiles Γ^u . If the resulting $BTG_t^u \cup BTG_{union}^u$ (a temporal profile union Ω_t^u) has at least one cycle (meaning an inconsistency) we have detected a preference change at t . Algorithm 1 formalizes this idea.

On line 4, the union of two better-than graphs corresponds to Ω_t^u computation. The preference revision operation (line 7) consists of transforming BTG_{union}^u into acyclic by removing the oldest edges. We implement the strategy proposed in Cadilhac et al. (2015) to obtain a consistent and updated set of preference relations. According to Cadilhac et al. (2015) a *preference revision* is a sequence of two operations: *downdating* the existing preferences to a maximal subset that is consistent with the new preference, followed by adding the new preference to the result. So, new preferences take priority over old ones.

The size of the observation period determines if we are tracking short-term or long-term preference events. As example of real events, we can cite new product releases and special personal occasions such as birthdays (Xiang et al. 2010). The window W adjusts this feature. On line 11, updating BTG_{union}^u means forget preferences inferred before $t - |W|$.

Remarking on Algorithm 1 time complexity analysis, the time to build a better-than graph (line 3) is $O(P)$, where P is the number of temporal preference relations in Γ_t^u , which in the worst case is the combination $C_{|A|,2}$, for A being the finite set of objects in the domain (the nodes). On line 4, unifying two graphs costs $2O(|A| + |P|)$ where A and P are the set of nodes and edges, respectively. The time to detect if a directed graph is acyclic (line 5) is $O(|A| + |P|)$. Preference revision (line 7) takes $O(|A|)$, which is the time to compute a maximal independent set in graphs. The last operation is to perform a graph update (line 10) which in the worst case is $O(|A| + |P|)$. Hence, *PrefChangeDetection*, in the worst case, has complexity of $O(|T| \times 5(|A| + |P|))$, which is equivalent to $O(|T| \times |P|)$, for $|A| < |P|$.

Algorithm 1 PrefChangeDetection

Input: User u , a vector Γ^u of size $|T|$ containing u 's temporal profiles for each $t \in T$, a window W over Γ^u
Output: A vector δ^u of size $|T|$ containing u 's preference changes for each $t \in T$

```

1:  $BTG_{union}^u \leftarrow \emptyset$  //  $BTG_{union}^u$  is the accumulated better-than graph of  $u$ 
2: for all  $t \in T$  do
3:   build better-than graph  $BTG_t^u$  from  $\Gamma^u[t]$ 
4:    $BTG_{union}^u \leftarrow BTG_{union}^u \cup BTG_t^u$ 
5:   if  $BTG_{union}^u$  is not acyclic then
6:      $\delta^u[t] = 1$ 
7:     revise  $BTG_{union}^u$  //remove cycle maintaining more recent preferences
8:   else
9:      $\delta^u[t] = 0$ 
10:  update  $BTG_{union}^u$  according to  $W$  //remove from  $BTG_{union}^u$  all  $\succ_{t'}$  s.t.  $t' < t - |W|$ 
11: return  $\delta^u$ 

```

4 Evolving social networks

In this section, we explore how to track the structural evolution of a social network. Our contributions are three-fold. Firstly, we discuss that representing and consequently, analyzing social interactions as a temporal network can be more suitable to our dynamic scenario than using static networks concepts. Next, we propose the idea of node event detection, which is the action of detecting some remarkable fact from a node viewpoint in relation to the whole evolving network. The solution is based on change-points in node centrality values. Finally, we design an algorithm able to process the evolving network looking for node events.

4.1 Temporal networks versus static networks

We explore two different representations of social networks: as a static network and as a temporal network (Holme and Saramaki 2012; Pereira et al. 2016a). The static network structure is a traditional approach where temporal aspects are negligible and its evolution is analyzed just as a set of graphs snapshots over time (Pereira et al. 2016a). On the other hand, in temporal networks the information of when interactions between nodes happen is taken into account.

Definition 6 (*Static Network or Aggregate Network*) A static network $G_s = (V, E)$ is a set E of edges registered among a set of nodes V during an observation interval $[0, T]$. An edge between two nodes $u, v \in V$ is represented by $e = (u, v)$.

Definition 7 (*Temporal Network*) A temporal network (or temporal graph) $G_t = (V, E)$ is a set E of edges registered among a set of nodes V during an observation interval $[0, T]$. An edge between two nodes $u, v \in V$ is represented by $e = (u, v, t)$, where t ($0 \leq t \leq T$) is the time at which the contact occurred. Edges can also be called contacts.

Example 2 Consider the temporal and the aggregate networks in Fig. 3. They represent a social network, where nodes are users and edges are interactions (for example, retweets) between two users. Suppose that node A has a high impact information to spread on the network. If we analyze the network from the aggregate network perspective, the information will reach node F . This is not true for the temporal network, as A just interacted at time t_3 with B and after that, it is not possible to reach F from B .




Fig. 3 Temporal network versus static (aggregate) network. In the static network, an information can flow from *A* to *F*. In the temporal network, *A* will never reach *F* with any information

The analysis of centrality metrics is inherent to the particular network representation we are using. Remarking on Example 2, there is a path between nodes *A* and *F* on static network, but not on the temporal network. This implies in different values for the same centrality metric. The problem of evolving centralities in temporal networks is addressed in Pereira et al. (2016a). Basically, in static scenarios node centrality metrics like closeness and betweenness (Zafarani et al. 2014) are computed considering the concept of shortest paths on graphs. Moving to a temporal network representation, temporal node centrality metrics now should take into account the fastest paths. In Sect. 6, we show that *static betweenness* and *temporal betweenness* node centralities have different behavior patterns according to the network representation and consequently, they correlate with user preferences in different ways. The same is true for *static closeness* and *temporal closeness*.

4.2 Node event detection

Nodes behavioral dynamics are non-stationary, that is, they change or fluctuate over time. For instance, the structure induced by emails for a given user *u* may change during the working hours. Perhaps this user serves as a coordinator at work and therefore, during the day her email activity represents a structural behavior such as the center of a star (node with larger number of incoming or outgoing edges).

We are proposing the notion of node event detection, i.e., to spot change-points in an evolving network at which one node deviates from its normal behavior. A node event in the above mentioned email network can represent that *u* is responsible for a sudden bug in a critical system of the company. Detecting a node event is the action of detecting some remarkable fact or occurrence in someone’s life. Our proposal is based on change-points in node centrality values.

4.2.1 Change-point scoring functions

We introduce change-point scoring functions which take values between 0 and 1 where a higher value indicates a change-point. For all functions, we denote $C_t^m(v)$ the centrality metric *m* of a node *v* at time *t*, for *m* being any centrality measure like closeness, betweenness, degree, Katz, PageRank etc. We also consider a window *W* containing past summarized centrality values.

Average score Given a node *v*, we compare the current centrality value $C_t^m(v)$ with the arithmetic mean of the $|W|$ past centrality values inside the window *W*. Formally, we define the average score as

$$\Pi_t(v) = \frac{|C_{past}^m(v) - C_t^m(v)|}{\max(C_{past}^m(v), C_t^m(v))} \tag{3}$$

where $C_{past}^m(v)$ is the average of previous centrality values stored in W , defined as

$$C_{past}^m(v) = \text{avg}(C_{t-|W|}^m(v), \dots, C_{t-1}^m(v)) \quad (4)$$

The denominator factor from Eq. 3 is responsible for normalizing the average score. This is a baseline score very common to non-stationary analysis. Detecting events with the average score simply means that node v changed its role on the network in relation to its own previous behavior, but any additional information like the whole network or v 's neighbors is considered.

Ranking score This approach is founded on change-point in rankings (Wei and Carley 2015). The idea is to maintain a ranking R_t containing all the nodes on the network ordered according to their centrality metrics values for each time instant t . Based on the variation of these metric values and consequently ranking positions from recent past (past positions stored in window W) to current time, we detect changes. Formally, given the current set of nodes V , we consider a ranking R_t containing all nodes in V ranked in descending order according to their centrality values at time t . We define $pos_t(v)$ as the position of node v in R_t , i.e., $C_t^m(u) > C_t^m(v)$ iff $pos_t(u) > pos_t(v)$, for $u, v \in V$. The *ranking score* $\Lambda_t(v)$ is the acceleration of node v in the ranking position from the past to current instant time t :

$$\Lambda_t(v) = \frac{|pos_t(v) - pos_{past}(v)|}{\max(pos_t(v), pos_{past}(v))} \quad (5)$$

where $pos_{past}(v)$ is the average of previous positions of v in rankings $R_{t-|W|}, \dots, R_{t-1}$.

Here we consider the target node v centrality evolution in relation to the evolution of the other nodes on the network. By using this score, we can detect changes that are specific to v , evidencing its changing behavior in contrast to the continuous behavior of the remaining nodes. This is important to distinct cases of bursts, for example, where the whole network is impacted and not necessarily we have a specific node change-point. So, the *ranking score* remains stable.

Definition 8 (Node event) Given an evolving network $\mathcal{N} = (V, E)$ and a target node $v \in V$, a node event $\varepsilon_t(v)$ for v at time t is said to be occurred if the score for change-point detection is greater than the threshold θ . In other words, we have:

$$\varepsilon_t(v) = \begin{cases} 1 & \Theta_t(v) > \theta \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

for Θ assuming any of the change-point scores: Π (average) or Λ (ranking).

4.3 NodeEventDetection algorithm

The most common way of processing evolving networks is by assuming they are edge streams. To detect node events we use a sliding window strategy based on time instants. Thus, as time flies, the oldest stream objects are forgotten and only the most recent edges are considered for updating centrality values. In the following, we formally describe this process. Algorithm 2 is a sketch for detecting node events on an evolving network \mathcal{N} .

Definition 9 (Edge stream) Consider a time domain T as an ordered set of discrete time instants $t \in T$. An edge stream is a continuous and temporal sequence of objects $S = E_1 \dots E_r$, such that each object $E_i = (u, v, t)$ corresponds to an interaction (or a contact) from node u to node v at t , for $t \in T$.

Algorithm 2 NodeEventDetection

Input: Target node v , an edge stream $E_1 \dots E_r$ observed during $|T|$ time instants, a window W over the edge stream, threshold θ

Output: A vector ε^v of size $|T|$ containing v 's node events for each $t \in T$

```

1:  $V \leftarrow \emptyset, E \leftarrow \emptyset, \mathcal{N} = (V, E)$ 
2:  $t_{current} \leftarrow t$  //  $t$  is the first time instant
3: for each incoming edge stream object  $E_i = (u, z, t)$  do
4:    $E \leftarrow E \cup \{E_i\}$ 
5:    $V \leftarrow V \cup \{u, z\}$ 
6:   update  $C_t^m(v)$  or  $pos_t(v)$  according to the change-point scoring function  $\Theta$ 
7:   compute summary values for  $v$  at  $t$ 
8:   if  $t > t_{current}$  then
9:      $t_{current} \leftarrow t$ 
10:    if  $\Theta_t(v) > \theta$  then
11:       $\varepsilon^v[t] = 1$  //node event detected
12:    else  $\varepsilon^v[t] = 0$ 
13:    slides  $W$ 
14:     $E \leftarrow E - \{(a, b, t') | t' < t - |W|\}$  //refresh  $\mathcal{N}$ 
15:    refresh summary values
16: return  $\varepsilon^v$ 

```

Window strategy We adopted a sliding time-based window of temporal extent $|W|$ and progression step of 1 time instant $t \in T$. According to our definition, for the same discrete time instant t the edge stream can have many edge stream objects. For example, on a Twitter interaction network, considering 1-day time instants, we can receive several edge stream objects per day. This window strategy is a good choice as it allows for the detection of node events (i) without much processing effort, (ii) taking advantage of scoring functions semantics and (iii) considering the rapidly evolving characteristic of online social networks.

Remark that the window slides over two structures: edge stream objects and summary values. The stream objects are nothing more than the network evolving over time. Thus, having a sliding window over such objects means that centrality metrics used for event detection will always be calculated on an upgraded network, where old edges are discarded. In the same way, values summarized in memory during stream processing are being forgotten as they become older and leave the window cover. As we will present, the summarization is also done in function of time instants.

Computing centrality values On line 6 we update node centralities values in function of the new incoming edge. We follow the greedy strategy described in Pereira et al. (2016a) when computing temporal centralities or static centralities, depending on the network representation being considered.

Summarizing values Each change-point scoring function requires different statistics summarized in memory. But the idea is the same: maintain for each node $|W| + 1$ values according to the scoring function. For average score we maintain centralities values C_t^m and $|W|$ past values; for ranking score, ranking positions pos . In this way, line 7 calls a computation referent to current values (at t) and line 15 refreshes values by forgetting old statistics outside the sliding window and computing average *past* values.

Table 1 Summary of networks statistics

	Twitter network	Jam network
Domain	Brazilian news in Twitter	Social music network
Time span	08/08/2016–11/09/2016	08/26/2011–09/26/2015
# nodes	292,310	54,393
# temporal edges	1,392,841 (retweets)	1,667,335 (likes)
Avg static path length	12.31	7.63
Avg temporal path length	5 (day granularity)	2 (week granularity)

Remarking on *NodeEventDetection* time complexity analysis, the most costly operation is on line 6 when computing centrality values. From Pereira et al. (2016a), the cost is $O(2(V \times P))$, for P being the average number of paths between two nodes. For ranking score, there is the additional cost of $O(V \log V)$ to order the ranking. The costs for $\Theta_t(v)$ (line 10) and to refresh summary values (line 15) are negligible. To refresh \mathcal{N} (line 14) the complexity is $O(V + E)$. In the end, the average complexity for each incoming object is $O(2(V \times P) + V \log V + V + E)$. As $|V| > |P|$ in real-world networks, we have $O(V \log V + E)$, which is a high time consumption solution. In fact, we address this issue as future work (see Sect. 7).

5 Methodology

In order to correlate user preferences changes and node events in temporal social networks we need a dataset (1) containing the information of when links occur in the network (temporal network topology) and (2) some semantic information from which user preferences can be extracted (network content). We chose two datasets to perform experiments, one based on Twitter data and the second based on the social music website This Is My Jam.²

5.1 Twitter dataset

Folha de São Paulo (or Folha, for short) is one of the most influential newspapers in Brazil. Taking advantage of the fact that Twitter is widespread in the country, we performed our analysis over the news domain on the Twitter social network. We collected a large body of tweets from Folha over the course of 94 days. Our data collection strategy was as follows. First, we used Twitter's streaming API to collect all tweets related to the newspaper (user @folha). Thus, our dataset consists of tweets concerning the news tweeted by Folha, the retweets and all inherent information mentioning these news. Next, we built the following interaction network: nodes are Twitter users. An edge (u_1, u_2, t) represents that u_2 retweeted at t some text originally posted by u_1 .³ In all, we collected 1,771,435 tweets, 150,822 of which were retweeted at least once. Table 1 summarizes statistics of the crawled network.⁴

² <http://www.thisismyjam.com>.

³ We consider retweets and quote-status that are retweets with comments.

⁴ Dataset available at <http://www.lsi.facom.ufu.br/~fabiola/temporal-networks>.

Table 2 Examples of some topics identified by LDA from Twitter data and respective keywords manually assigned to them for better interpretability

Keyword	Topic (top-5 words)
Olympic games	Rio, olimpiada, brasil, jogos, metro
Lava Jato	Moro, lula, cunha, juiz, sergio
USA elections	Trump, hillary, eua, midia, federais
Lower house speaker	Cunha, camara, maia, presidente, governo
Odebrecht	Via, folha_com, odebrecht, caixa, milhoes

5.1.1 Extracting preferences

Probabilistic topic models such as LDA have been applied to extract and represent users' profile in different application scenarios, e.g., Web search and recommendation (Agarwal and Chen 2010; Liu 2015; Christidis et al. 2010). In this work we follow this trend to profile users by applying LDA as we do not have explicit preferences elicited in our dataset. Thus, in order to discover what users are talking about on the network we performed topic modeling with the LDA algorithm (Blei et al. 2003).

Every interaction (or retweet) between two users is associated with a textual content. We treat each such tweet (textual information) as a document, and the aggregation of all users' interactions considering the entire observation period forms a text corpus. Based on this corpus we perform LDA to extract 50 topics such that each document (tweet) is represented by a topic distribution. According to Wallach et al. (2009) choosing a larger k for LDA does not significantly affect the quality of the generated topics. The extra topics can be considered noise. However, choosing a small k may not separate the information precisely. Thus, we varied k from 20 to 80 and from empirical observations we selected $k = 50$ topics.

We analyzed the interpretability of the topics and manually assigned a keyword describing each topic. On Table 2 there are some examples of mined topics and their respective assigned keywords. Following this, we manually grouped these 50 keywords into 10 more general topics, as detailed on Table 3. The reason to group topics into more general ones is to provide better interpretability as these final 10 topics are the domain of preferences. Thus, $A = \{\textit{politics, international, corruption, sports, security, education, entertainment, economy, religion, others}\}$ is the set of objects in the domain on which we extract user preferences and each tweet is labeled with one object $o \in A$.

To extract pairwise preferences for each user we use the following strategy: if user u tweets (or retweets) about o at time t , then u has more interest in o over the remaining topics in domain at that moment. We also considered a weight $w_t^u(o)$ based on the number of tweets posted at the same time on a particular topic o . In this case, the top posted topic is preferred over others, the second top posted topic is preferred over the remaining ones and so on. Formally, we have: $\Gamma_t^u = \{o \succ_t^u o' \mid w_t^u(o) > w_t^u(o') \text{ and } o, o' \in A\}$. Noteworthy here is that the time t being considered depends on the time granularity in question, which can be of 1 day or 1 month, for instance. Therefore, a user can post many tweets at the same t .

Table 3 Manually grouping topic keywords into 10 more general topics

General topic	Keywords describing a topic
Politics	Pro-PT day, coup, lower house speaker, Dilma, Marina Silva, elections, Doria, Temer, INSS, PEC, strike
International	Venezuela, USA elections
Corruption	Moro, Lula, Odebrecht, lava jato, triplex, delação
Sports	Olympic Games, Football
Security	Violence, policy, popular manifestations
Education	High school, ENEM
Entertainment	Youtuber, book, show
Economy	Petrobras, inflation rate
Religion	Pope, Universal
Others	Press, journalism, curses

Example 3 As example, let us suppose that *John* posts 4 times about *corruption* (*c*), 3 times about *sports* (*s*), 2 times about *politics* (*p*) and 1 time about *international* (*i*) on time 3. The temporal preferences of John on 3 are: $\Gamma_3^{John} = \{c \succ_3^{John} s, s \succ_3^{John} p, p \succ_3^{John} i, i \succ_3^{John} security, i \succ_3^{John} education, i \succ_3^{John} entertainment, i \succ_3^{John} economy, i \succ_3^{John} religion, i \succ_3^{John} others\}$, besides those temporal preferences obtained from transitive closure of Γ_3^{John} omitted for better presentation.

Figure 4 illustrates samples of the evolving network. As we presented in Sect. 2 there are different strategies to extract user preferences from social networks. We chose the use of topic modeling in order to handle network content and then correlate the evolution patterns of these preferences with evolution patterns of centrality metrics. In Pereira et al. (2016b) we used a different technique to extract preferences mostly based on network topology (number of followers/followees). By considering topics, we improve the impact of our findings as extracting preferences from topics is based on network content.

5.2 This is my jam dataset

This Is My Jam (TIMJ) was an online social music network where users could share their favorite songs with their followers. Only one song could be shared at a time – the current *jam*, which lasted for up to one week in users' statuses. Furthermore, as a social network, users could like each other's jam. TIMJ dataset was released by Jansson et al. (2015). We built a temporal network based on users' likes, where nodes are Jam users and an edge (u_1, u_2, t) means that u_2 liked u_1 's jam posted at t . In this way, the directed edges represent the music influence flow. Jam network features are summarized in Table 1.

5.2.1 Extracting preferences

User preferences were extracted based on music genres. Originally, the TIMJ dataset does not contain jams genre annotations. Jansson et al. (2015) mapped the TIMJ dataset to




Fig. 4 Snapshots of samples of the evolving interaction network. Nodes are Twitter users. One tie from user u_1 to u_2 means that u_2 retweeted at t some text originally posted by u_1 . The colors represent topics that users are talking about at t . The samples were built by filtering nodes with degree between 50–22,000 and edges representing the 4 most popular topics. Each snapshot corresponds to 1 day time-interval. This figure highlights the *edges* evolving aspect. Nodes are not evolving for better visualization (Color figure online)

the Million Song Dataset (MSD) (Bertin-Mahieux et al. 2011) – a million popular collection of music tracks and their metadata. From these music tracks, we considered the ground truth CD2 from Schreiber (2015) to obtain song-level genre annotations. Only the songs present in the ground truth were taken into account in our analysis. As result, the final set of preference domain is composed by 15 elements: $A = \{rock, pop, country, electronic, reggae, rnb, metal, jazz, punk, folk, latin, world, rap, blues, newage\}$ and we got 528,787 jams annotated with the underlying genre $o \in A$.

The pairwise preferences for each user are extracted from the current jam genre. If user u posted a jam annotated with genre o at time t , then u clearly prefers o over the remaining genres in the domain at that moment. As in Twitter dataset, we considered the weight $w_t^u(o)$ based on the number of times the same genre appeared in u 's status during the time granularity being taken into account.

Example 4 As example, let us suppose that *Mary* posted 3 *rock* jams and 2 jams of *pop* on time $t = 15$. The temporal preferences of *Mary* at t are: $\Gamma_{15}^{Mary} = \{rock \succ_{15}^{Mary} pop, pop \succ_{15}^{Mary} country, pop \succ_{15}^{Mary} electronic, \dots, pop \succ_{15}^{Mary} newage\}$, besides those temporal preferences obtained from transitive closure of Γ_{15}^{Mary} .

5.3 Discussion

Though our analysis is limited to the Twitter news and social music domains due to the availability of public datasets, we expect our results to generalize to other items like movies, videos, books, vacation packages, shopping etc., which are fairly susceptible to social influence effects. In both domains, the user preferences were extracted based on the content being shared by the users whereas the temporal networks were built based on the interaction of the users with their friends. Moreover, our proposed method behavior will not be affected if users' preferences are estimated from completely independent external sources, as social networks invariably model users behaviors.

6 Experimental evaluation

The main goal of experiments is to investigate the correlation between preference changes δ (Def. 5) and node events ε (Def. 8) on Twitter and Jam temporal networks.⁵ All algorithms were implemented in Java language using Gephi API⁶ as foundation. All the experiments run over a server equipped with Intel(R) Xeon(R) CPU @ 2.40GHz on 140GB RAM, twenty cores and Linux Ubuntu operating system.

6.1 Experimental environment

Centrality Metrics We consider two centrality measures: *betweenness* and *closeness*. These measures have different meanings and our objective is to stress to what extent their evolution correlate with preference changes.

According to Zafarani et al. (2014), in closeness centrality, the intuition is that the more central nodes are, the more quickly they can reach other nodes. Formally, these nodes should have a small average shortest path length to other nodes. The smaller the average shortest path length, the higher the centrality for the node. The betweenness centrality characterizes how important nodes are in connecting other nodes. For a node v , compute the number of shortest paths between other nodes that pass through v .

Change-point scores and preference changes For node events detection we consider three different scores: the proposed approaches (1) average score Γ and (2) ranking score Λ , and (3) the baseline approach of Akoglu and Faloutsos (2010) which we call Z score. In this baseline approach, authors also propose to spot change-points on a time-varying graph from which many nodes deviate from their common behavior. It is the work more related to ours due to two aspects: (i) the change-point based approach and (ii) the temporal dynamics of the network. The idea is to characterize a node with several features so that it becomes a multi-dimensional point. Z score is computed in function of the dot-product between the current feature-vector \mathbf{v} and a typical feature-behavior \mathbf{r} , which is the average of past feature-vectors.

For preference change detection we implement our proposed approach δ described in Sect. 3.

Social network modeling We compare *static networks* with *temporal networks*. The difference is that in the temporal scenario we consider temporal paths (fastest paths, as discussed in Sect. 4) when computing centrality metrics, while in the static scenario we consider shortest paths. In temporal networks the temporal order is taken into account, while in static networks it is not. Note that despite static networks do not consider edges labeled with time instants, they are analyzed over time, considering also the sliding window. The difference between both approaches is essentially that inside the window being analyzed, time instants are considered (temporal) or not (static) when computing nodes centrality.

Datasets We vary the time granularity of the social temporal networks Twitter and Jam. In Jam network, time granularities are month, semester and year. In Twitter network, we consider day, week and month. Thus, in all we have six social networks related with news and music domains.

Window size $|W|$ The solutions we propose for the problem of preference change and node events detection are highly sensitive to the size of the observation window W . We vary the window size with values of 2, 4 and 7 time units. This size is related to the desired semantics we wish to analyze. If we are interested in tracking short-term events, then short sizes fit

⁵ Source codes available at <http://www.lsi.facom.ufu.br/~fabiola/temporal-networks>.

⁶ <https://gephi.org/toolkit/>.

Table 4 Experimental environment

Feature	Variation	Default
Dataset	Jam-month, jam-semester, jam-year Twitter-day, twitter-week, twitter-month	Jam-month Twitter-week
$ W $	2, 4, 7	2
θ (Γ and Π)	0.1, 0.2, 0.5	0.2
θ (Z)	0.01, 0.015, 0.04	0.015
Centrality metric	Betweenness, closeness	Closeness
Change-point score Θ	Ranking Λ , average Π , Z	Average Π
SN Modeling	Static, temporal	Temporal

better. For instance, preferences over the domains of news or restaurants have a high rate of change. On the other hand, long sizes are more appropriate when the events are not frequent, for example preferences about musics and movies. Twitter-month does not vary for values 4 and 7 because it does not contain more than 3 months. The same occur with jam-year because it is limited to 4 time steps (4 years).

Threshold θ Adjusts the intensity of node events we are looking for, varying from smooth to drastic events. In our experiments we explore how this intensity impacts on correlations with preference changes. From some observations in our data, we detected that Z score has lower levels of θ in comparison to the other scores. Thus, we consider different ranges according to scores. To setup Z values, we varied from 0.01 to 0.05 in a 0.001 granularity in order to observe the amount of detected events for the default features described above. After, we chose the following values to conduct the remainder of the experiments based on diversity: 0.01, 0.015 and 0.04. For ranking and average, the procedure was the same, varying from 0.1 to 0.5, and the final values are: 0.1, 0.2 and 0.5.

Table 4 summarizes values considered in our experiments.

6.2 Performance evaluation

The results in Fig. 5 correspond to runtimes of the Algorithm 1 for all datasets and different window sizes $|W|$. According to Algorithm 1 complexity analysis, detecting preference changes costs $O(|T||P|)$, which is related to the number of temporal preference relations P and the time interval T being analyzed—the longer T , the more costly the algorithm will be. In Fig. 5 we refer to the runtime accumulated for all users in the datasets. Twitter contains more users than Jam. Twitter-day network contains the largest interval $T = 94$ and jam-year has a low number of users as well as a short time interval $T = 4$. The window size is related to P . As $|W|$ increases, more temporal preference relations can be extracted, impacting on the runtime.

Remarking on Algorithm 2, the runtimes to detect node events are depicted in Fig. 6. For the sake of simplicity, we do not present ranking score runtime information. Ranking and average scores have the same computational complexity. *NodeEventDetection* performance is directly related with network size, which means that the more nodes and edges in a network, more paths between nodes can be detected. In all scenarios, temporal networks are more time-consuming than static networks counterpart. In fact, when considering temporal




Fig. 5 Performance evaluation of the algorithm *PrefChangeDetection*. Runtimes refer to the time elapsed to process all users of the corresponding dataset (Color figure online)

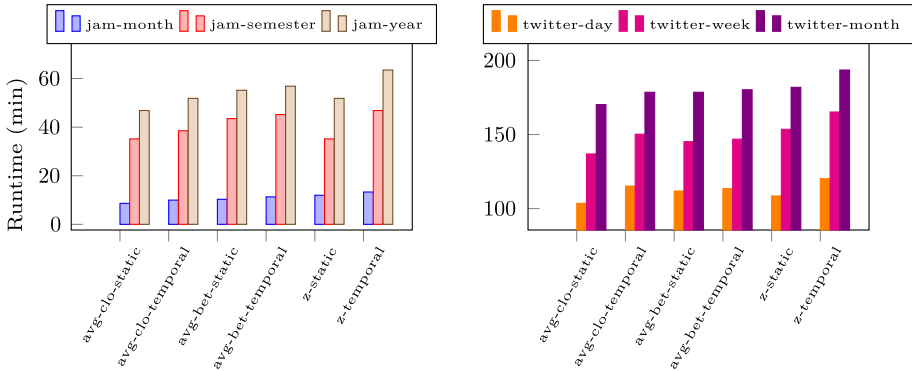


Fig. 6 Performance evaluation of the algorithm *NodeEventDetection*. Runtimes refer to the time elapsed to process all users of the corresponding dataset

order, there are more paths than when time is not taken into account. Comparing centrality runtime behavior, we conclude that computing closeness centrality is faster than computing betweenness centrality (Brandes 2001). This difference also impacts on the high runtime elapsed by Z score, which covers both centralities.

6.3 Analyzing network and preference evolution

Taking into account the set of parameters and possible scenarios to stress, we first perform observations taken from both specific nodes/users and the whole network evolving behavior. In the following, we detail important evidences extracted from these observations.

6.3.1 The evolving networks

In the first analysis we compare, quantitatively, all change-points scores averaged over all users for each time step, also varying centrality metrics. Default setup was considered for the remaining features. The results are presented in Figs. 7 and 8, for Twitter and Jam networks, respectively. This experiment reflects networks’ global behavior. The most important observation is that values for *ranking* and *average* are high in contrast to Z indicating that we should consider different values for θ when detecting events, otherwise *ranking* and *average* scores will detect much more node events than Z, not reflecting the reality. This behavior can be explained by the fact that Z score is more complex and consider a set of centrality measures (closeness and betweenness in this case) to describe a node while *ranking* and *average*

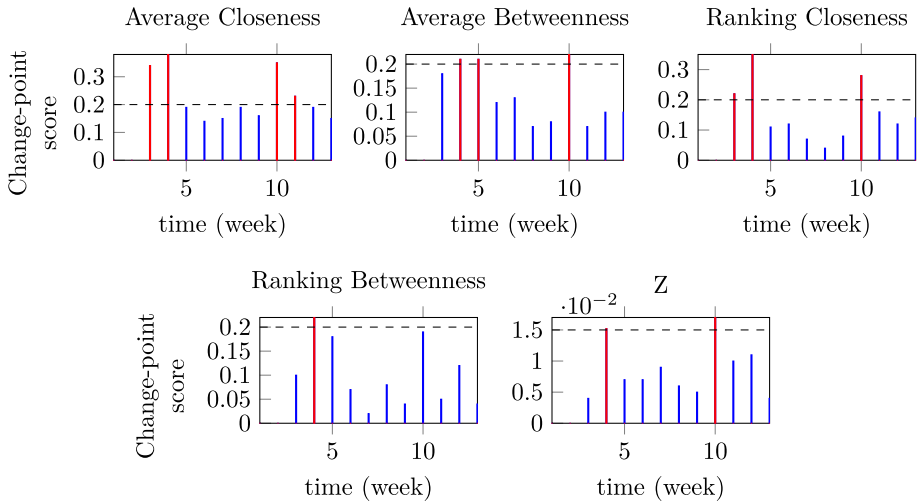


Fig. 7 Change-point scores averaged over all users in Twitter-week network for $|W| = 2$ and temporal modeling. Red bars are the change-points detected for $\theta = 0.2$ (average and ranking) and $\theta = 0.015$ (Z) (Color figure online)

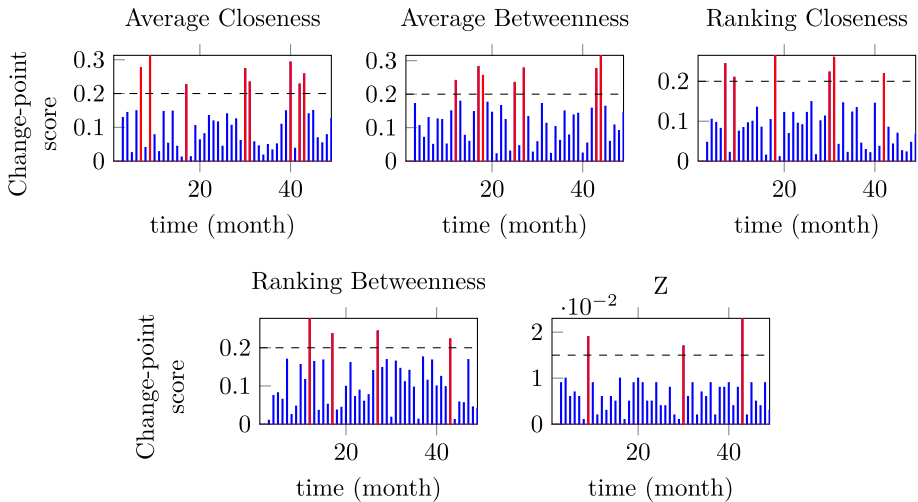


Fig. 8 Change-point scores averaged over all users for $|W| = 2$ and temporal modeling the Jam-month network. Red bars are the change-points detected for $\theta = 0.2$ (average and ranking) and $\theta = 0.015$ (Z) (Color figure online)

are computed with respect to only one centrality measure. Concerning centrality metrics, we vary *average* and *ranking* for closeness and betweenness. Quantitatively, change-point scores values remain in the same range independent to the centrality metric. The number of events detected is different for each centrality metric which is expected, as they have different meanings and thus vary according to different changes in the structure of the network.

Qualitatively speaking, the change-points detected occurred on similar moments for *average*, *ranking* and *Z* in both datasets. These observations give us confidence in terms of the time instants the events occurred, independent to the centrality metric and the change-point

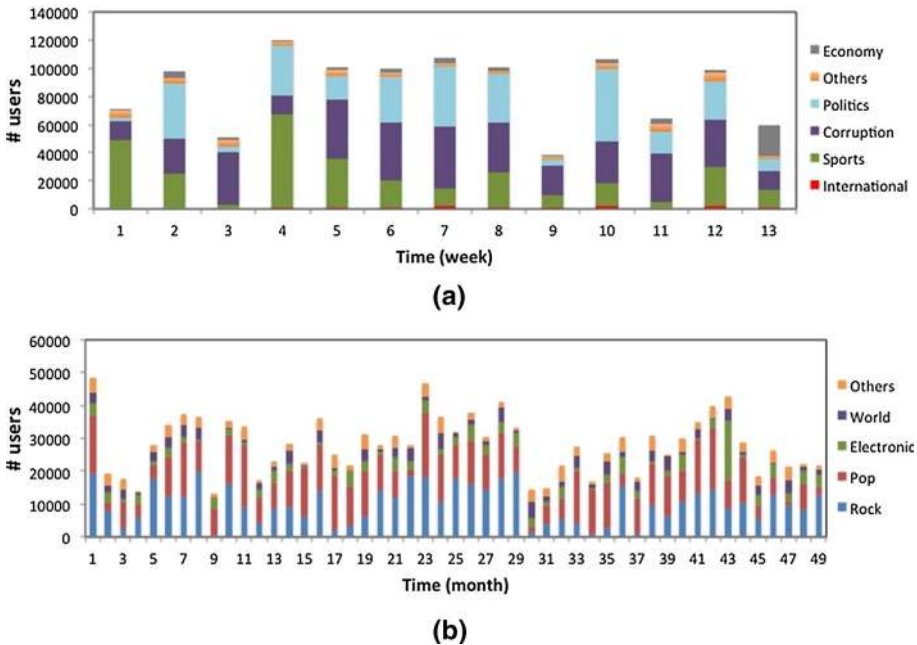


Fig. 9 Most preferred topics over time considering all users in the **a** twitter-week and **b** jam-month networks (Color figure online)

score strategy. It is an open question to define which change-point score fits better in a given scenario. The difficulty is related to the lack of a ground truth when analyzing social media data as users are scattered all across the globe (Zafarani and Liu 2015). We glimpse that the variety of scenarios we propose for detecting node events can be further stressed and used to define evaluation metrics. Remark that here we just perform an analytical comparison among the events as our focus is on correlating the detected change-points with preference changes, not on defining the highest accuracy for the node event detection task.

6.3.2 Preference dynamics

We analyze how preferences evolve in both networks considering a global perspective. Results are depicted in Fig. 9. In twitter-week network the topics *sports*, *corruption* and *politics* are the most preferred during the whole period. Comparing weeks 3 and 4, the number of users preferring *sports* over the others have increased. The same behavior can be observed for weeks 9 and 10 regarding *politics*, and 12 and 13 for *economy*. These change points occur around the same time instants detected on previous experiments, specially considering average closeness setup (Fig. 7).

Jam-month network users mostly prefer *rock* and *pop*. A pattern deviation can be observed on months 9, 30, 33, 34 and 43 when users mostly prefer genres different from *rock* and *pop*. Again, we can establish a comparison between these time steps and those detected on Fig. 8.

In order to illustrate a local perspective of preference change process, in Fig. 10 we show a given user u 's better-than graphs (BTGs) in two different moments of twitter-day network (u id = 58488491). On Aug 21st u preferences were *corruption* and *politics* over *sports* and then *sports* over the remaining topics. On Aug 22nd, new preferences $sports >_{Aug22}^u politics$ and $sports >_{Aug22}^u corruption$ appeared, causing a preference change event. After the revision,



Fig. 10 Better-than graphs representing u_3 preferences on days Aug 21 and Aug 22. Aug 21 was the end date of Olympic Games in Rio de Janeiro

the resulting acyclic BTG represents u preferences on Aug 22nd. Considering that Aug 21st was the end date of Olympic games in Rio de Janeiro, probably u had been influenced by this trending topic on the network.

6.3.3 User preferences and network evolution

In the last analysis we observe the relationship between change behaviors considering all nodes/users. Figure 11 depicts comparisons among all scoring strategies in relation to the percentage of nodes that change their behavior from twitter-week and jam-semester temporal networks. The first observation is that for all scores the percentages maintain a pattern with low deviation. This indicates coherency on scoring strategies. We can also observe that Z score detected fewer changes than *average* and *ranking*. Moreover, betweenness centrality detected a higher number of changes than closeness. From these observations, we were able to ascertain high levels of confidence concerning change-point scores and centrality metrics.

From the preference evolution viewpoint, the percentage of users that change their preferences is very similar to the percentage of nodes change-points previously discussed. On average, 36% and 27% of users changed their preferences on a weekly and semiannually basis, respectively.

6.4 Relating preference changes and node events

There are many directions to explore from the evidences presented in the previous section: (i) to what extent evolving networks are related with user preference dynamics (UPD)? (ii) Which centrality metrics should be used in order to analyze UPD? (iii) Which change-point scores should be considered and (iv) what is the best network modeling for analyzing UPD: static or temporal? To address these points we formulated the following research questions: *Q1: Is there a relationship between user preference changes and centrality-based node events in evolving social networks?*

In fact, in Twitter domain, preference changes and the network structure are both based on retweets. However, a preference change is based not only on quantitative retweets, but also on retweets text which define the preference domain of 10 topics. As a counterexample, suppose that a user u at time t_1 retweeted 3 times on the topic a and 1 time on the topic b . Then, at time t_2 , u retweeted 1 time on the topic a . Considering our retweet-defined network, the number of u 's incoming edges at t_1 is 3 while at t_2 is 1. This could imply on u 's centrality (betweenness or closeness) change. However, there is not a u 's preference change. In Jam

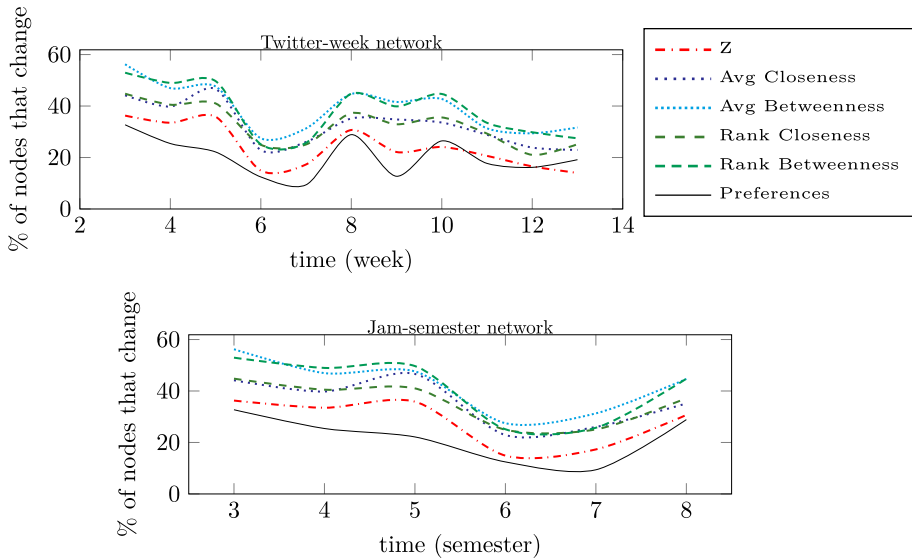


Fig. 11 Percentage of nodes/users that change centrality/preference at each time step in (top) Twitter-week and (bottom) Jam-semester temporal networks (Color figure online)

domain, preference changes and the network structure are not built on the basis of the same actions. Preferences are based on what users listen. The network is defined over what users explicitly like from other users. Under this perspective, a correlation is not straight and we investigate it in this research question.

We use Pearson Correlation Coefficient (PCC) to evaluate if there is a linear correlation between δ and ε and the strength of this correlation. For each user u of our observation period, we compute $PCC(\delta^u, \varepsilon^u)$ considering a population of the whole observation period (94 twitter-day, 13 twitter-week, 3 twitter-month and 49 jam-month, 8 jam-semester, 4 jam-year). Then we averaged these correlation values $PCC_{avg}(\delta, \varepsilon)$ over all users.

We explore several scenarios for each of the six social networks – *twitter-day*, *twitter-week*, *twitter-month*, *jam-month*, *jam-semester*, *jam-year*, in order to stress the time granularity effect. We also vary the parameter θ according to respective scores range (see Table 4). This parameter indicates that the closer to 1 more significant are the centrality changes that are being considered. Then, we vary the window size $|W|$ to explore long-term and short-term impact of the events on the correlation strength of variables. When considering smooth variations (low θ) more events were detected.

Each scenario compares PCC values in relation to betweenness and closeness centralities for average and ranking scores, and in relation to Z score with betweenness and closeness being used to describe a node. Figures 12 and 13 illustrate our results highlighting the comparison between static and temporal networks correlation strengths.

In all scenarios δ and ε associate significantly (as compared to the corresponding critical values – in all scenarios critical values are lower than 0.1). Our null hypothesis H_0 is that there is no linear correlation between δ and ε , i.e. $PCC = 0$. Two random variables (with no correlation) would have a 90% probability of p-value greater than a critical value. We observe a strong correlation between change events in user preferences and in centrality metrics in most scenarios.

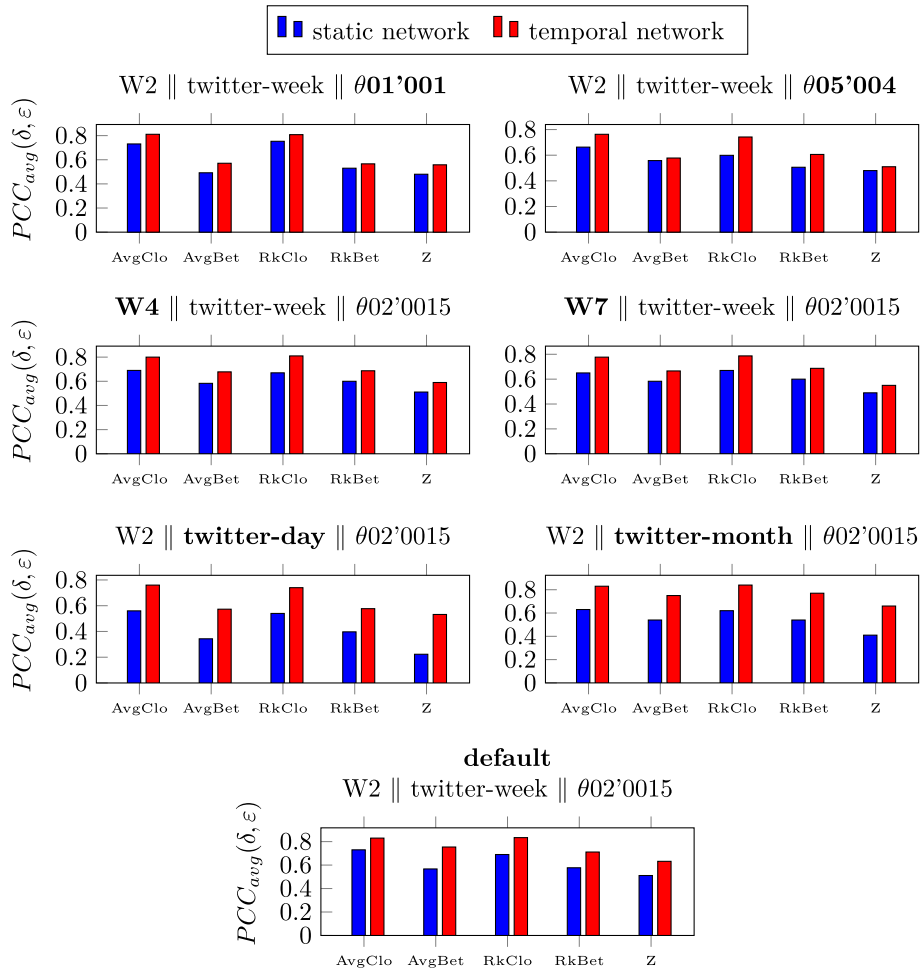


Fig. 12 PCC between preference changes and centrality-based node events for Twitter dataset (Color figure online)

PCC values for Jam networks are higher than Twitter networks comparing similar scenarios. This can be explained by the preference extraction strategies and inherent noise. In fact, the Jam preference semantic based on music genres is more accurate than the topic modeling strategy used for preference extraction from Twitter. Moreover, besides users mostly retweet their preferences, they can retweet due to other reasons (Metaxas et al. 2015), while in general users listen what they prefer (Moore et al. 2013).

Q2: Are temporal networks more suitable than static networks for analyzing user preference dynamics?

Across all scenarios modeling our network with temporal information made difference. The more time instants, the greater the difference of PCC values in temporal networks against static networks. For instance, in Twitter default scenario the higher PCC in the static network is 0.71 while the same value corresponds to the lower PCC in the temporal network. Thus, temporal networks are statistically more suitable than static networks for analyzing UPD.

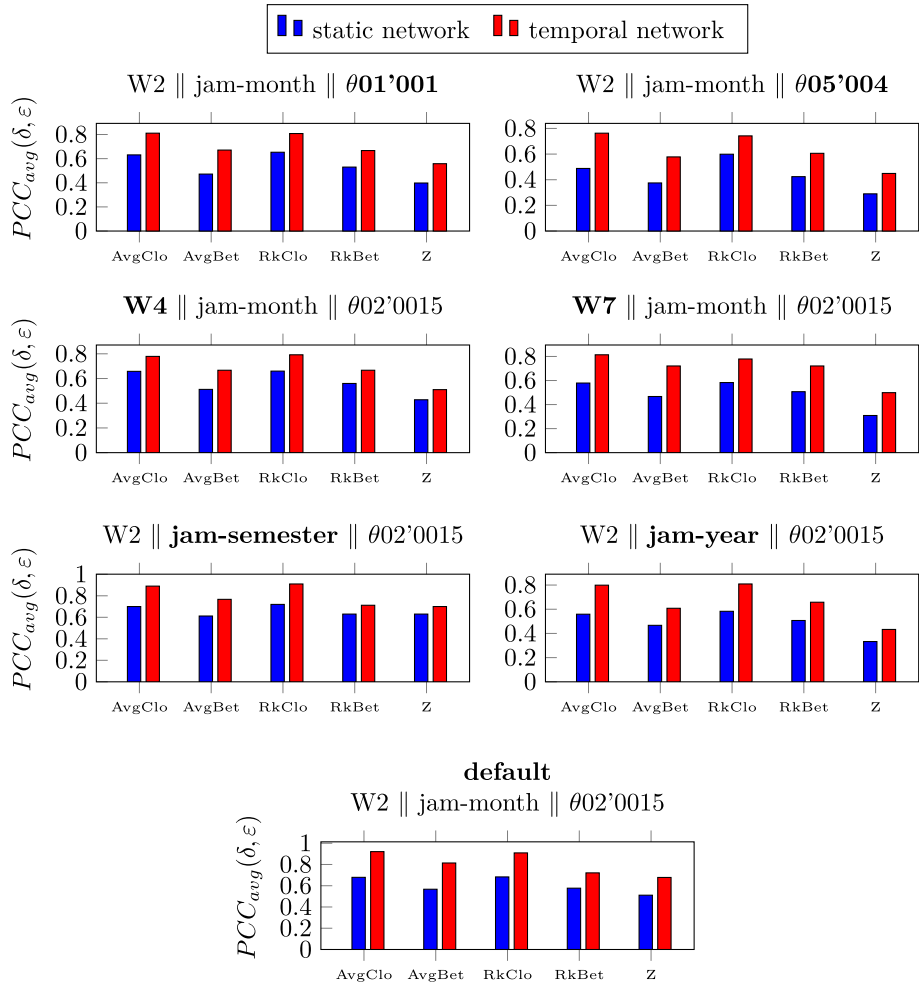


Fig. 13 PCC between preference changes and centrality-based node events for This Is My Jam dataset (Color figure online)

The results obtained so far can be explained by the phenomena of information propagation and inherent consequences of homophily and influence. The main difference between temporal and static networks is that temporal networks take into account the contact sequence between nodes to compute paths (Pereira et al. 2016a) and this has an impact on different centrality measures. The related work Guille and Hacid (2012) discusses about relation among preferences and information propagation in social networks. The aspects described in our motivating example (Sect. 1) could illustrate that preferences are directed by information flow on the social network. Finally, temporal networks represent information flow more realistically.

Q3: Considering closeness and betweenness, what change-point score and respective centrality metrics should be used when analyzing UPD?

If we analyze correlation values comparing change-point scores, we find that *average* and *ranking* are more correlated than *Z*. However, there is no consensus in relation to the best change-point score. For instance, considering Jam default settings *average* has stronger

correlations than *ranking*, but in Twitter default the behavior is the opposite. Despite *Z* score is induced by structural measures as in *average* and *ranking*, the combination of betweenness and closeness measures to describe a node did not result in stronger correlations than considering them separately. The centralities are conceptually different and not necessarily when one is highly correlated the other will be, decreasing *Z* score performance.

Now observing centrality metrics we conclude that closeness is more suitable when correlating UPD and node events, considering average and ranking. The closeness centrality measures the inverse total distance to all other nodes and is high for nodes that are close to all others. Similarly, for temporal networks, the idea is to measure how quickly a node may on average reach other nodes. In this work, we hypothesize that user preference dynamics are related to her social network evolution, given the aspect of social influence and, consequently, network structural changes. Thus, we conclude that this relationship of a node quickly reaching others is the most important aspect we should consider when addressing the UPD problem.

7 Conclusion and future work

We have investigated the interplay between temporal dynamics of user preferences and her social network evolving over time. The first step was to define what are user preference dynamics (UPD). We have proposed a temporal preference model able to describe user preferences over time through user profiles. Moreover, we have defined a strategy based on inconsistency to detect changes on temporal preferences as time flies. As a solution for analyzing preference dynamics, we considered temporal social networks, i.e., social networks where the order that interactions occur are taken into account when computing network structural metrics. We explored the idea of centrality-based node event detection in order to identify significant changes of a node on the network. Finally, we joined our proposals and performed an experimental evaluation focused on the main goal: the interplay between user preferences and social networks over time. We have discovered that there is a strong correlation between preference change events and centrality-based node events, specially when considering temporal networks (temporal node centrality). Moreover, we have concluded that closeness centrality is more suitable when correlating UPD and node events than betweenness. By correlating changes on preferences and changes on node centralities we move towards understanding how *content* and *topology* evolve on a social network.

Many lines of research remain open for future works. A limitation of our work is the lack of a ground truth to determine (i) whether detected node events indeed are events for a given user and (ii) whether extracted preferences really reflect users' tastes. In fact, evaluation without a ground truth in social media research is a pressing need (Zafarani and Liu 2015). Another direction is to investigate research questions that arise when analyzing UPD on social networks. For instance, is there a causality relation? Who are the most influential users when a preference change is detected? Finally, we glimpse the need of online and incremental algorithms for streaming graphs (Kas et al. 2013; Aggarwal and Subbian 2014). The algorithms we designed in this paper do not support online stream processing of temporal networks.

Acknowledgements This work was supported by the research project "TEC4Growth - Pervasive Intelligence, Enhancers and Proofs of Concept with Industrial Impact/NORTE-01-0145-FEDER-000020", financed by the North Portugal Regional Operational Programme (NORTE 2020). This work was also supported by the Brazilian Research Agencies CAPES and CNPq. GMBO is also grateful to Fapemig support.

References

- Abbasi, M. A., Tang, J., & Liu, H. (2014). Scalable learning of users' preferences using networked data. In *Proceedings of the 25th ACM conference on hypertext and social media* (pp. 4–12). New York, NY, USA: ACM. HT '14.
- Agarwal, D., & Chen, B. C. (2010). flda: Matrix factorization through latent dirichlet allocation. In *Proceedings of the third ACM international conference on web search and data mining* (pp. 91–100). New York, NY, USA: ACM. WSDM '10.
- Aggarwal, C., & Subbian, K. (2014). Evolutionary network analysis: A survey. *ACM Computing Surveys*, 47(1), 10–36.
- Aggarwal, C. C., & Subbian, K. (2012). Event detection in social streams. In *12th SIAM international conference on data mining* (pp. 624–635). USA.
- Akoglu, L., & Faloutsos, C. (2010). Event detection in time series of mobile communication graphs. In *Proceedings of 27th army science conference*, no. 3 in 18.
- Althoff, T., Jindal, P., & Leskovec, J. (2017). Online actions with offline impact: How online social networks influence online and offline user behavior. In *Proceedings of the tenth ACM international conference on web search and data mining* (pp. 537–546). New York, NY, USA: ACM. WSDM '17.
- Arias, M., Arratia, A., & Xuriguera, R. (2014). Forecasting with twitter data. *ACM Transactions on Intelligent Systems and Technology*, 5(1), 8:1–8:24.
- Bertin-Mahieux, T., Ellis, D. P., Whitman, B., & Lamere, P. (2011). The million song dataset. In: *Proceedings of the 12th international conference on music information retrieval (ISMIR)*.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., & Hwang, D. U. (2006). Complex networks: Structure and dynamics. *Physics Reports*, 424(4–5), 175–308.
- Brandes, U. (2001). A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25, 163–177.
- Cadilhac, A., Asher, N., Lascarides, A., & Benamara, F. (2015). Preference change. *Journal of Logic, Language and Information*, 24(3), 267–288.
- Christidis, K., Apostolou, D., & Mentzas, G. (2010). Exploring customer preferences with probabilistic topics models. In *Preference learning workshop, ECML/PKDD*.
- Cordeiro, M., & Gama, J. (2016). *Online social networks event detection: A survey* (pp. 1–41). Cham: Springer International Publishing.
- Cordeiro, M., Sarmiento, R. P., & Gama, J. (2016). Dynamic community detection in evolving networks using locality modularity optimization. *Social Network Analysis Mining*, 6, 15. <https://doi.org/10.1007/s13278-016-0325-1>.
- de Amo, S., Diallo, M. S., Diop, C. T., Giacometti, A., Li, D., & Soulet, A. (2015). Contextual preference mining for user profile construction. *Information Systems*, 49, 182–199.
- Eberle, W., & Holder, L. (2016). Identifying anomalies in graph streams using change detection. In *KDD workshop on mining and learning in graphs (MLG)*.
- Guille, A., & Hacid, H. (2012). A predictive model for the temporal dynamics of information diffusion in online social networks. In *Proceedings of the 21st international conference on world wide web* (pp. 1145–1152). New York, NY, USA: ACM. WWW '12 Companion.
- Guille, A., Hacid, H., Favre, C., & Zighed, D. A. (2013). Information diffusion in online social networks: A survey. *SIGMOD Record*, 42(2), 17–28.
- Hansson, S. O. (1995). Changes in preference. *Theory and Decision*, 38(1), 1–28.
- Holme, P. (2014). Analyzing temporal networks in social media. *Proceedings of the IEEE*, 102(12), 1922–1933.
- Holme, P., & Saramaki, J. (2012). Temporal networks. *Physics Reports*, 519(3), 97–125.
- Ide, T., & Kashima, H. (2004). Eigenspace-based anomaly detection in computer systems. In *Proceedings of the tenth ACM SIGKDD international conference on knowledge discovery and data mining* (pp 440–449) KDD '04.
- Imran, M., Chawla, S., & Castillo, C. (2016). A robust framework for classifying evolving document streams in an expert-machine-crowd setting. In *Proceedings of the 18th international conference on data mining (ICDM)*.
- Jansson, A., Raffel, C., & Weyde, T. (2015). This is my jam—data dump. *16th International Society for Music Information Retrieval Conference Late Breaking and Demo Papers*.
- Kapoor, K. (2014). Models of dynamic user preferences and their applications to recommendation and retention. Ph.D. thesis, University of Minnesota.

- Kapoor, K., Srivastava, N., Srivastava, J., & Schrater, P. (2013). Measuring spontaneous devaluations in user preferences. In *Proceedings of the 19th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1061–1069). KDD '13.
- Kas, M., Wachs, M., Carley, K. M., & Carley, L. R. (2013). Incremental algorithm for updating betweenness centrality in dynamically growing networks. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining* (pp. 33–40). New York, NY, USA: ACM. ASONAM '13.
- Liu, F. (2011). *Reasoning about preference dynamics* (1st ed., Vol. 354). Netherlands: Springer.
- Liu, X. (2015). Modeling users' dynamic preference for personalized recommendation. In *Proceedings of the 24th international joint conference on artificial intelligence (IJCAI'15)* (pp 1785–1791).
- Lou, J. K., Wang, F. M., Tsai, C. H., Hung, S. C., Kung, P. H., & Lin, S. D. (2013). Modeling the diffusion of preferences on social networks. In *Proceedings of the 2013 SIAM international conference on data mining* (pp. 605–613).
- Macropol, K., Bogdanov, P., Singh, A.K., Petzold, L., & Yan, X. (2013). I act, therefore i judge: Network sentiment dynamics based on user activity change. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining* (pp. 396–402). ASONAM '13.
- Metaxas, P., Mustafaraj, E., Wong, K., Zeng, L., O'Keefe, M., & Finn, S. (2015). What do retweets indicate? results from user survey and meta-review of research. *Ninth international AAAI conference on web and social media (ICWSM)* (pp. 658–661).
- Moore, J., Chen, S., Turnbull, D., & Joachims, T. (2013). Taste over time: The temporal dynamics of user preferences. In *Proceedings of the 14th international society for music information retrieval conference*.
- Nicosia, V., Tang, J., Mascolo, C., Musolesi, M., Russo, G., & Latora, V. (2013). *Temporal Networks*. Springer Berlin Heidelberg, Berlin, Heidelberg, chap Graph Metrics for Temporal Networks, pp. 15–40.
- Oliveira, M., Guerreiro, A., & Gama, J. (2014). Dynamic communities in evolving customer networks: An analysis using landmark and sliding windows. *Social Network Analysis and Mining*, 4(1), 208.
- Pereira, F. S. F. (2015). Mining comparative sentences from social media text. In *Workshop on interactions between data mining and natural language processing (DMNLP) co-located with European conference on machine learning and principles and practice of knowledge discovery in databases (ECML/PKDD)* (pp. 41–48).
- Pereira, F. S. F., & de Amo, S. (2015). Mineracao de preferencias do usuario em textos de redes sociais usando sentencas comparativas. In *Symposium on knowledge discovery, mining and learning (KDMiLe)* (pp. 94–97).
- Pereira, F. S. F., Amo, S., & Gama, J. (2016a). Evolving centralities in temporal graphs: a twitter network analysis. In *17th IEEE international conference on mobile data management (MDM), 2016*.
- Pereira, F.S.F., de Amo, S., & Gama, J. (2016b). On using temporal networks to analyze user preferences dynamics. In *Discovery science: 19th international conference, DS 2016, Bari, Italy, 2016*.
- Rafailidis, D., & Nanopoulos, A. (2014). Modeling the dynamics of user preferences in coupled tensor factorization. In *Proceedings of the 8th ACM conference on recommender systems* (pp. 321–324). ACM.
- Ranshous, S., Shen, S., Koutra, D., Harenberg, S., Faloutsos, C., & Samatova, N. F. (2015). Anomaly detection in dynamic networks: A survey. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(3), 223–247.
- Rossetti, G., Guidotti, R., Miliou, I., Pedreschi, D., & Giannotti, F. (2016). A supervised approach for intra-/inter-community interaction prediction in dynamic social networks. *Social Netw Analys Mining*, 6(1), 86.
- Schlitter, N., & Falkowski, T. (2009). Mining the dynamics of music preferences from a social networking site. In *International conference on advances in social network analysis and mining, 2009. ASONAM '09*. (pp 243–248).
- Schreiber, H. (2015). Improving genre annotations for the million song dataset. In *Proceedings of the 16th international society for music information retrieval conference, ISMIR* (pp. 241–247).
- Sun, J., & Tang, J. (2011). A survey of models and algorithms for social influence analysis. In C. C. Aggarwal (Ed.), *Social Network Data Analytics* (pp. 177–214). US: Springer.
- Sun, Y., Li, H., Councill, I. G., Lee, W. C., & Giles, C. L. (2008). Measuring user preference changes in digital libraries. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management* (pp 1497–1498). CIKM '08.
- Tan, C., Tang, J., Sun, J., Lin, Q., & Wang, F. (2010). Social action tracking via noise tolerant time-varying factor graphs. In *Proceedings of the 16th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1049–1058). KDD '10.
- Thimm, M. (2013). Dynamic preference aggregation under preference changes. In *Proceedings of the fourth workshop on dynamics of knowledge and belief (DKB'13)*.

- Wallach, H.M., Mimno, D., & McCallum, A. (2009). Rethinking lda: Why priors matter. In *Proceedings of the 22nd international conference on neural information processing systems* (pp. 1973–1981). NIPS'09.
- Wei, W., & Carley, K. M. (2015). Measuring temporal patterns in dynamic social networks. *ACM Transactions on Knowledge Discovery from Data*, 10(1), 9:1–9:27.
- Wilson, N. (2004). Extending cp-nets with stronger conditional preference statements. In *Proceedings of the 19th National Conference on Artificial Intelligence* (pp. 735–741). AAAI'04.
- Wu, H., Cheng, J., Huang, S., Ke, Y., Lu, Y., & Xu, Y. (2014). Path problems in temporal graphs. *Proceedings of the VLDB Endowment*, 7(9), 721–732.
- Wu, L., Ge, Y., Liu, Q., Chen, E., Long, B., & Huang, Z. (2016). Modeling users' preferences and social links in social networking services: A joint-evolving perspective. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 279–286).
- Xiang, L., Yuan, Q., Zhao, S., Chen, L., Zhang, X., Yang, Q., & Sun, J. (2010). Temporal recommendation on graphs via long- and short-term preference fusion. In *Proceedings of the 16th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 723–732). New York, NY, USA:ACM. KDD '10.
- Zafarani, R., & Liu, H. (2015). Evaluation without ground truth in social media research. *Communications of the ACM*, 58(6), 54–60.
- Zafarani, R., Abbasi, M. A., & Liu, H. (2014). *Social media mining: An introduction*. New York, NY, USA: Cambridge University Press.
- Zhang, J., Wang, C., & Wang, J. (2014a). Learning temporal dynamics of behavior propagation in social networks. In *Proceedings of the twenty-eighth AAAI conference on artificial intelligence* (pp. 229–236). AAAI'14.
- Zhang, J., Wang, C., Wang, J., & Yu, J. X. (2014b). Inferring continuous dynamic social influence and personal preference for temporal behavior prediction. *Proceedings of the VLDB Endowment*, 8(3), 269–280.
- Zhang, Y., Zhou, J., & Cheng, J. (2011). Preference-based top-k influential nodes mining in social networks. In *2011 IEEE 10th international conference on trust, security and privacy in computing and communications* (pp. 1512–1518).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.