

Graph-based Rumour Detection for Social Media

Nguyen Thanh Tam

Ecole Polytechnique Federale de Lausanne
Lausanne, Switzerland

ABSTRACT

Social platforms became a major source of rumours. While rumours can have severe real-world implications, their detection is notoriously hard: Content on social platforms is short and lacks semantics; it spreads quickly through a dynamically evolving network; and without considering the context of content, it may be impossible to arrive at a truthful interpretation. Traditional approaches to rumour detection, however, exploit solely a single content modality, e.g., social media posts, which limits their detection accuracy. In this paper, we cope with the aforementioned challenges by means of a multi-modal approach to rumour detection that identifies anomalies in both, the entities (e.g., users, posts, and hashtags) of a social platform and their relations. Based on local anomalies, we show how to detect rumours at the network level, following a graph-based scan approach.

1 INTRODUCTION

Social platforms became widely popular as a means for users to share content and interact with other people. Due to their distributed and decentralised nature, content on social platforms is propagated without any type of moderation and may thus contain incorrect information. Wide and rapid propagation of such incorrect information quickly leads to *rumours* that may have a profound real-world impact. For instance, in April 2013, there was rumour about two explosions in the White House, injuring also Barrack Obama [59]. The rumour was fuelled by content posted using a hacked Twitter account associated with a major new agency. The resulting panic had major economic consequences, such as a \$136.5 billion loss at the stock market. This incident highlights the need for early and accurate *rumour detection*, in particular on social platforms.

It is notoriously hard to detect rumours [47]. Posts on social platforms are short and lack semantics. For instance, tweets have a limited number of characters, and comprise slang and spelling mistakes. Hence, traditional techniques to assess the credibility of (long, well-written) documents are of limited use for social platforms. Also, user interactions at unprecedented scale lead to rumours spreading quickly. Earliness of rumour detection is as important as detection accuracy. Moreover, social platforms are dynamic. Content is posted continuously, so that rumour detection cannot exhaustively collect data before giving results, but needs to work with streaming data. Finally, posts on social platforms are contextual. A post in isolation may not provide sufficient information for rumour detection. Instead, modalities such as user backgrounds, hashtags, cross-references, and user interactions must be considered to improve detection accuracy.

Several debunking services such as *snoopes.com* have been established to expose rumours and misinformation. They harness collaborative user efforts to identify potential rumours, which are then verified by experts. Due to such manual processing, the number of potential rumours that can be assessed is limited and significant time is needed for verification, which motivated work on automated rumour detection. Given the short length of posts on social platforms, rumour detection is often approached by grouping posts that relate to a single event [27]. This does not work in an online setting, though, since the posts related to an event are not available a priori.

Traditional rumour detection techniques tend to rely solely on the textual information of posts, potentially combined with features on post authors and their relations. However, focusing on one or two modalities of posts on social platforms is insufficient. For instance, users posting rumour-related content are often ignored by other users, which is not directly visible in features that capture solely the characteristics of a single user. In another example, posts circulating among a group of users that believe in conspiracy theories are likely to refer to rumours. Without information from outside the group, it is impossible to know whether these posts are related to a rumour.

Against this background, we argue for a novel approach to rumour detection that identifies anomalies on social platforms by comparing data *between peers* and *with the past*. Such anomalies can be observed for different modalities (e.g., users, tweets) and at varying levels of granularity. For example, a sudden increase or decrease in the number of followers of a user may be related to the user spreading rumours. Also, within a group of users, the credibility of one user being significantly lower than their peers may stem from the propagation of rumours. Moreover, relations between entities (e.g., users, posts, hashtags, links) may hint at anomalies, e.g., differences in time and location mentioned in a tweet and in a linked article.

In this paper, we present models and methods to realise the idea of detecting rumours based on anomalies. To this end, we follow a data management approach: We ground rumour detection in algorithms that work on a generic graph representation of social data, thereby achieving a solution that is applicable for any type of social platform. We first show how to identify anomalies locally, by assessing entities and relations of a social platform in comparison to their peers and to their past. Yet, acknowledging the inherent randomness of social platforms, anomalies are then viewed at a broader scale. To conclude on the spread of rumours, which is deemed more important than their classification [47], we incorporate the vicinity of local anomalies.

Our contributions and the structure of the paper (following a discussion of some background in §2) are summarised as follows:

- *Social Platform Model and Rumour Detection* (§3). Based on a model for social platforms, we develop a general process to detect rumours based on local and global anomalies.

- *Local Anomaly Detection (§4)*. We propose a non-parametric method for anomaly detection at the level of individual entities, based on differences between (i) current and past observations related to an entity, and (ii) the entity and its peers.
- *Global Anomaly Detection (§5)*. We lift anomaly detection to groups of entities, taking into account relations between them.

An evaluation of our approach is presented in §6. We review related work in §7 and conclude in §8.

2 BACKGROUND

Anomalies in social media. Abnormal propagation of information on social platforms can be classified as different types of anomalies, including hypes, fake news, satire news, disinformation, misinformation, and rumours [61]. For hypes, information is propagated in cascades that accidentally ‘blow-up’ on social platforms, e.g., related to popular events. Rumours, in turn, originate from the fact that people tend to exaggerate what they dislike [4]. Their veracity needs to be assessed, which is commonly done by assigning a trust score to entities, such as users and posts [1].

Here, we focus on detecting rumours. While hypes and rumours share some characteristics, they differ in how information is propagated. In hypes, information is spread randomly and chaotically. As revealed in a recent survey [47], however, rumours are propagated in a channelled manner, spreading ‘farther, faster, and deeper’ through interactions of actual users rather than bot accounts.

Type of anomalies differ in their sets of indicative signals. For example, detection of hypes (e.g., breaking news) focuses on peak volume of social posts and sharing activities [35, 36]. Spam detection of online reviews, in turn, uses user signals, such as average rating, number of reviews, and selectivity [51]. Our approach for rumour detection looks at inconsistency signals, exemplified below.

Twitter as an example. While we use Twitter as an example of a social platform throughout the paper, our model is applicable to other social platforms [39], as it is based on a universal graph representation (§3), generic statistical measures to compute anomalies (§4), and a graph-based anomaly detection algorithm (§5).

Consider a snapshot of Twitter social graph, as shown in Fig. 1. It includes users, tweets, hashtags, and linked articles. Each entity has different features, e.g., a user has a registration date and a number of followers. Entities are connected by relations. For instance, the relation between a tweet and an article indicates that the content of the tweet contains a link to that article. Moreover, each relation has an attribute value, e.g., the tweet-article relation has an attribute that indicates the difference between the publication dates of the tweet and the article, respectively.

Rumours are often manifested in anomalies related to entities and their relations. In Fig. 1, one may observe that the highlighted user has a registration date that is significantly newer than those of related users. At the same time, the number of followers is very high, compared to the historical record of the user. Other entities in this example are also suspicious, due to anomalies. For the highlighted tweet, the number of retweets is suddenly higher than in the past, as is the number of mentions for the highlighted linked article.

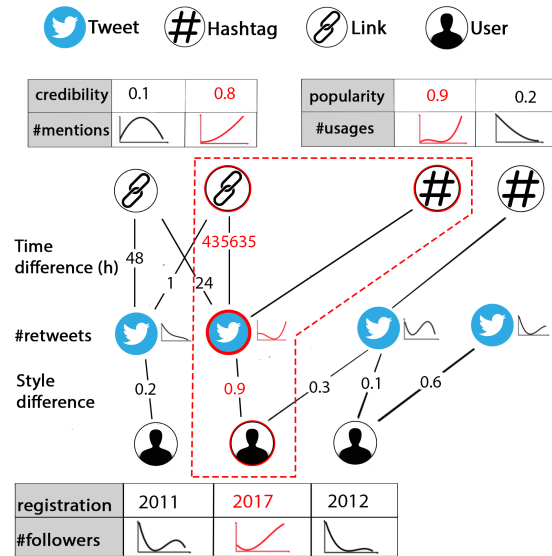


Figure 1: Multi-modal social graph

The above local anomalies provide a first signal for rumour detection. Yet, in isolation, these signals are not reliable. For instance, a user sparking a hype will also experience a sudden increase in the number of followers. We therefore need to consider *global anomalies* that comprise connected entities for which local anomalies have been observed. In the example, a rumour-related user is expected to post a rumour-related tweet, which links to a rumour-related article. Moreover, these connections between entities are also meaningful for rumour detection. For instance, in Fig. 1, the time difference between the highlighted tweet and linked article is suspicious, as is the difference between the regular linguistic style of this user (derived from past tweets) and the style of this particular tweet.

In this work, we provide the methods to realise the above idea: We exploit local anomalies and, based thereon, global anomalies among the entities of a social platform to reliably detect rumours.

3 MODEL AND APPROACH

Below, we present a model to capture entities of a social platform and their relations (§3.1). We then define the rumour detection problem (§3.2) and outline our approach to address it (§3.3).

3.1 A Model of Social Platforms

A social platform comprises many entities that are linked to each other by relations.

Entities (nodes). Our model comprises entities of specific types, i.e., modalities, such as tweets, links, users, and hashtags. Entities are modelled using feature vectors, where the features depend on the entity type. For the example in Fig. 1, each user has registration date and number of followers as features. While we limit the discussion to the above modalities in the remainder of this paper, our model is generic in the sense that further modalities such as images and videos [25] can be incorporated.

Relations (edges). Characteristics of entities in isolation are not sufficient to detect rumours. The relations between them provide a richer picture and thus can be expected to be beneficial for rumour

detection. Each relation is also modelled by a feature vector, which is specific to the the type (or modality) of the relation. For the example in Fig. 1, each tweet-article relation has the time difference between the publication times of tweets and linked articles.

Multi-modal social graph. A *multi-modal social graph*, or *social graph*, is composed of modalities, entities, and relations between entities. We denote by $D = \{D_1, \dots, D_n\}$ a set of entity types, while $V = V_1 \cup \dots \cup V_n$ is a set of entities, such that V_i is the set of entities of type D_i . Similarly, $C \subseteq [D]^2 = \{C_1, \dots, C_m\}$ is a set of relation types ($[D]^2$ being the 2-element subsets of D), $E = E_1 \cup \dots \cup E_m$ are sets of relations, where E_i is the set of relations of type C_i .

Based thereon, a social graph is defined as $G = (Q, V, E, f)$, where $Q = D \cup C$ is called the set of modalities of G . The feature information f of entities and relations is used to capture rumour signals in a social graph. Formally, $f = \{f_1, \dots, f_{n+m}\}$ is a set of mapping functions, where $f_i : Q_i \rightarrow \mathbb{R}^{q_i}$ defines an q_i -dimensional feature vector $f_i(x)$ for each element x of the modality Q_i .

The notion of a social graph enables us to address rumour detection with techniques for data management. As such, the developed algorithms are also applicable to data of social platforms that can be transformed to a graph representation [22, 41, 44, 58].

3.2 Rumour Detection

In a social graph, rumours materialise for a subset of its entities. The definition of this subset is not known, so that its identification is referred to as the rumour detection problem. That is, there is some (unknown) function that assigns truth values to entities (regular or rumourous), which shall be approximated.

PROBLEM STATEMENT. Given a social graph $G = (Q, V, E, f)$ and a ground-truth set $R^* \subseteq Q$, the rumour detection problem is to find a label function $l : Q \rightarrow \{1, 0\}$ to categorize which entities are rumourous, such that detection coefficient is maximized:

$$\frac{|R^* \cap R|}{|R^* \cup R|} \quad \text{with } R = \{x \in Q \mid l(x) = 1\}.$$

While the above definition is independent of the type of entity that is considered rumourous, in the remainder, we focus on the detection of rumourous tweets. The reason being that there is no clear-cut truth function to label other entities. For example, users may spread rumours in some tweets, but propagate regular information in others.

3.3 Approach Overview

Addressing the above problem requires us to overcome the trade-off between accuracy and completeness, which is difficult [8]. A common strategy is to first focus on completeness and subsequently optimize the accuracy of rumour detection. Filtering out false positives is often easier than finding additional true positives.

Following this line, we first strive for completeness by collecting all rumourous signals in data features: The more anomalous a feature of a tweet, the more rumourous it is. However, such a feature-based approach alone will not yield high accuracy of rumour detection. Since there is always randomness and noise in the data of a social platform, we conclude that a tweet is rumourous only if it is part of a rumourous graph structure. For example, in Fig. 1, the highlighted subgraph denotes such a structure for the




Figure 2: Rumour as Anomaly Detection Process

respective tweet, capturing rumourous context related to a user, hashtag, and linked article.

Retrieving all rumour signals from a social graph, we then reduce false positives by cross-checking between the signals, while incorporating their contexts. More precisely, we use the structural information of a social graph (i.e. relations between entities) to find a subgraph that is most rumourous. The tweets contained in this subgraph are then considered to be the actual rumour.

Rationale. Our approach is driven by the following observations:

- Identifying solely individual rumourous tweets ignores the rumour structure, i.e., it neglects that a cluster of rumourous tweets denotes a single rumour. Hence, rumour detection shall incorporate the co-occurrence of rumourous tweets as part of a rumour.
- Identifying rumours solely on the level of tweets neglects the interplay of modalities in rumour propagation. A social graph defines complex relations between entities, so that the identification of rumourous tweets, e.g., leads to the identification of rumourous users, hashtags, and links. Hence, the structure of a social graph shall be exploited to assess the propagation of rumourous information. This way, the need to detect explicit events by aggregating entities is eliminated, which is a common first step in traditional rumour detection [37].

Framework. Against this background, we design a two-step rumour detection process, illustrated in Fig. 2. In a first step, we aim to detect local anomalies in entities and relations. In a second step, these local anomalies and the relations in the graph enable the detection of rumours at the subgraph level. Below, we summarise the two steps, while their details are given in §4 and §5, respectively.

Local anomaly detection. First, we design a function that assigns an anomaly score to each entity. We argue that an anomaly scoring shall satisfy the following requirements:

- (R1) *Completeness:* In order to eliminate false negatives in rumour detection, the identification of anomalies in the data shall be comprehensive. That is, complementary angles to identify deviations from expected observations should be considered.
- (R2) *Uniformity:* For entities of all modalities, there shall be a uniform scoring domain (independent of the number of features), with a uniform ordering (lower value indicating more rumourousness), and a uniform distribution (scores are uniformly distributed in $[0, 1]$). The latter is important as thresholding for rumour detection is challenging for non-uniform distributions.
- (R3) *Non-parametric:* We assume that features follow an unknown baseline distribution. It is estimated based on the data and serves to assess the level of anomalousness per entity.

Global anomaly detection. Second, we rely on the detected local anomalies and aim at the detection of global anomalies, which indicate rumours. This shall incorporate the following requirements:

- (R4) *Cross-checking*: In order to avoid false positives, rumourousness between neighbouring entities shall be cross-checked in the social graph. As content on social platforms is dynamic and rumours may propagate very quickly, a rumourous entity is expected to affect its neighbours immediately. Hence, global anomaly detection shall consider the context of local anomalies.
- (R5) *Structuredness*: Any algorithmic solution to detect global anomalies shall acknowledge the structure of rumours. The ‘rumour-related’ parts of a social graph, in terms of rumourous information that jointly denotes a rumour, shall be detected.
- (R6) *Non-parametric*: The scoring of a global anomaly shall not assume any prior distribution of local anomaly scores. This supports multi-modality and robustness to different datasets.

4 LOCAL ANOMALY DETECTION

This section is devoted to the computation of local anomaly scores in a social graph. Guided by the above requirements (R1, R2, R3), we first show how to construct features for identifying rumours (§4.1). Then, we introduce history-based anomaly scores (§4.2) and similarity-based anomaly scores (§4.3). Based thereon, a unified anomaly score is derived for each graph element (§4.4).

4.1 Features to Identify Rumours

Feature engineering is the only domain-specific step of our approach, which we illustrate here for the case of Twitter. We distinguish history-based and similarity-based features. The former capture differences between the current and past state of an entity. The latter help to cross-check the differences between entities and relations of the same type. Specifically, we consider the following features per modality:

- User: The *registration age* and *credibility score* are considered indicators for rumours, since users spreading rumours tend to create new accounts to hide their identity. Moreover, sudden changes in the *frequency* of status updates, the number of *followers*, and the number of *#friends* may be related to rumours.
- Tweet: We consider *keywords* and the *linguistic style*. Tweets that are subjective or emotional are more likely to be rumour-related as they aim to provoke strong emotions to promote sharing. Also, the number of *retweets* may indicate rumours.
- Link: Articles linked in tweets may indicate rumours, which we assess based on the *credibility score* and *linguistic style* of the linked source and article, respectively. Furthermore, the number of *mentions* over time is used as a feature.
- Hashtag: The *popularity*, as measured by a semantic ranking [6], and sudden changes in the number of *usages* of a hashtag are expected to be rumour-related.

We further consider the features of relations between entities:

- Tweet-Link: The *time*, *location*, and *event* mentioned in a tweet may be different from the respective details given in the linked article. Also, the *linguistic style* of the tweet may be different from the one of the linked article.

- User-Tweet: The linguistic *style* of a tweet may differ from the regular style of the user.
- User-Link: The *source* linked in a tweet is anomalous.
- User-Hashtag: The hashtag is *novel*, i.e., it has not been used by the user before.
- Link-Hashtag: The hashtag has been *mentioned* in the linked article very frequently.

While some of the features are static (similarity-based), others are dynamic (history-based), so that they are derived from time snapshots using streaming APIs, such as [29]. We compute the features using established methods, whose details are described in §6.2.

Using the above features independently may lead to false positives. For instance, although rumours usually have a specific linguistic style, the reverse is not always true as, e.g., news about tragedies also adopt an emotional style. To mitigate such effects, we consider the above diverse set of features, which addresses requirement R1.

4.2 History-based Scoring

An anomaly score may be based on the differences between the current and past values of a feature vector. To this end, we establish a baseline distribution for each attribute to represent the normal behaviour, in the absence of any rumour. Then, based on the baseline distribution and the current feature values, we estimate an empirical p-value to measure the anomalousness of a feature. Aggregating these values, we assess the anomalousness of an entity or relation.

Deriving historic data. To derive historic values of features of entities or relations, we apply a temporal window. For an entity or relation x , the historic data is denoted by $X_t = \{x_1, \dots, x_t\}$, where all x_i are temporal snapshots of x . This way, historic data of the same length is considered for different history-based features of x , which enables the integration of features with varying temporal properties. Yet, t is not fixed across entities or relations, so that historic data of different lengths may be incorporated for different modalities. Note that collecting historic data is straight-forward for common platforms. Details on our data collection can be found in §6.2.

Anomaly score of a history-based feature. Our computation is based on the following null hypothesis: If there is no rumour and we select a random observation from the past, how likely is it that its value is greater than or equal to the current one? Based on historic data, the anomaly score of a feature $j \in [1, q_i]$ of an element (entity or relation) $x \in Q_i$ at timestamp t is defined as the statistical confidence degree (i.e., the p-value, the lower the better):

$$p_T(f_{i,j}(x_t)) = \frac{|\{x_r \in X_{t-1} : f_{i,j}(x_r) \geq f_{i,j}(x_t)\}|}{|X_{t-1}|} \quad (1)$$

where $f_{i,j}(x_t)$ refers to the j -th component of the feature vector $f_i(x_t)$ of an element x at timestamp t . In other words, the p-value is computed based on the number of past values $f_{i,j}(x_r)$ that are greater than the current observation $f_{i,j}(x_t)$. This is a *non-parametric* statistical measure (addressing requirement R3), since it does not assume any prior distribution on the historic data [12].

History-based anomaly score. The non-parametric p-value of an entity or relation x specifies its anomaly score based on historic observations. We aggregate these anomaly scores as follows [12]:

$$p_T(x_t) = \frac{|\{x_r \in X_{t-1} : p_{min}(x_r) \leq p_{min}(x_t)\}|}{|X_{t-1}|} \quad (2)$$

where $p_{min}(x_r) = \min_{j=1 \dots q_i} p(f_{i,j}(x_r))$. That is, at each timestamp, we compute the minimum value over all features. Then, the anomaly score $p_T(x_t)$ is the number of past minimum feature values $p_{min}(x_r)$ that are less than the current minimum feature value $p_{min}(x_t)$.

The reason for using *min* for the aggregation is to avoid false negatives, where some features are anomaly-significant, whereas others are not. Moreover, we do not consider the minimum p-value over all features at a single timestamp directly, since elements can have different numbers of features. Rather, our idea is to cross-check the scores between different timestamps across features, so that our aggregation yields *uniform* scores over all entities and relations, regardless of their modality, which addresses requirement R2 [12].

4.3 Similarity-based Scoring

Anomalousness can also be quantified by differences between entities and relations of the same type. For instance, the linguistic style of a tweet is a static property, that often lacks historic data, but may be a strong indicator of rumours. We therefore establish a baseline for features of static properties, as detailed below.

Anomaly score of a similarity-based feature. The null hypothesis of this case is summarised as: If there is no rumour, how likely does a randomly selected set of observations for a feature of different elements (entities or relations) of the same modality would have values greater than the considered element. We capture the null distribution of a feature of an element x of modality Q_i using the feature values of its peers ($x' \in Q_i$). Then, the p-value of a similarity-based feature $j = 1 \dots q_i$ of an element x is defined as follows:

$$p_S(f_{i,j}(x)) = \frac{|x' \in Q_i : f_{i,j}(x') \geq f_{i,j}(x)|}{|Q_i|} \quad (3)$$

That is, the p-value is computed based on the number of values $f_{i,j}(x')$ from other elements of the same modality that are greater than the value of the current element, $f_{i,j}(x)$. This p-value is also non-parametric (as defined by requirement R3), since it does not assume any prior distribution on the elements.

Similarity-based anomaly score. Again, based on the p-value of a similarity-based feature of an element x , the similarity-based anomaly score of x is defined as follows:

$$p_S(x \in Q_i) = \frac{|x' \in Q_i : p_{min}(x') \leq p_{min}(x)|}{|Q_i|} \quad (4)$$

where $p_{min}(x') = \min_{j=1 \dots q_i} p_S(f_{i,j}(x'))$. For each element, we compute the minimum value over all features. Then, the anomaly score of an element is the number of elements such that the minimum feature value of the current element is larger than their minimum feature values. As above, we choose *min* as an aggregation function to avoid outliers. We also aggregate across elements rather than features of a single element only. This yields *uniform* anomaly scores of elements from different modalities (requirement R2).

4.4 Unified Scoring

As both entities and relations show history-based and similarity-based features, we combine the respective anomaly scores:

$$p(x) = \min\{p_T(x), p_S(x)\} \quad (5)$$

where $p_T(x) = 1$, if x has no history-based features, and $p_S(x) = 1$, if x has no similarity-based features. Again, *min* is used in the aggregation to avoid outliers.

We note that $p_T(\cdot)$ and $p_S(\cdot)$ are uniformly distributed in $[0, 1]$ under the assumption that, in the absence of rumours, (i) the current observations are interchangeable with observations in the past; and (ii) the current observations of an element are interchangeable with observations from other elements. Based thereon, the probability that $f_{i,j}(x_r) \geq f_{i,j}(x)$ and $f_{i,j}(x') \geq f_{i,j}(x)$ is 0.5, which makes $p_T(f_{i,j}(x))$ and $p_S(f_{i,j}(x))$ follow a uniform distribution in $[0, 1]$. Also, the minimum of p-values from different features are interchangeable with past minimum values or from other peers, so that $p_T(x)$ and $p_S(x)$ are uniformly distributed in $[0, 1]$.

The *uniform* distribution of p-values is important: It enables us to handle the heterogeneity of a social graph, as different elements and modalities are mapped to the same domain of p-values. Moreover, the model facilitates the integration of multiple features for a single user, tweet, link, or hashtag, without a priori knowledge on the importance of feature for rumour detection. Finally, the overall p-value is non-parametric, since it does not assume any prior distribution, but integrates any correlation of p-values of different features.

5 GLOBAL ANOMALY DETECTION

Guided by the requirements for global anomaly detection (R4, R5, R6), we introduce the notion of an anomaly graph (§5.1), before turning to the computation of the anomalousness of a subgraph (§5.2), and the detection of a most anomalous subgraph (§5.3).

5.1 Anomaly Graph

Rumour detection using solely local information is not reliable. Local anomalies may be outliers (false positives), as features on social platforms are often noisy [29] and there are no clear-cut thresholds to filter false positives. Hence, rumour detection shall incorporate information from several elements (entities and relations) of a social graph, each providing a different view on a rumour and, thus, potentially reinforcing each other. A global view is further valuable to differentiate between anomalies that stem from the random nature of social platforms from those that originate from rumours. Finally, the propagation of rumourous information in a social graph helps to understand the rumour structure.

Formally, using the local anomaly detection, each element (entity or relation) in a social graph is associated with a p-value of being rumour-related. Given a social graph $G = (Q, V, E, f)$, this yields an anomaly graph $A = (Q, V, E, p)$, where $p : Q \rightarrow [0, 1]$ is a mapping that assign anomaly scores to entities or relations. This anomaly graph is the starting point for the identification of global anomalies, which materialise as subgraphs of the anomaly graph.

5.2 Anomalousness of a Subgraph

Rumour structure. Given an anomaly graph $A = (Q, V, E, p)$, a rumour structure is a subgraph of A that is *induced* and *connected*, which are standard graph properties [14]. Connectedness is required to cross-check anomaly scores between different elements. The subgraph shall be induced as we shall consider all relations between connected entities as a whole to eliminate false positives.

The anomalousness of a rumor structure is assessed based on:

- *Direct connections*, i.e., the relations (edges) of the graph. While both entities and relations are assigned anomaly scores, we need to conclude on the anomalousness of entities only (e.g., a tweet may be rumourous, while it is not meaningful to consider a tweet-link relation as rumourous). Hence, anomaly scores of a relation and its endpoints need to be unified.
- *Indirect connections* hold between entities that are connected by a path (of length larger than one) in the graph. The longer the path, the smaller the effect of the entities on each other, though.

Anomaly Hypergraph. To incorporate the above aspects, we propose to transform the anomaly graph to an anomaly hypergraph. The idea is to replace every two entities and the relation between them by a hypernode, which represents the collective information on the entities and the relation, while also providing an aggregated view on their anomaly scores. The hypernode inherits all further relations of the two original entities, i.e., it is connected to all entities to which the original entities had been connected. Formally, given two entities $v_1, v_2 \in V$ and a relation $e = \{v_1, v_2\} \in E$ of an anomaly graph $A = (Q, V, E, p)$, we define the respective hypernode as $v_H = \{v_1, v_2, e\}$ with an anomaly score:

$$p_H(v_H) = \max\{p(v_1), p(v_2), p(e)\} \quad (6)$$

Since $p(\cdot)$ is uniformly distributed in $[0, 1]$, $p_H(\cdot)$ also follows a uniform distribution in $[0, 1]$. Here, using *max* for aggregation reduces the chance of false positives, following requirement R4.

Processing all pairs of entities that are connected by a relation in the anomaly graph $A = (Q, V, E, p)$ as detailed above yields an anomaly hypergraph $H = (Q_H, V_H, E_H, p_H)$, with $Q_H \subset [Q]^2$ being a set of modalities, V_H being a set of hypernodes, $E_H \subseteq [V_H]^2$ being a set of edges, and p_H being a mapping function that assigns a anomaly score to each hypernode. Fig. 3 illustrates this construction.




Figure 3: Hypergraph construction

Anomalousness measurement. Using the hypergraph H , we strive for a *connected* subgraph S that shows the highest level of

anomaly. Since the hypernodes already include the original relations, it is straightforward to revert a subset of connected hypernodes to an induced connected subgraph of the original anomaly graph.

To this end, we first measure the anomalousness of a subgraph, acknowledging the structure of rumours, see requirement R5. We employ the idea of scan statistics [24], which computes the statistical significance of a subgraph S being anomalous without assuming any prior distribution of the subgraph [12]:

$$P(S) = \max_{0 < \alpha \leq \alpha_{max}} \phi(\alpha, |V_\alpha(S)|, |V(S)|) \quad (7)$$

where α_{max} is the maximum statistical significance level ($\alpha_{max} = 0.05$ indicates that the value is *at least* 95% statistical significant), $V(S)$ is the node set of S , $V_\alpha(S) = \{v \in V(S) : p_H(v) \leq \alpha\}$ is the set of nodes in S with anomaly scores that are significant at the confidence level $\alpha > 0$.

To maximize the detection coefficient (see §3.2), function $\phi(\cdot)$ shall favour the propagation of rumours, meaning that ‘insignificant’ nodes ($V(S) \setminus V_\alpha(S)$) are also accepted as long as they are connected with enough ‘significant’ entities ($V_\alpha(S)$). This is motivated by the dynamic nature of a rumour: Anomaly scores of rumourous entities vary over time and may not be significant at the same time. Moreover, function $\phi(\cdot)$ shall be *non-parametric* (requirement R6), i.e., a function that compares the observed number of α -significant p-values $|V_\alpha(S)|$ to the expected number of α -significant p-values $\mathbb{E}[|V_\alpha(S)|]$. Since our p-values are uniformly distributed in $[0, 1]$, we have $\mathbb{E}[|V_\alpha(S)|] = \alpha|V(S)|$. Therefore, we can directly compare $|V(S)|$ and $|V_\alpha(S)|$ as follows [7]:

$$\phi(\alpha, |V_\alpha(S)|, |V(S)|) = |V(S)| \times KL\left(\frac{|V_\alpha(S)|}{|V(S)|}, \alpha\right) \quad (8)$$

where KL is the Kullback-Leibler divergence defined as $KL(x, y) = x \log(x/y) + (1-x) \log(\frac{1-x}{1-y})$. Since $KL(x, y) \geq 0$, it follows that $P(S) \geq 0$ (the higher, the more anomalous). Based thereon, our goal is to detect subgraphs as large as possible (via $|V(S)|$), that have a high confidence level of anomalousness (via $|V_\alpha(S)|/|V(S)|$).

5.3 Detection of a Most Anomalous Subgraph

Detecting a rumour structure in an anomaly graph $A = (Q, V, E, p)$ is equivalent to finding a connected subgraph with maximal anomalousness in the anomaly hypergraph $H = (Q_H, V_H, E_H, p_H)$:

$$\arg \max_{S \in \mathcal{S}(H)} P(S) \quad (9)$$

where $\mathcal{S}(H)$ contains all possible connected subgraphs of H .

As the above problem is computationally expensive [12], we develop an approximation solution that scales to real-world social graphs. In the context of online social platforms, we argue that such a detection algorithm needs to satisfy two additional requirements:

- *Extensibility.* In practice, multiple rumours may occur at the same time. Hence, we consider a threshold as a relaxation parameter. We then aim at detecting all subgraphs in the anomaly graph that have an anomalousness value above this threshold. Such a threshold may be set based on rumours detected and verified in the past.




Figure 4: Illustration of Alg. 1

Algorithm 1: Anomalous Subgraphs Detection

input : An anomaly graph $A = (Q, V, E, p)$,
a retain threshold τ (for streaming version),
a coverage level of anomaly K (default = 5),
a specified number of hops Z (default = $\log(|V|)$)

output: A sorted list of subgraphs \mathbb{S}

- 1 Construct anomaly hypergraph $H = (Q_H, V_H, E_H, p_H)$ from A ;
- 2 Sort the nodes in H by anomaly score;
- 3 $\alpha_{max} = 0.05, \mathbb{S} = \mathbb{C} = \emptyset$;
- 4 **for** $q \in [1, \dots, |Q_H|]$ **do**
- 5 **for** $k \in [1, \dots, K]$ **do**
- 6 $R = \{v_k\}, v_k$ is the k -th most anomalous node in V_H of modality q ;
- 7 **for** $z \in \{1, \dots, Z\}$ **do**
- 8 $H' = \{v \in V_H \setminus R : \exists v' \in R, \{v, v'\} \in E_H\}$;
- 9 $\langle S, P(S) \rangle = \text{bestNeighbourhood}(H', R, \alpha_{max})$;
- 10 **if** $S \setminus R \neq \emptyset$ **then** $R = S$;
- 11 **else break**;
- 12 $\mathbb{S} = \mathbb{S} \cup \{R\}$;
- 13 **for** $S \in \mathbb{S}$ **do**
- 14 **if** $P(S) \geq \tau$ **then** $\mathbb{C} = \mathbb{C} \cup \{S\}$; // candidate rumours
- 15 **return** \mathbb{S} ;

- *Incremental processing*: To cope with continuous data generated by social platforms, detection shall be incremental, incorporating new data as it arrives.

An Extensible and Incremental Algorithm. Due to the inherent complexity of Eq. 9, we present an approach to approximate a solution, see Alg. 1 (extended from [12]). It takes as input an anomaly graph and a detection threshold, and returns a sorted list of the most anomalous subgraphs that satisfy the threshold. The solution to Eq. 9 is simply the top-1 in the list. Moreover, in the light of the rumour detection problem (§3.2), only the tweet nodes of the output graph may be considered. Since multiple rumours may spread simultaneously on social platforms, however, we include a coverage level K as an input parameter, to cover rumours with smaller anomalousness values.

Our algorithm first expands the subgraphs from a seed node to their neighbours, before greedily optimising the anomaly score for the subgraphs. Specifically, we construct a hypergraph H (line 1), in which each hypernode has an anomaly score, as detailed above. We sort the hypernodes by these scores as this later improves the run-time of the scan statistics subproblem. We then select a root node (line 6), determine its neighbourhood (line 8), and find the subgraph in this neighbourhood with the highest anomaly score (line 9) (extended from [31]). The latter greedily retains nodes in the increasing order of p-values (the smaller, the better). Then, we continue to expand the subgraph until our root node set is equal to the most anomalous node set (line 10), i.e., it cannot be expanded

further to increase the anomaly score. This guarantees that the subgraph is connected and its anomaly score is maximal [12].

Fig. 4 illustrates the core step of extending the neighbourhood of a root node and finding the optimal subgraph in Alg. 1 (line 6- 10).

6 EMPIRICAL EVALUATION

We evaluated our approach with a large real-world dataset obtained from Twitter. Below, we introduce our experimental setting (§6.1), data collection methodology (§6.2). We show that our approach outperforms baseline methods for rumour detection in terms of effectiveness (§6.3).

6.1 Experimental Setting

Metrics. We use the following evaluation metrics:

- The detection *coefficient*, first proposed in [40], can be seen as a combination of precision and recall applied to a graph setting. R^* is defined as the set of rumour-related entities, whereas R is the set of entities labelled by a rumour detection technique. Then, the measure is defined as:

$$\text{Coefficient} = \frac{|R^* \cap R|}{|R^* \cup R|}$$

Baselines. State-of-the-art rumour detection [61] is not applicable in our context, as it aims at learning a classification model based on a collection of entities that have been labelled with rumours. Such a collection is typically extracted by a pre-processing step that crawls the data related to a particular event, thereby assuming that the extracted elements can be labelled accordingly. As a result, the performance of these approaches strongly depends on the accuracy of such pre-processing [10, 26]. In our work, we progressively detect rumour-related entities by scanning abnormal signals (entities with high anomaly scores) in the social graph.

This fundamental difference in the taken approach is also reflected in the employed evaluation measures. Existing rumour detection techniques are evaluated using machine learning metrics, applied per rumour. This is not possible for our approach, so that we rely on the detection coefficient, applied per graph entity. In a broad sense, most rumour detection techniques focus on maximizing accuracy, instead of striving for a balance of accuracy and completeness.

Against this background, we consider several baseline methods. We implemented these methods based on the respective papers.

- *Decision* [9]: A decision tree classifier that is based on the Twitter information credibility model. The decision tree is constructed based on several hand-crafted features.
- *Nonlinear* [50]: An SVM-based approach that uses a set of hand-crafted features, selected for the tweets to classify.
- *Rank* [59]: A rank-based classifier that aims to identify rumours based on enquiry tweets.

In addition, we also compare our approach with methods based on homogeneous graphs that contain only a single modality. For instance, a tweet graph contains only tweets, while edges between tweets represent that tweets stem from the same user, have retweet relations, or share a keyword. We constructed four such homogeneous graphs, for users, tweets, links, and hashtags, respectively.

Parameters. We set the statistical significance level $\alpha_{max} = 0.05$ (i.e. the result is guaranteed to be *at least* 95% confidence). The coverage level K in Alg. 1 has been varied, so that we can detect multiple rumours at the same time.

Experimental environment. All results have been obtained on an Intel Core i7 system (2.8 Ghz, 32GB RAM).

6.2 Data Collection

Rumour collection. Snopes is a world-leading rumour-debunking service. Unlike other organizations such as Politifact and Urbanlegends, it is considered to be objective when evaluating the veracity of rumours [2, 45]. Snopes editors investigate each rumour along different dimensions and provide an argumentative report as shown in Table 1. For example, the claim describes the rumour succinctly and the rating represents its truth value according to the fact-checker.

Table 1: Information about a rumour.

Attribute	Example
id	trump-aid-puerto-rico
date	10/2/2017
genesis tweet	[...] President Trump has dispatched 140 helicopters [...]
sources of veracity	press reports, local officials, organizations
rating	MIXTURE [3]

Multi-model social graph construction. Twitter is a large social platform with tweets covering various domains such as politics and crime. It is frequently used by users to express their opinions in a timely manner, e.g., by retweeting others, which provides insights into how rumours propagate. These characteristics make Twitter data particularly suitable for evaluating rumour detection methods.

Datasets. The collected data comprises 4 million tweets, 3 million users, 28893 hashtags, and 305115 linked articles, revolving around 1022 rumours from 01/05/2017 to 01/11/2017. This period was chosen as it contains several rumours, e.g., related to the Las Vegas shooting and information published by the US administration. Our data spans over 20 different domains, available at [5]. Here, we report results for the most popular ones:

- *Politics*: rumours related to all political issues.
- *Crime*: rumours related to criminology and incidents, such as the Las Vegas shooting.

Each of the datasets is a full view of the social graph. The modelled entity types, relation types, and features are summarised in §4.1.

6.3 Effectiveness of Rumour Detection

We evaluate the detection coefficient of our approach versus the baseline methods in Fig. 5 for the domains *Politics* and *Crime* (the same trends emerge for the other domains). We vary the amount of rumours contained in the dataset, i.e. data sparsity, by randomly removing some rumours, so that the remaining rumours cover 30%, 60%, 100% of the original count.

In general, our approach outperforms the baseline methods in the detection of rumour-related tweets. For instance, taking the results of the *Politics* dataset, when considering 30% of the rumours, our approach achieves a coefficient of 0.82, whereas the best baseline method achieves solely a coefficient of 0.62.




Figure 5: Rumour Detection Coefficient across datasets

7 RELATED WORK

Rumour detection. While there is a large body of work on rumour detection on social platforms, surveyed in [60], little has been done to exploit multiple modalities to detect rumours. Most work leverages only textual data such as tweets [9, 17, 59]; whereas others consider different data entities such as users and hashtags but still treat them as additional features or textual data only [15, 20, 27]. Techniques based on hand-crafted features [9, 32, 33, 50, 52, 55, 59] are grounded in an ad-hoc definition of features, which are expected to be strong indicators of rumours. Recently, deep features based on temporal dependencies of the posts have been proposed [27]. While this approach achieves high detection accuracy, it first requires the detection of an explicit event and thus depends on the accuracy of this event detection step [54]. There are further approaches [28, 48] that take into account how rumours propagate. However, these techniques require large collections of tweets to conduct the respective analysis. As such, they cannot be expected to yield small lag times in the detection of rumours and are not well-suited for a streaming setting. Our approach is the first to leverage not only the textual data, but also other modalities.

Anomaly detection. Anomaly detection can be classified into point or group-based techniques [57]. Point-based anomaly detection aims to detect individuals, for which the behaviour is different from the general population [21, 23, 38]. Group-based anomaly detection, in turn, strives for groups of individuals that collectively behave differently compared to some population [11–13, 30, 49, 56]. However, none of the above techniques has been applied to rumour detection. While [12] addresses a similar use case, it neglects the anomalies related to feature differences between entities. Our technique is the first one for group-based anomaly detection that simultaneously identify anomalies in all features, entities, and relations. Most of the work on anomaly detection in general and rumour detection in particular focuses on accuracy. Here, we define the detection coefficient to capture the balance between accuracy and completeness, which is optimised by our approach.

Information networks. There exists various graph-based models for data of social platforms, referred to as information networks [18, 34, 39, 46, 53]. Some models capture real-world entities, such as users and posts [42], while others represent derived data elements, such as topics [43]. Existing work on anomaly detection in information networks focuses on modelling the propagation patterns of known phenomena [16, 61] or classifies known events [19, 59]. This setting is orthogonal to our work, since we strive for the detection of phenomena that emerge on social networks, but are not known a priori.

8 CONCLUSION

This paper proposed an approach for rumour detection that is grounded in the anomalies of a social graph. Unlike traditional approaches that focus only on accuracy, we optimised the detection coefficient, which represents the trade-off between accuracy and completeness. We presented a two-step detection approach that detects anomalies at the local and global level. While the former increases the completeness of detection by reducing false negatives, the latter optimises the detection accuracy by reducing false positives. Our experiments showed that our method is effective and efficient, detecting rumours early and accurately.

REFERENCES

- [1] [n. d.]. <https://www.engadget.com/2018/08/21/facebook-rates-user-trustworthiness/>.
- [2] [n. d.]. <https://www.networkworld.com/article/2235277/data-center/data-center-fact-checking-the-fact-checkers-snoopes-com-gets-an-a.html>.
- [3] [n. d.]. <https://www.snoopes.com/fact-check/trump-aid-puerto-rico/>.
- [4] [n. d.]. <https://www.theverge.com/2018/8/21/17763886/facebook-trust-ratings-fake-news-reporting-score>.
- [5] [n. d.]. <http://tiny.cc/pls2qy>.
- [6] Piyush Bansal, Somay Jain, and Vasudeva Varma. 2015. Towards semantic retrieval of hashtags in microblogs. In *WWW*. 7–8.
- [7] Robert H Berk and Douglas H Jones. 1979. Goodness-of-fit test statistics that dominate the Kolmogorov statistics. *Probability theory and related fields* (1979), 47–59.
- [8] Michael Buckland and Fredric Gey. 1994. The relationship between recall and precision. *Journal of the American society for information science* 45, 1 (1994), 12–19.
- [9] Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *WWW*. 675–684.
- [10] Sylvie Cazalens, Julien Leblay, Philippe Lamarre, Ioana Manolescu, and Xavier Tannier. 2018. Computational fact checking: a content management perspective. *PVLDB* 11, 12 (2018), 2110–2113.
- [11] Varun Chandola, Arindam Banerjee, and Vipin Kumar. 2007. Outlier detection: A survey. *Comput. Surveys* (2007).
- [12] Feng Chen and Daniel B Neill. 2014. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *KDD*. 1166–1175.
- [13] Kaustav Das, Jeff Schneider, and Daniel B Neill. 2009. *Detecting anomalous groups in categorical datasets*. Carnegie Mellon University.
- [14] Reinhard Diestel. 2018. *Graph theory*. Springer Publishing Company, Incorporated.
- [15] Chi Thang Duong, Quoc Viet Hung Nguyen, Sen Wang, and Bela Stantic. 2017. Provenance-based rumor detection. In *ADC*. 125–137.
- [16] Adrien Friggeri, Lada A Adamic, Dean Eckles, and Justin Cheng. 2014. Rumor Cascades. In *ICWSM*.
- [17] Aditi Gupta, Ponnurangam Kumaraguru, Carlos Castillo, and Patrick Meier. 2014. Tweetcred: Real-time credibility assessment of content on twitter. In *SocInfo*. 228–243.
- [18] Nguyen Quoc Viet Hung, Hoyoung Jeung, and Karl Aberer. 2013. An Evaluation of Model-Based Approaches to Sensor Data Compression. *TKDE* (2013), 2434–2447.
- [19] Nguyen Quoc Viet Hung, Nguyen Thanh Tam, Ngoc Tran Lam, and Karl Aberer. 2013. An Evaluation of Aggregation Techniques in Crowdsourcing. In *WISE*. 1–15.
- [20] Nguyen Quoc Viet Hung, Duong Chi Thang, Matthias Weidlich, and Karl Aberer. 2015. Minimizing Efforts in Validating Crowd Answers. In *SIGMOD*. 999–1014.
- [21] Alexander Ihler, Jon Hutchins, and Padhraic Smyth. 2006. Adaptive event detection with time-varying poisson processes. In *KDD*. 207–216.
- [22] Xin Jin, Cindy Xide Lin, Jiebo Luo, and Jiawei Han. 2011. SocialSpamGuard: A Data Mining-Based Spam Detection System for Social Media Networks. *PVLDB* 4, 12 (2011), 1458–1461.
- [23] Wen-Hua Ju and Yehuda Vardi. 2001. A hybrid high-order Markov chain model for computer intrusion detection. *Journal of Computational and Graphical Statistics* 10, 2 (2001), 277–295.
- [24] Martin Kulldorff. 1997. A spatial scan statistic. *Communications in Statistics-Theory and methods* (1997), 1481–1496.
- [25] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a social network or a news media?. In *WWW*. 591–600.
- [26] Julien Leblay, Ioana Manolescu, and Xavier Tannier. 2018. Computational fact-checking: Problems, state of the art, and perspectives. In *The Web Conference*.
- [27] Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *IJCAL*. 3818–3824.
- [28] Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning. In *ACL*. 708–717.
- [29] Fred Morstatter, Jurgen Pfeffer, Huan Liu, and Kathleen M Carley. 2013. Is the Sample Good Enough? Comparing Data from Twitter’s Streaming API with Twitter’s Firehose. In *ICWSM*.
- [30] Krikamol Muandet and Bernhard Schölkopf. 2013. One-class support measure machines for group anomaly detection. *arXiv preprint arXiv:1303.0309* (2013).
- [31] Daniel B Neill. 2012. Fast subset scan for spatial pattern detection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74, 2 (2012), 337–360.
- [32] Quoc Viet Hung Nguyen, Kai Zheng, Matthias Weidlich, Bolong Zheng, Hongzhi Yin, Thanh Tam Nguyen, and Bela Stantic. 2018. What-if Analysis with Conflicting Goals: Recommending Data Ranges for Exploration. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 89–100.
- [33] Thanh Tam Nguyen, Quoc Viet Hung Nguyen, Kai Zheng, Matthias Weidlich, Bolong Zheng, Hongzhi Yin, and Bela Stantic. 2018. *Exploring Data Partitions for What-if Analysis*. Technical Report. EPFL.
- [34] Thanh Toan Nguyen, Thanh Cong Phan, Chi Thang Duong, Quoc Viet Hung Nguyen, and Bela Stantic. 2017. *Probabilistic Schema Covering*. Technical Report. EPFL.
- [35] Alexandra Olteanu, Carlos Castillo, Nicholas Diakopoulos, and Karl Aberer. 2015. Comparing events coverage in online news and social media: The case of climate change. In *ICWSM*. 288–297.
- [36] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. CrisisLex: A lexicon for collecting and filtering microblogged communications in crises. In *ICWSM*. 376–385.
- [37] Vu Pham, Théodore Bluche, Christopher Kermorvant, and Jérôme Louradour. 2014. Dropout improves recurrent neural networks for handwriting recognition. In *ICFHR*. 285–290.
- [38] Matthias Schonlau, William DuMouchel, Wen-Hua Ju, Alan F Karr, Martin Theus, and Yehuda Vardi. 2001. Computer intrusion: Detecting masquerades. *Statistical science* (2001), 58–74.
- [39] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and S Yu Philip. 2017. A survey of heterogeneous information network analysis. *TKDE* 29, 1 (2017), 17–37.
- [40] Skyler Speakman, Yating Zhang, and Daniel B Neill. 2013. Dynamic pattern detection with temporal consistency and connectivity constraints. In *ICDM*. 697–706.
- [41] Yizhou Sun, Charu C Aggarwal, and Jiawei Han. 2012. Relation strength-aware clustering of heterogeneous information networks with incomplete attributes. *PVLDB* 5, 5 (2012), 394–405.
- [42] Yizhou Sun and Jiawei Han. 2012. Mining heterogeneous information networks: principles and methodologies. *Synthesis Lectures on Data Mining and Knowledge Discovery* 3, 2 (2012), 1–159.
- [43] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale information network embedding. In *WWW*. 1067–1077.
- [44] Io Taxisidou and Peter Fischer. 2013. Realtime analysis of information diffusion in social media. *PVLDB* 6, 12 (2013), 1416–1421.
- [45] Nguyen Thanh Tam, Matthias Weidlich, Hongzhi Yin, Bolong Zheng, Nguyen Quoc Viet Hung, and Bela Stantic. 2019. User Guidance for Efficient Fact Checking. *PVLDB* 12, 8 (2019), 850–863.
- [46] Nguyen Thanh Toan, Phan Thanh Cong, Duong Chi Thang, Nguyen Quoc Viet Hung, and Bela Stantic. 2018. Bootstrapping uncertainty in schema covering. In *ADC*. 336–342.
- [47] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- [48] Ke Wu, Song Yang, and Kenny Q Zhu. 2015. False rumors detection on sina weibo by propagation structures. In *ICDE*. 651–662.
- [49] Liang Xiong, Barnabás Póczos, Jeff G Schneider, Andrew Connolly, and Jake VanderPlas. 2011. Hierarchical probabilistic models for group anomaly detection. In *AISTATS*. 789–797.
- [50] Fan Yang, Yang Liu, Xiaohui Yu, and Min Yang. 2012. Automatic detection of rumor on Sina Weibo. In *KDD*. 13.
- [51] Junting Ye, Santhosh Kumar, and Leman Akoglu. 2016. Temporal opinion spam detection by multivariate indicative signals. In *ICWSM*. 743–746.
- [52] Hongzhi Yin, Hongxu Chen, Xiaoshuai Sun, Hao Wang, Yang Wang, and Quoc Viet Hung Nguyen. 2017. SPTF: A Scalable Probabilistic Tensor Factorization Model for Semantic-Aware Behavior Prediction. In *ICDM*. 585–594.
- [53] Hongzhi Yin, Zhiting Hu, Xiaofang Zhou, Hao Wang, Kai Zheng, Nguyen Quoc Viet Hung, and Shazia Wasim Sadiq. 2016. Discovering interpretable geo-social communities for user behavior prediction. In *ICDE*. 942–953.
- [54] Hongzhi Yin, Nguyen Quoc Viet Hung, Zi Huang, and Xiaofang Zhou. 2018. Joint Event-Partner Recommendation in Event-based Social Networks. In *ICDE*. 1–12.
- [55] Hongzhi Yin, Xiaofang Zhou, Bin Cui, Hao Wang, Kai Zheng, and Nguyen Quoc Viet Hung. 2016. Adapting to User Interest Drift for POI Recommendation. *TKDE* (2016), 2566–2581.
- [56] Rose Yu, Xinran He, and Yan Liu. 2015. Glad: group anomaly detection in social media analysis. *TKDD* 10, 2 (2015), 18.

- [57] Rose Yu, Huida Qiu, Zhen Wen, ChingYung Lin, and Yan Liu. 2016. A survey on social media anomaly detection. *ACM SIGKDD Explorations Newsletter* 18, 1 (2016), 1–14.
- [58] Fan Zhang, Wenjie Zhang, Ying Zhang, Lu Qin, and Xuemin Lin. 2017. OLAK: an efficient algorithm to prevent unraveling in social networks. *PVLDB* 10, 6 (2017), 649–660.
- [59] Zhe Zhao, Paul Resnick, and Qiaozhu Mei. 2015. Enquiring minds: Early detection of rumors in social media from enquiry posts. In *WWW*. 1395–1405.
- [60] Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2017. Detection and Resolution of Rumours in Social Media: A Survey. *arXiv preprint arXiv:1704.00656* (2017).
- [61] Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *CSUR* 51, 2 (2018), 32.