Behavior Mining Methods for Dynamic Risk Analysis in Social Media Communication

Jan Ole Berndt

TriLabS@CIRT, Wirtschaftsinformatik I, Universität Trier, Germany
berndt@uni-trier.de

Abstract. Spreading information through social media can be beneficial in crisis situations as well as harmful for a person's or company's reputation. Which and how information is spread depends on network structures and individual behaviors. Simulation-based methods are suitable to systematically explore potential system behavior, its potentials, and its risks. This paper introduces behavior mining from social media to develop such simulations. It provides an integrated analysis workflow and gives an overview of applicable methods for each process step.

Keywords: Social Media · Risk Analysis · User Behavior · Data Mining.

1 Introduction

Communication in social media plays a major role for crisis, risk, and reputation management. Crucial information as well as misinformation spreads rapidly through online social networks which can be either beneficial (e.g., for crisis management) or harmful (e.g., endangering reputation) [1, 23]. In fact, these processes can lead to severe cascading effects in economy, politics, and other domains [3, 25]. Consequently, understanding communication processes in social media is crucial for making the right decisions in the event of a crisis.

Simulation-based methods are particularly suitable to systematically explore potential system behavior which emerges from individual interactions between media users [5]. The dynamics in social networks are highly dependent on network structures and individual behaviors. Even for very simple behavioral patterns, the overall result on the system level can become completely unpredictable [11]. Therefore, it is necessary to use realistic user interconnections and communication patterns in social media simulation for obtaining meaningful results. While there is a wealth of network analysis and computer linguistic approaches to social media, less research has been conducted to infer patterns of actor behavior from social media data.

This contribution introduces behavior mining from social media data. Its goal is to provide an overview of available methods and a workflow incorporating them to develop simulations as a method for dynamic risk analysis in networked communication. To that end, Section 2 further elaborates on systemic risks in social media and Section 3 presents the process of behavior mining with a discussion of available methods for each process step. Finally, Section 4 concludes on the findings of this paper.

2 Analyzing Systemic Risks in Social Media

In social media, users are interconnected in complex networks of friendship, acquaintance, and general interest. Formally, these networks can be modeled as graphs with media users as nodes and their connections as edges [27]. Communication takes place along these edges and becomes visible to nodes connected to the active user. For instance, Twitter[1] users can follow each other, resulting in all followers of a specific user to get notified about that user's activities. By reacting to or forwarding messages, communication cascades through the network.

From the perspective of systemic risks, the effects of cascading communication can be either beneficial or hazardous. Consequently, existing work on social media and risks focuses on the following two main lines of research.

1. Social media usage for risk management and crisis communication
2. Social media endangering reputation or spreading misinformation

As a means for risk management, social media provide communication infrastructures which allow for spreading information rapidly to those who are affected by a crisis. Prime examples are disasters like earthquakes, epidemic diseases, or plane crashes [26, 15]. In these instances, it is necessary to provide information about the situation and advice for recommended action to those immediately endangered, their relatives, as well as any helpers. In addition, planners and risk managers need to listen to their audience and take account of their concerns in crisis communication [28]. Thus, risk and crisis management requires channels for multi-directional communication with large numbers of people.

Using social media has been proposed for providing information and monitoring the situation in the event of a crisis [1]. To that end, researchers have derived best practices for pre-event management, collaboration with the public, and communication strategies from real-world examples of successful social media usage in crisis situations [28]. To support these, sentiment analysis and emotion detection techniques are available for monitoring public opinion and adapting communication strategies accordingly [17, 15]. These techniques are complemented with static analyses of the underlying network structures in social media to identify the most influential users who can serve as multipliers for spreading information [10]. Information spreads throughout a network in a process of so-called social contagion [20]. The most common approach to analyze such a process is the SIR model which groups users into those being potentially attentive to information (or: susceptible, S), those actively spreading the information (or: infected, I), and those being already informed but no longer active (or: recovered, R) [8]. Based on a given network structure, this model allows for dynamic analyses of information flows in social media by means of simulation. Moreover, such an approach facilitates optimizing communication policies by identifying the best set of users to be informed first in the event of a crisis for maximizing the reach of that information. This kind of influence maximization has been applied to various fields including marketing and public health [16, 31].

[1] https://twitter.com/

Nevertheless, crisis communication using social media also bears risks in itself. Among other challenges, it is important to prevent rumors and the spreading of misinformation which can lead to panic, to avoid information overload, and to ensure that information cannot be abused for criminal purposes [30]. This requires careful planning of whom to address with which information in what way. Even outside crisis situations, these are challenges companies and individuals face when utilizing social media for marketing, political communication, or personal interests. While the speed of information diffusion can be beneficial, it can also become harmful and potentially uncontrollable in case of rumors and misinformation being spread (fake news) or in mass protests (storms of protest or Twitterstorms) [2, 23]. These phenomena endanger a company's or person's reputation and can even lead to severe systemic effects in economy and politics, ranging from a loss of revenue to social upheaval [3, 25].

In order to understand the aforementioned phenomena, computer simulation has been proposed as a method for dynamic analysis of network effects and individual behavior [5]. Such a simulation can be used to develop communication strategies by anticipating potential reactions and their effects throughout a social network. However, these reactions can hardly be captured by simplified contagion models. When attempting to affect the content and flow of information, individual motivations and behavioral dispositions that drive communication must be considered. Depending on the composition of these individual behaviors, simulated communication processes can vary drastically between negotiations of differing opinion as well as pure protest and resentment [21]. Consequently, simulations for analyzing and addressing systemic risks in social media require representations of communicative behavior to produce meaningful results. To that end, the following section introduces behavior mining for extracting individual communicative patterns from social media data as a foundation for simulation.

3 Behavior Mining from Social Media Communication

To generate realistic representations of communicative behavior, a method is required to identify and derive this kind of behavior from social media data. Indeed, there are various approaches readily available to mine such data for different purposes. In particular, event mining methods are popular for risk assessment applications in which, e.g., social unrest is predicted [12]. These methods focus on user groups, combinations of communicated contents, as well as communication frequencies that indicate the targeted activities. Additionally, particular user groups are identified according to their influencing potential for customer relationship management and marketing purposes [29]. This can be achieved by clustering social graphs in which these groups occur as densely interconnected users with similar interests. For both applications, geo-spatial clustering is used to narrow down a particular area of events or marketing activities [4].

While event mining and geo-spatial clustering are highly relevant for risk identification, the methods applied in that context are less useful for extracting individual patterns of communication. A specific event or location influences the

Feature Extraction	Data Preprocessing	User Clustering	Behavior Extraction
Content-based	Profiling	Centroid-based	Prototypical user
Metadata-based	Outlier detection	Density-based	Inductive logic programming
		Distribution-based	Decision-tree learning
		Hierarchical	

Table 1. Behavior mining process (left to right) with methods for each process step.

topic of certain conversations, but that topic can be discussed in different ways, depending on the behavioral patterns of participating users [21]. Therefore, another approach is necessary. Such an approach to behavior mining must identify particular groups of users with similar behaviors and then extract these behaviors as decision-making rules for actors in a simulation. However, in order to achieve this, features to describe and discriminate behavioral patterns have to be identified. This results in a behavior mining process consisting of the following four consecutive steps and the associated methods shown in Table 1.

1. **Feature Extraction:** Identify properties of communication processes that allow for distinguishing between different behavioral patterns.
2. **Data Preprocessing:** Prepare the available data for automated analysis of user activities according to the identified features.
3. **User Clustering:** Group users according to similarities between their communicative behaviors using the extracted features.
4. **Behavior Extraction:** Identify prototypical behaviors for each identified user group in the form of condition-action rules.

The following sections outline those tasks and discuss available methods for each of these process steps in detail with respect to risk analysis in social media.

3.1 Feature Extraction

Features for characterizing user behavior can either be based on communication content or on metadata about the interaction process. The former includes *discourse topics* and *sentiment* expressed throughout a conversation [17]. To extract these features from raw data, computer linguistic methods are required. The most simple way to model topics in social media is using hashtags with which users provide keywords for describing their contributions. Since not all messages contain hashtags, they can be complemented with other distinctive words. Their occurrence frequencies in different messages form a topic model for a conversation [6]. On social media platforms like Twitter, more structured arguments are rare [5]. Hence, subjects that attract an individual's attention and their opinion toward it suffice as content-based features in most cases.

While content-based features are important for analyzing different types of discourses, a wide range of activity patterns can simply be observed in metadata-based analyses. Social media metadata covers the activities of all observed users in an abstracted form. It includes the *time* a message is sent or published, its *sender*, potential *receivers* based on the underlying social network graph, any explicitly mentioned *addressees* (e.g., so-called @-mentions), as well as the *type of message* (i.e., an original contribution, a reply to another message, or a forwarded message) [7]. This data is readily available both through programming interfaces of social media platforms and in existing data sets for social media analysis [5, 19]. From metadata, composite features like inter-activity times, activity type frequencies, and activity type sequences can be derived. These specify patterns of behavior which correspond to prototypical actor types and roles in social media.

3.2 Data Preprocessing

Before user behaviors can be extracted, social media data must be preprocessed. Gathering data or using existing data sets results in a collection of single communication events or a graph of interconnected events and users. The aforementioned features either describe the nature of individual events or they derive statistics across several of them. Consequently, the data must be profiled to bring the extracted features into a *one line per user* format. In that format, all features are listed for the corresponding user who's behavior they characterize. This is a prerequisite for identifying similarities and differences between users with respect to the extracted features of their communicative behavior.

However, there can be a wide range of user behaviors out of which some might be incomparable to others. These can disturb further analyses because they will not fit into any other group of behaviors. For instance, unusually small inter-activity times are an indicator for either a *corporate account* or a *social bot* [9]. While such accounts have the potential to crucially impact communication in social media, it is impossible to sensibly fit a single instance of them into any other user group. Hence, it is necessary to detect such outliers and exclude them from the clustering step [14]. They can still be included again later as individual *single entity clusters* to analyze their special impact.

3.3 User Clustering

User clustering groups social media users with common characteristics together while distinguishing between groups that differ with respect to one or more features. The resulting clusters reflect specific communicative roles which users adopt or particular topics they are interested in. These aspects drive the communication process and have crucial impact on information diffusion. For example, a user adopting the role of a *producer* will primarily introduce original content that can start communication or steer the process into new directions. Contrastingly, *communicators* and *networkers* will add to the existing content and primarily distribute information [13]. Hence, these user groups exhibit characteristic behavior patterns by which they can be identified and distinguished.

To derive characteristic patterns from a data set, *centroid-based clustering* is particularly suitable. This method groups data points around a prototypical instance representing common and distinguishing features of the group [18].

Consequently, it allows for directly identifying typical behavioral patterns of, e.g., a *producer* or a *networker* among the potentially wide variations of those behaviors. Even if users are similar in their activities and underlying motivations, they still differ from each other. These variations produce noise which *density-based clustering* as well as *distribution-based clustering* methods can handle. Density-based algorithms distinguish between high and low density regions, whereas distribution-based approaches attempt to match data points to statistical distributions [18]. However, sometimes several user groups can be characterized by similar behaviors while differing in particular aspects; i.e., they overlap to a certain extent. For instance, *communicators* and *networkers* can differ in their communication contents, but they both mainly distribute information [13]. When overlapping user behaviors can be subsumed under more generic groups, *hierarchical clustering* is useful. This method creates a taxonomy of groups by starting with each data point as an individual cluster and then grouping similar ones together [18]. Nonetheless, its computational complexity makes it difficult to apply that method to the large amounts of data present in social media analysis. Thus, which clustering method is applied best highly depends on the analyzed communication process and the user population participating in it.

3.4 Behavior Extraction

While user clustering identifies *what* behavioral patterns exist in social media, it does not derive *how* these patterns are generated. Therefore, behavior extraction has the task to find decision rules which map specific events to communicative actions. These mappings then determine under which conditions, e.g., a message is forwarded or replied to. Hence, they govern whether crucial information can reach its target audience, whether misinformation can spread, whether a mass protest can emerge, or whether nothing significant happens.

To extract behaviors, specific events must be identified that lead to the same reactions within a user group. If a centroid-based clustering method was used before, a prototypical user for each behavior pattern is already available. That user's activities must be grouped by their types or the topics they refer to. Then, they can be related to circumstances under which they occur. For instance, messages from particular other users that are regularly forwarded or topics that frequently provoke replies to observed contents. Rules for such behaviors can be derived by *inductive logic programming* which hypothesizes on logical inference rules based on background knowledge and a set of examples [22]. For social media analysis, background knowledge covers the different possible events, activities, and observations, whereas the examples are given by their respective co-occurrences. However, since user behavior is not strictly deterministic, these can be inconsistent which provides a challenge for logic-based methods.

As an alternative, behavioral rules can be derived by means of *decision-tree learning* [24]. This method can classify events in social media according to the activities they provoke in a group of users by ranking influential factors for distinguishing between these activities. The paths from a decision-tree's root to each leaf form a set of rules for the users' activity selections analogous to the inference rules provided by inductive logic programming. All users belonging to the same cluster can act as a sample for this classification method. Thus, decision-tree learning does not require a single prototypical user, but extracts activity patterns for a whole group of similarly behaving individuals. These patterns explain how communication takes place in social media which is a prerequisite for developing simulations to aid crisis communication and risk management.

4 Conclusions

This paper has pointed out systemic risks of social media with respect to crisis and reputation management. To handle these risks, it is necessary to analyze communication processes which result from individual behaviors of interconnected users. To that end, the paper has introduced behavior mining for analyzing such communication. It has provided an overview of the corresponding workflow and discussed methods for each process step according to their suitability for social media. Hence, it has laid the foundation for establishing behavior mining as an approach to systemic risk analysis in networked communication.

Nevertheless, the process of behavior mining is still work in progress which needs to be applied and evaluated with real-world data. While there are existing studies on crisis and risk management in social media as well as cluster analyses of network structures and communication flows [15, 5], applying the presented integrated workflow to these and other examples is subject to future work.

References

1. Alexander, D.E.: Social media in disaster risk reduction and crisis management. Sci. Eng. Ethics **20**(3), 717–733 (2014)
2. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. J. Econ. Perspectives **31**(2), 211–36 (2017)
3. Aula, P.: Social media, reputation risk and ambient publicity management. Strategy & Leadership **38**(6), 43–49 (2010)
4. Bao, J., Zheng, Y., Wilkie, D., Mokbel, M.: Recommendations in location-based social networks: a survey. GeoInformatica **19**(3), 525–565 (2015)
5. Berndt, J.O., Lorig, F., Timm, I.J., Barth, C., Bucher, H.J.: A systematic approach to agent-based dynamic analysis of social media communication. Intern. J. Advances in Internet Technology **10**(1&2) (2017)
6. Blei, D.M.: Probabilistic topic models. Commun. ACM **55**(4), 77–84 (2012)
7. De Domenico, M., Lima, A., Mougel, P., Musolesi, M.: The anatomy of a scientific rumor. Scientific reports **3**, 2980 (2013)
8. Dodds, P.S., Watts, D.J.: A generalized model of social and biological contagion. J. Theor. Biol. **232**(4), 587–604 (2005)
9. Ferrara, E., Varol, O., Davis, C., Menczer, F., Flammini, A.: The rise of social bots. Comm. ACM **59**(7), 96–104 (2016)
10. Freeman, L.C.: Centrality in social networks conceptual clarification. Soc. networks **1**(3), 215–239 (1978)

11. Garee, M.J., Chan, W.K.V., Wan, H.: Regression-based social influence networks and the linearity of aggregated belief. In: Proceedings of the 2018 Winter Simulation Conference. pp. 941–952. IEEE (2018)
12. Gupta, D.K., Muthiahz, S., Maresy, D., Ramakrish, N.: Forecasting civil strife: An emerging methodology. In: 3rd Intern. Conf. Human and Social Analytics (HUSO 2017). pp. 1–6 (2017)
13. Haas, S., Trump, T., Gerhards, M., Klingler, W.: Web 2.0: Nutzung und Nutzertypen. Mediaperspektiven **4**, 215–222 (2007)
14. Hodge, V., Austin, J.: A survey of outlier detection methodologies. Artif. Intell. Rev. **22**(2), 85–126 (2004)
15. Hoste, V., Van Hee, C., Poels, K.: Towards a framework for the automatic detection of crisis emotions on social media : a corpus analysis of the tweets posted after the crash of germanwings flight 9525. In: 2nd Intern. Conf. Human and Social Analytics (HUSO 2016). pp. 29–32 (2016)
16. Huang, H., Shen, H., Meng, Z., Chang, H., He, H.: Community-based influence maximization for viral marketing. Appl. Intell. (online first), 1–14 (2019)
17. Hutto, C.J., Gilbert, E.: Vader: A parsimonious rule-based model for sentiment analysis of social media text. In: 8th Intern. AAAI Conf. Weblogs and Social Media. pp. 216–255 (2014)
18. Jain, A.K.: Data clustering: 50 years beyond k-means. Pattern Recogn. Lett. **31**(8), 651 – 666 (2010)
19. Kunegis, J.: KONECT - The Koblenz Network Collection. In: Proc. Int. Web Observatory Workshop. pp. 1343–1350 (2013)
20. Levy, D., Nail, P.: Contagion: A theoretical and empirical review and reconceptualization. Genet. Soc. Gen. Psych. **119**(2), 253–283 (1993)
21. Lorig, F., Rodermund, S., Berndt, J., Timm, I.: Modeling and simulation of complex agents for analyzing communication behavior in social media. Intern. J. Advances in Internet Technology **11**(1&2), 11–20 (2018)
22. Muggleton, S., De Raedt, L.: Inductive logic programming: Theory and methods. J. Logic Program. **19**, 629–679 (1994)
23. Ott, L., Theunissen, P.: Reputations at risk: Engagement during social media crises. Public Relat. Rev. **41**(1), 97–102 (2015)
24. Safavian, S.R., Landgrebe, D.: A survey of decision tree classifier methodology. IEEE T. Syst. Man Cyb. **21**(3), 660–674 (1991)
25. Shirky, C.: The political power of social media: Technology, the public sphere, and political change. Foreign Aff. pp. 28–41 (2011)
26. Signorini, A., Segre, A.M., Polgreen, P.M.: The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. PloS one **6**(5), e19467 (2011)
27. Vega-Redondo, F.: Complex Social Networks. Cambridge University Press (2007)
28. Veil, S.R., Buehner, T., Palenchar, M.J.: A work-in-process literature review: Incorporating social media in risk and crisis communication. J. Conting. Crisis Man. **19**(2), 110–122 (2011)
29. Wang, X., Tang, L., Gao, H., Liu, H.: Discovering overlapping groups in social media. In: 10th Intern. Conf. Data Mining (ICDM). pp. 569–578. IEEE (2010)
30. Wendling, C., Radisch, J., Jacobzone, S.: The use of social media in risk and crisis communication. OECD Working Papers on Public Governance (24) (2013)
31. Yadav, A., Chan, H., Xin Jiang, A., Xu, H., Rice, E., Tambe, M.: Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In: AAMAS 2016. pp. 740–748 (2016)