Lecture Notes in Social Networks

Nitin Agarwal · Nima Dokoohaki
Serpil Tokdemir *Editors*

Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining

Springer

Lecture Notes in Social Networks

More information about this series at http://www.springer.com/series/8768

Nitin Agarwal • Nima Dokoohaki • Serpil Tokdemir
Editors

Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining

Springer

Editors
Nitin Agarwal
Collaboratorium for Social Media
and Online Behavioral Studies (COSMOS)
Information Science Department
University of Arkansas at Little Rock
Little Rock, Arkansas, USA

Nima Dokoohaki
Intellectera AB
Stockholm, Sweden

Serpil Tokdemir
Collaboratorium for Social Media
and Online Behavioral Studies (COSMOS)
University of Arkansas at Little Rock
Little Rock, Arkansas, USA

ISSN 2190-5428 ISSN 2190-5436 (electronic)
Lecture Notes in Social Networks
ISBN 978-3-319-94104-2 ISBN 978-3-319-94105-9 (eBook)
https://doi.org/10.1007/978-3-319-94105-9

Library of Congress Control Number: 2018952350

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Acknowledgments

This effort is funded in part by the U.S. National Science Foundation (IIS-1636933, ACI-1429160, and IIS-1110868), U.S. Office of Naval Research (N00014-10-1-0091, N00014-14-1-0489, N00014-15-P-1187, N00014-16-1-2016, N00014-16-1-2412, N00014-17-1-2605, and N00014-17-1-2675), U.S. Air Force Research Lab, U.S. Army Research Office (W911NF-16-1-0189), U.S. Defense Advanced Research Projects Agency (W31P4Q-17-C-0059), and the Jerry L. Maulden/Entergy Endowment at the University of Arkansas at Little Rock. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations. The researchers gratefully acknowledge the support.

Contents

About the Editors

Nitin Agarwal is a Distinguished Professor and Maulden-Entergy Endowed Chair of Information Science at the University of Arkansas at Little Rock. He is also the Director of the Collaboratorium for Social Media and Online Behavioral Studies (COSMOS). His research interests include social computing, deviant behavior modeling, mis/disinformation dissemination, computational propaganda analysis, group dynamics, social-cyber forensics, data mining, artificial intelligence, and privacy. His research has been supported by the U.S. National Science Foundation (NSF), the Army Research Office (ARO), the Office of Naval Research (ONR), the Air Force Research Laboratory (AFRL), the Defense Advanced Research Projects Agency (DARPA), and the Department of Homeland Security (DHS) with a total funding of over $10 million. He is a fellow of the prestigious International Academy, Research and Industry Association (IARIA). Dr. Agarwal received his doctorate at the Arizona State University in 2009 with outstanding dissertation recognition and was recognized as top 20 in their 20s by Arkansas Business. He has published over 100 peer-reviewed articles with several best paper awards and has been recognized as an expert in social, cultural, and behavioral modeling by several international news media organizations.

Nima Dokoohaki is a senior data scientist. He is currently affiliated with Intellectera, a data science research and development company where together with cofounders he develops and delivers solutions for consumer behavior modeling and analytics. In addition, he maintains collaboration with a research group at Software and Computer Systems department of Royal Institute of Technology (KTH) as an external advisor. His research interests include trust and privacy, applied machine learning, social computing, and recommendation systems. He received his Ph.D. in information and communications technology (ICT) in 2013. The main theme of his research was how to understand and leverage the notion of social trust so online service providers can deliver more transparent and privacy-preserving analytical services to their end users. His research has been backed by European projects funded from EU FP7 and Horizon 2020 framework programs, as well as distinguished public funding organizations including Swedish Research Council

and Vinnova. In 2014, he received a distinguished fellowship from the European Research Consortium for Informatics and Mathematics (ERCIM). He has published over 30 peer-reviewed articles. In addition to two best paper awards, he has been interviewed for his visible research, and his lecture has been broadcasted on Swedish public television. An ACM professional member, he is a certified reviewer for prestigious Knowledge and Information Systems (KAIS) as well as occasional reviewer for recognized international venues and journals.

Serpil Tokdemir is a research project analyst at the Office of Medicaid Inspector General (OMIG), Little Rock, Arkansas, USA. Dr. Tokdemir has a joint affiliation with the Collaboratorium for Social Media and Online Behavioral Studies (COSMOS) at UALR as research associate. Her work involves extracting raw data from Fraud and Abuse Detection System (FADS), cluster analysis, anomaly/outlier detection, predictive analysis and decision support systems, data visualization, content mining, and network analysis. Dr. Tokdemir obtained her Ph.D. from UALR in 2015 with support from U.S. National Science Foundation (NSF). Bringing together the computational modeling and social science theories, her dissertation explored the role of social media in coordinating online collective action in the context of Saudi Arabian Women's campaigns for right to gender equality. She has published several articles in this domain and won the most published student distinction by Engineering and Information Technology college at UALR. She obtained her Bachelor of Science in Computer Science from Marmara University, Istanbul, Turkey, in 2003. She completed her Master of Science (MS) in Computer Science from Georgia State University in 2006, Atlanta, Georgia, USA.

Part I
Emerging Social Issues

Intent Mining for the Good, Bad, and Ugly Use of Social Web: Concepts, Methods, and Challenges

Hemant Purohit and Rahul Pandey

Abstract The social web has empowered us to easily share information, express opinions, and engage in discussions on events around the world. While users of social media platforms often offer help and emotional support to others (the good), they also spam (the bad) and harass others as well as even manipulate others via fake news (the ugly). In order to both leverage the positive effects and mitigate the negative effects of using social media, intent mining provides a computational approach to proactively analyze social media data. This chapter introduces an intent taxonomy of social media usage with examples and describes methods and future challenges to mine the intentional uses of social media.

1 Introduction

The rapid adoption of social media has made the activity on online social networks (OSNs) an integral part of our daily lives. As per Pew Research Center survey,[1] nearly seven in every ten people in the USA use some type of OSNs (as of January 2018). The trend for the adoption of OSNs is not limited to the USA alone but worldwide, as evident from more than two billion monthly active users on Facebook across the world. The large scale of such digital connectivity comes with a medium to share information rapidly and interact with others virtually anywhere and anytime. Thus, OSNs facilitate an opportune playground for the users with varied intent (the purpose for an action), from helping others during disasters [36] to harassing and hate speech conversations [31] as well as manipulation with fake news

[1] http://www.pewinternet.org/fact-sheet/social-media/.

H. Purohit (✉) · R. Pandey
Department of Information Sciences and Technology, George Mason University, Fairfax, VA, USA
e-mail: hpurohit@gmu.edu; rpandey4@gmu.edu

N. Agarwal et al. (eds.), *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*, Lecture Notes in Social Networks, https://doi.org/10.1007/978-3-319-94105-9_1

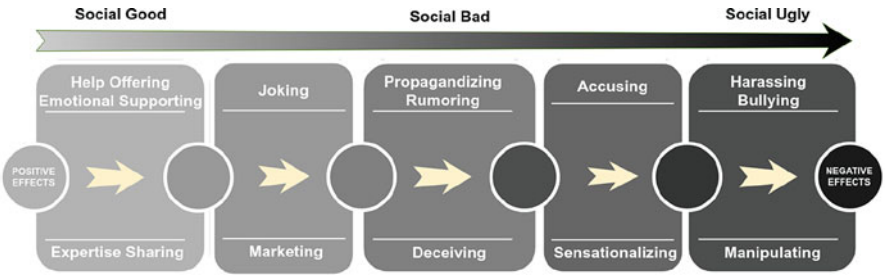

Fig. 1 A spectrum to demonstrate the variety of user intents existing on social media for the diverse uses from the social good to ugly

[45] and bots [16]. The good, bad, and ugly uses of OSNs have a profound impact on the evolution of our society. In fact, the gush[2] of fury and vitriol on the OSN companies in recent times for their inability to control the spread of disinformation [5, 28, 43] motivates the need to better understand the diverse uses of OSNs during real-world events.

We have seen many examples of the diverse usage of OSNs in the last decade, such as for coordination and self-organization during different types of social activism. For instance, #BlackLivesMatter [11] and #OccupyWallStreet [10] for justice and inequality, #Metoo [54] and #ILookLikeAnEngineer [26] for eradicating workplace harassment and stereotypes, etc. Likewise, OSNs have provided a valuable information exchange platform to support and rebuild communities after catastrophic natural disasters such as #HurricaneSandy [36] and #HaitiEarthquake [30] as well as enable community healing after man-made disasters such as mass shootings [19] and terror attacks [21]. Unfortunately, OSNs have also facilitated the amplification of malicious agenda such as to harass and bully others especially youth [8, 34], to spread disinformation for alternative narratives, as well as to manipulate public opinions via fake news during elections [2, 45].

One approach to understanding the nature and motives of information sharing on OSNs is to analyze the potential intent types (c.f. Fig. 1) associated with the OSN user interactions. Recognizing user intent helps collect evidences for interpreting and predicting potential actions and consequences—analogous to the problem of plan recognition in Artificial Intelligence [47]. We can model intent based on the content of the message shared, activity logs of the user sharing the message, and the link structure in OSNs that support the information flow of the message. Given the volume, variety, and velocity of information flowing on OSNs, computational approaches of intent mining provide a promising direction to help study the varied types of intent at large scale.

[2] https://journalistsresource.org/studies/society/internet/fake-news-conspiracy-theories-journalism-research.

Rest of the chapter provides an extensive overview of concepts, methods, and challenges in mining intent. In particular, Sect. 2 describes related concepts for a taxonomy of intent types, Sect. 3 provides an overview of different methods to process content, user, and network structure data for modeling intent, where the data may exist in different modalities—text, images, and videos. Finally, Sect. 4 describes the challenges in mining intent for future research directions to lead the society towards the good use of OSNs.

2 Concepts

This section describes first the concept of intent from multidisciplinary perspective, followed by describing the taxonomy of intent types for the diverse uses of OSNs, on a spectrum of positive to negative effects as shown in Fig. 1.

2.1 Intent: Multidisciplinary Perspective

Intent in the simplest form can be defined as a purpose for an action. In a more in-depth form, one can understand the broader view of intent from the concept of "intentional stance" proposed by the well-known philosopher and cognitive scientist Daniel Dennett [13]. Intentional stance is the highest abstract level of strategies for predicting and thereby, explaining and understanding the behavior of an entity (e.g., OSN user). Likewise, in Artificial Intelligence research community, the intent recognition problem has been studied for understanding the behavior of agents in the context of goal and plan recognition [47]. The power to recognize the plans and goals of other agents enables effective reasoning about the actions of the agents. In our context of the different uses of OSNs, a user can express desires and beliefs for certain intentionality in either message content or through his interactions and activities on OSNs. Therefore, a variety of factors can affect an individual's expression of intentionality through different information modalities. For example, "I wanna give #blood today to help the victims #sandy" shows the intent to donate blood for the desire to help and for the belief of resource scarcity to treat victims in the aftermath of Hurricane Sandy [36]. Intent can be expressed both explicitly and implicitly in a given content. Table 1 shows examples of messages with different intent types.

2.2 Intent Taxonomy

Given the diverse uses of OSNs and the endless possibilities of actions, the variety of intent behind the actions would be vast. Therefore, an approach to better understand

Table 1 Modified examples (for anonymity) of OSN messages from past events, with intent expressed in the textual content

Social media message	Intent [implication]
M1. *I want to send some clothes for hurricane relief #sandy*	Offering Help [community rebuilding and trust for collective action]
M2. *I support what u said about shooting here in Florida, Ill stand with u at any time. I am a retired teacher*	Emotional Supporting [community healing for psychological support]
M3. *You can find the @user student Developer Pack here: URL*	Expertise Sharing [improving learning from experiences of peers and mentors]
M4. *it's better to start 3rd-world war instead of letting Russia & assad commit #HolocaustAleppo*	Propagandizing [group-specific beliefs leading to echo chambers]
M5. *hi @user, we sincerely apologize for your inconvenience, in order to regain access to your account, please visit: URL*	Deceiving [financial frauds and stealth of personal information]
M6. *One of the suspects (according to BPD) is Sunil Tripathi. The missing Brown student NEWS reported on in March URL*	Rumoring [creating uncertainty in situational awareness for poor decision support]
M7. *you're a despicable whore*	Harassing [affecting mental health and physical well-being]
M8. *We are a people whose true lives begin after their death. #hijrah #jihad #shahadah*	Manipulating [shifting public attitude towards radicalized outfits]
M9. *DONT EVEN ASK EM WHO DEY WIT JUS BLOW EM FACES*	Bullying [threatening and creating fear and insecurity in the society]
M10. *there's a new drink called Sandy, it is a watered down Manhatten*	Joking [creating junk for some sections of the community]
M11. *No luck needed to #SAVE up to 60% off! Visit URL details of #vacation package*	Marketing [spamming in the information ecosystem]
M12. *white women have lied about rape against black men for generations*	Accusing [giving an alternative, supporting narrative to stereotypical groups]
M13. *There's no New Clinton, never has been. Shes same rape defending, racist, homophobic liar shes been for 70 yrs*	Sensationalizing [diverting from key issues and politicizing environment]

Intent can also be expressed using other information modalities (image, audio, or video); Fig. 2 shows an image example

the diverse uses and organize the associated intent for actions, we can consider a spectrum representation for the OSN uses with positive to negative effects. We also create a taxonomy of intent types as shown in Fig. 1. On the left side of the spectrum, social good uses of OSNs lead to the positive, enlightening effects of inspiration, cooperation, and trust in our society and strengthen the value of social networking in our lives. On the other hand, as we move towards the right end, social bad and ugly uses start to lead the negative effects of creating distrust, radicalization, and fear in our society. The social bad and ugly uses discredit and ruin the social networking values in our lives.

Fig. 2 Fake image shared during Hurricane Sandy 2012 (https://mashable.com/2012/10/29/fake-hurricane-sandy-photos/#Wc2mpf4QXgqV)

The proposed spectrum of Fig. 1 is flexible to extend the OSN uses as well as the intent taxonomy with different interpretations in the future. We broadly define five types of intent: (a) intent for the social good use, (b) intent for the social bad use, (c) intent for the social ugly use, (d) intent for the mixed social good and bad uses, and (e) intent for the mixed social bad and ugly uses.

The proposed intent types are described in the following with real examples of OSN messages in Table 1.

(a) *Intent for the Social Good Use*
People have good intentions and attitudes who believe in the social welfare and who would come forward to help others in the times of needs. OSNs facilitate a medium for such users to not only assist during disasters to provide emotional support and donations, but also, in general, offer help with the expertise to educate, inform, and caution-advice others. Illustrative intents in this category are:

- **Help Offering:** to express assistance to people in need of a resource or service. For example, message M1 in Table 1 shows a user offering clothing donations for disaster relief during Hurricane Sandy [36]. Likewise, users also offer to help with resources often, like blood donation [37].
- **Emotional Supporting:** to express care and sympathy for someone affected by an event. For example, message M2 in Table 1 shows support for the affected community of a mass shooting event. OSNs have played such roles in supporting a community for psychological well-being and caring of the affected people from depression and trauma [19].

– **Expertise Sharing:** to suggest or give advice to an information seeker based on expertise. For example, message M3 in Table 1 shows the answer to a user with a query to seek resources. OSNs provide hashtag and reply-based affordances for conversation chains, to allow expertise and knowledge sharing.

(b) *Intent for the Social Bad Use*
OSN users are not just humans but also social bots, who often participate for different motives in the conversations on social media. Both these types of users have contributed to create propaganda and spread the spamming content extensively in the recent years. Illustrative intents in this category are:

– **Propagandizing:** to create certain perception or belief towards an agenda of an organization or a group. For example, message M4 in Table 1 shows a strong justification for the government policies and attempts to convince the audience to believe in them [27].
– **Deceiving:** to spread spam or malicious content for a financial fraud or the purposeful misleading. For example, message M5 in Table 1 shows a clickbait and a potential scam for attracting readers to malicious sites related to buying some products and then stealing personal and financial information [25].
– **Rumoring:** to share unverified information aligned with emotions of someone that creates uncertainty. For example, message M6 in Table 1 shows a rumor indicating an emotionally charged message during Boston bombing and drawing everyone's attention to a misguided fact [46].

(c) *Intent for the Social Ugly Use*
Unfortunately, OSNs have become an avenue for conspiracy theories in recent years, where fake user accounts incite social tensions and radicalize others. Furthermore OSNs provide a medium to easily connect and converse with anyone that is abused (especially among youth) to engage in the online harassment and bullying, with the strong mental health implications. Illustrative intents in this category are:

– **Harassing:** to cause emotional distress to someone by insults, misogyny, or hateful messages and trolling for publicly shaming someone. For example, message M7 in Table 1 shows a sender harassing a receiver, which can lead to both mental and physical harm to the receiver [14].
– **Manipulating:** to purposefully divert a discourse to radicalize as well as politically or socially divide people. For example, message M8 in Table 1 shows how a potential member of a terror group can influence others and boost their recruitment drives [17].
– **Bullying:** to threaten or intimidate for creating a fear among a recipient. For example, message M9 in Table 1 shows a message of a gang member involved in illegal activities who threatens the rival gang, creating a fear in the social environment of the local region [3].

(d) *Intent for the Mixed Social Good and Bad Uses*
OSN users come from all sections of our society and their participation motives can range from personal to commercial usage. In this case, not all members of the society would benefit from all the activities of such users (e.g., a repetitive irrelevant advertisement) and therefore, the OSN use can be considered as mixed. Illustrative intents in this category are:

- **Joking:** to ridicule for fun or make a mockery of some event, object, or person. For example, message M10 in Table 1 shows a user making fun of Hurricane Sandy that may be amusing to some but contributes to the information overload on others, such as emergency services who would be working hard to monitor OSN streams for situational awareness [36].
- **Marketing:** to promote and advertise a product or service for selling. For example, message M11 in Table 1 shows a brand user creating a marketing pitch to attract more buyers that may be useful to some users who are looking to buy a travel package but a spam for those who are not traveling [12].

(e) *Intent for the Mixed Social Bad and Ugly Uses*
Users on OSN platforms may hold specific beliefs and may be associated with specific ideological identities such as political, religious, and social activist groups. Thus, their propaganda activities on OSNs can be motivated to meet the purpose of those belief and ideologies, however giving rise to echo chambers, which are the drivers of conspiracies. Illustrative intents in this category are:

- **Accusing:** to accuse someone and doubt publicly for creating an alternative reality. For example, message M12 in Table 1 shows a user trying to develop a narrative by accusing a female rape victim publicly and, thus, trying to undermine the key social issue of rape myths [4, 39].
- **Sensationalizing:** to provoke the audience to divert to an issue for frightening and politicizing the environment. For example, message M13 in Table 1 shows how a social issue can be mixed with a political context and divert the focus in a conversation away from the social issue (i.e., against rape) [40].

3 Methods

This section presents different types of methods to mine intent types described in the previous section.

Early research in online intent mining was focused on search engines, question-answering and product review forums, ad recommendation systems as well as spam detectors in information networks. For search systems, the key challenge was to understand information seeking intent of users in the queries on search engines using logs and give the relevant results to the users. Although, user query intent covers only a few categories of the broad variety of intents possible for uses of OSNs. In particular, query intent can be navigational, informational, or transactional

information to meet a user's information requirement [23], but the intent types in a social environment relate to communication and engagement with others in a conversation for different purposes, such as offering help or manipulating others. For question-answering and product forums, the possible intent types are centered around information seeking and knowledge sharing. For the ad recommendation systems, the commercial intent of buying and selling are priorities. For spam detection in networks, researchers focus on modeling patterns of malicious behavior but there are other types of intent possible for OSNs. Additionally, researchers have investigated intent across different modalities of information than the textual content, such as fake images [20] for rumors (see Fig. 2). Literature shows intent modeling in OSNs based on content of a message, user profile activities over time, and the network of user interactions as well as structural links of friendship and trust. We describe the methods under three major categories of content-based, user-based, and network-based approaches.

3.1 Content-Based Intent Mining

This type of methods solve the problem of inferring intent from a given instance of a message content shared on an OSN. Inferring intent from content is challenging due to possibilities of multiple natural language interpretations in a given text message. Therefore, to make the intent mining problem computationally tractable, prior research primarily exploited the text classification problem format [38]. Although it is different from the well-studied text analytics tasks of topic classification (focused on the subject matter) as well as opinionated text classification of sentiment or emotion (focused on the current state of affairs). For instance, in a message "people in #yeg feeling helpless about #yycflood and wanting to help, go donate blood," the task of topic classification focuses on the medical resource "blood," the task of sentiment and emotion classification is focused on the negative feeling expressed for being helpless. In contrast, intent classification concerns the author's intended future action, i.e., "wanting to help/donate." Therefore, the choice of feature representation is different across the tasks (e.g., adjectives are considered important for capturing sentiment and emotion, and likewise, verbs are important for indicating intent or action). Given the complexity to understand intent from natural language, researchers have explored various classifier designs using both rule-based systems and machine learning techniques.

Rule-based approaches are appropriate for small-scale data while for the large-scale data with intent labels, machine learning approaches can be leveraged. We summarize few approaches from the literature for brevity. Among rule-based classification approaches, Ramanand et al. [41] created rules for transactional (buying–selling) wishes in the product review text (e.g., "*<modal verb ><auxiliary verb >{window of size 3} <positive opinion word>*") and Purohit et al. [37] created rules for help-seeking and offering behavior during disasters (e.g., *" (Pronoun except you = yes)* $\wedge$ *(need/want = yes)* $\wedge$ *(Adjective = yes/no)* $\wedge$

(Thing = yes)" for seeking help about a "Thing" such as food). Among machine learning approaches, we can develop a classifier for detecting credible information messages to undermine potential rumor intent, such as Castillo et al. [7] proposed a classification method using the diverse features from message content, posting and re-tweeting behavior of users, and from citations to external sources. The key challenge of classification methods is to design good features that can efficiently capture the intent representation. Hollerit et al. [22] created a binary classifier for buying–selling posts on Twitter by exploring n-grams and POS tags-based features, and Carlos and Yalamanchi [6] proposed a supervised learning classifier for commercial intent based on features grounded in speech act theory. Purohit et al. [36, 38] proposed pattern-aided supervised classification approaches to identify the intent of help-seeking or offering during disasters, by combining the features from a bag-of-tokens model with patterns extracted from a variety of declarative and psycholinguistic knowledge sources. Likewise, Nazer et al. [33] proposed a system for identifying help-seeking request intent during disasters by combining content-based and context-based features such as the device type of a message source and location. While creating an exhaustive set of user-defined features from the user-generated content of social media can be challenging, researchers also explored deriving some valuable data-driven features for better generalization. Wang et al. [50] proposed a semi-supervised learning approach using the link prediction task in a graph of the tweet and intent-specific keyword nodes, in order to categorize intent tweets into different categories of general interests such as food and drink, travel, and goods and services. Given the possible lack of sufficient labeled data in an application domain, one can also use the transfer learning paradigm. Among such approaches, Chen et al. [9] built a combined classifier based on two classifiers trained on different source and target domains, in order to identify cross-domain intentional posts of commercial value (buying/selling) in discussion forums. Likewise, Ding et al. [15] proposed a convolutional neural network-based method for identifying user consumption intent for product recommendations, by transferring the mid-level sentence representation learned from one domain to another by adding an adaptation layer. Pedrood and Purohit [35] proposed sparse coding-based feature representation for efficient transfer learning to detect intent of help-seeking or offering in the future disaster event by exploiting data of historic disaster events.

3.2 User Profile-Based Intent Mining

This type of methods solve the problem of inferring the intent of a user by exploiting the patterns of activities or messages of the user in his historic profile data. It is similar to the idea of personalized recommender systems, which exploit all the historical data of a user to create his interest profile. The primary focus of these types of methods for OSNs is to model malicious user behavior such as spamming behavior to identify spammer networks or specific orientation towards

some beliefs. We explain a few approaches for brevity. The majority of such methods for learning and modeling the user behavior from historic data rely on machine learning techniques given the possibility to leverage large-scale historic data.

Intent mining literature has different methods from supervised to unsupervised learning for modeling user behavior, by leveraging features of all modalities such as text and images as well as temporal patterns of user activities. For instance, Jin et al. [25] created a detection system for users with malicious intent (spamming) by using both image and textual content features from the historic user profiles as well as the social network features. Lee et al. [29] created a supervised classifier to identify malicious content polluters using a diverse set of features from historic profile data including demographics, social network structure, the content of messages, as well as temporal behavior patterns in the activity. Among unsupervised learning approaches, Mukherjee et al. [32] proposed a method to exploit observed behavioral footprints of fake reviewers using a Bayesian framework. Furthermore, Ferrara et al. [16] review different methods for social bot detection using both feature-based and graph-based and crowdsourcing-based approaches.

Beyond the bot users, human users are also involved in the social bad and ugly uses of OSNs, such as with the intents of bullying and threatening others. Squicciarini et al. [44] proposed an approach to study both the detection of cyberbullies and the identification of the pairwise interactions between OSN users, who contributed in spreading the bullying intent. Salawu et al. [42] provide an extensive survey of the state of the art cyberbullying detection approaches. Likewise, Balasuriya et al. [3] studied the problem of detecting gang member profiles on Twitter that often share messages with the threatening intent, by proposing a method of supervised classification with diverse features of tweet text, profile information, usage pattern of emoji symbols, as well as additional information from the descriptions and comments on the external links of YouTube videos. On the other side of the OSN use spectrum, we can also model the user behavior in general for understanding the intent of non-malicious kind. For instance, Tomlinson et al. [48] proposed a method to detect a user's long-term intent and analyzing differences across cultures in expressing intent. Authors captured the latent cultural dimensions via the Singular Vector Decomposition technique. Such methods can be valuable for large-scale studies to assist multidisciplinary research at the intersection of social, humanities, and computing sciences.

3.3 Network-Based Intent Mining

Methods in this category focus on inferring the intent of a user by exploiting a given network structure of social relationships of the user in an OSN. The patterns of network structure can inform the membership to spam communities as well as information propagation cascades with the distinctive signatures of fake or rumor spreading intent. The network-based approaches have an advantage of being

language independent of the content, although they have to deal with a challenge of acquiring the network structure for any data modeling. Social network analysis methods are valuable for extracting the structural patterns. We summarize some of these approaches next.

The malicious users whether social bots or spammers or even radicalized group users often form community structures in the network, for sharing content and giving others a deceiving perception of general users. For instance, Ghosh et al. [18] investigated Twitter network for link farming—an approach to acquire a large number of follower links—by studying nearly 40,000 spammer accounts suspended by Twitter. Their analysis showed that the link farming is very common, where a majority of the links are acquired from a small fraction of Twitter users that are themselves seeking links. Likewise, a study conducted by Al-khateeb et al. [1] discovered cyber propaganda campaigns against NATO's Trident Juncture Exercise 2015 using social network analysis. There are also approaches for combining both features of network structure and content or user interaction patterns. Yu et al. [53] proposed a subgroup detection method to identify deceptive groups from their conversations, by combining linguistic signals in the content of interactions and signed network analysis for dynamic clustering. Among the approaches to model information propagation for identifying the intent of users, Starbird [45] studied the network generated from the common URL domains in the potentially malicious user messages on Twitter, which contained alternative narratives about mass shooting events and discovered the patterns of different domains and how they connect to each other. A model proposed by Wu and Liu [52] for the propagation of messages in OSNs infers embeddings of users with network structures as well as represents and classifies propagation pathways of a malicious intent message. Jiang et al. [24] provide an extensive survey of the approaches for malicious intent behavior detection across the categories of traditional spam, fake reviews, social spam, and link farming.

On the other side of the spectrum in Fig. 1 for positively using OSNs also, researchers have designed network-based approaches to glean intent of social good to help others. Welser et al. [51] identified key roles of Wikipedia editors such as substantive experts and vandal fighters by extracting patterns from edit histories as well as egocentric networks of users. Likewise, Tyshchuk et al. [49] presented a methodology that combined natural language processing and social network analysis to construct a network of actionable messages, for discovering communities and extracting leaders with a social good intent to help.

In summary, the approaches described above provide an overview of how one can study a variety of intent types in OSN uses by leveraging the message content, user profile history, and the social network structure.

4 Challenges and Future Research Directions

The use of OSNs in the future is going to be dependent on how OSN providers address the concerns of intent related to the social bad and ugly uses, which have given the perception that social networks are broken.[3] It is an open question—how we can create OSN platform affordances that would help manage both accountability of user activities and verification for trusted user networks, while discouraging the actors with social bad intents.[4] Similarly, it will be very important to boost the OSN uses with social good intents, such that we can still preserve some level of trust for OSN uses in the society. We describe some of these challenges in the following that future researchers can build on:

- **Profiling anonymous identities.** The cases for bullying and harassing intents often include harassers with anonymous profiles. The impact of such virtual anonymity leads to a lack of accountability and trust, due to the abuse of the medium of information sharing on OSNs. The anonymous users can spread information with malicious agenda but still remain unaccountable for the consequential effects. We need to address the challenge of understanding content and interaction patterns of such anonymous profiles for designing efficient user profiling methods.
- **Transforming social bots.** It is not clear how many users of OSNs are actually human users versus social bots, some of those present a threat to the information ecosystem of our society. While existing methods of bot detection provide some capability at scale to detect the bots, it is not clear beyond suspending them if we could alternatively transform the behavior of these bots. For example, teaching the intent behavior of social good as opposed to social bad (e.g., as observed in 2016, for the Microsoft chatbot[5]) could present an interesting opportunity to the human-in-the-loop Artificial Intelligence research.
- **Fixing erroneous spreading of malicious intent.** Sometimes the OSN users rapidly spread unverified, fake information due to emotional provocation such as after looking at an image of a disaster-affected site, although without a malicious goal. In this case, even if the user would like to change his course of action, the current OSN affordances only allow deletion of content for that individual user but the effect on the network is not handled effectively. Future research can also investigate this challenge of how to fix the issue of controlling message propagation.
- **Hybrid information filtering.** OSNs have been criticized lately to control what information a user can see, based on their content filtering and ranking algorithms. It leads to the formation of echo chambers with negative consequences.

[3] https://www.technologyreview.com/s/610152/social-networks-are-broken-this-man-wants-to-fix-them.

[4] https://datasociety.net/output/dead-reckoning/.

[5] http://www.bbc.com/news/technology-35890188.

There is a need for fairness and diversity in representation of information shown to a user such that the resulting content covers the varied intents of a story. It should further de-prioritize strongly subjective content and also provide an opportunity to the user to change the prioritization.

To conclude, this chapter presented a detailed overview of different uses of OSNs on the spectrum of social good to social ugly and also introduced an intent taxonomy. It further described intent mining methods and future challenges, which can help discover the varied types of intent behind the uses of OSNs.

Acknowledgements The authors thank Professor Amit Sheth at Kno.e.sis Center, Wright State University for valuable feedback and US National Science Foundation (NSF) for partially supporting this research on intent mining through grant award IIS-1657379. Opinions in this chapter are those of the authors and do not necessarily represent the official position or policies of the NSF.

References

1. Al-khateeb, S., Hussain, M.N., Agarwal, N.: Social cyber forensics approach to study twitter's and blogs' influence on propaganda campaigns. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, pp. 108–113. Springer, Berlin (2017)
2. Allcott, H., Gentzkow, M.: Social media and fake news in the 2016 election. J. Econ. Perspect. **31**(2), 211–36 (2017)
3. Balasuriya, L., Wijeratne, S., Doran, D., Sheth, A.: Finding street gang members on twitter. In: 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 685–692. IEEE, New York (2016)
4. Boux, H.J., Daum, C.W.: At the intersection of social media and rape culture: how Facebook postings, texting and other personal communications challenge the real rape myth in the criminal justice system. Univ. Illinois J. Law Technol. Policy **1**, 149 (2015)
5. Brown, J.: Is social media bad for you? the evidence and the unknowns (2018). http://www.bbc.com/future/story/20180104-is-social-media-bad-for-you-the-evidence-and-the-unknowns
6. Carlos, C.S., Yalamanchi, M.: Intention analysis for sales, marketing and customer service. In: Proceedings of COLING 2012: Demonstration Papers, pp. 33–40 (2012)
7. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on twitter. In: Proceedings of the 20th International Conference on World Wide Web, pp. 675–684. ACM, New York (2011)
8. Chatzakou, D., Kourtellis, N., Blackburn, J., De Cristofaro, E., Stringhini, G., Vakali, A.: Mean birds: detecting aggression and bullying on twitter. In: Proceedings of the 2017 ACM on Web Science Conference, pp. 13–22. ACM, New York (2017)
9. Chen, Z., Liu, B., Hsu, M., Castellanos, M., Ghosh, R.: Identifying intention posts in discussion forums. In: Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies, pp. 1041–1050 (2013)
10. Conover, M.D., Ferrara, E., Menczer, F., Flammini, A.: The digital evolution of occupy wall street. PLos One **8**(5), e64679 (2013)
11. De Choudhury, M., Jhaver, S., Sugar, B., Weber, I.: Social media participation in an activist movement for racial equality. In: ICWSM, pp. 92–101 (2016)
12. De Vries, L., Gensler, S., Leeflang, P.S.: Popularity of brand posts on brand fan pages: an investigation of the effects of social media marketing. J. Interact. Mark. **26**(2), 83–91 (2012)
13. Dennett, D.C.: The Intentional Stance. MIT Press, Cambridge (1989)

14. Dinakar, K., Jones, B., Havasi, C., Lieberman, H., Picard, R.: Common sense reasoning for detection, prevention, and mitigation of cyberbullying. ACM Trans. Interact. Intell. Syst. **2**(3), 18 (2012)
15. Ding, X., Liu, T., Duan, J., Nie, J.Y.: Mining user consumption intention from social media using domain adaptive convolutional neural network. In: AAAI vol. 15, 2389–2395 (2015)
16. Ferrara, E., Varol, O., Davis, C., Menczer, F., Flammini, A.: The rise of social bots. Commun. ACM **59**(7), 96–104 (2016). http://doi.acm.org/10.1145/2818717
17. Gates, S., Podder, S.: Social media, recruitment, allegiance and the Islamic state. Perspect. Terrorism **9**(4), 107–116 (2015)
18. Ghosh, S., Viswanath, B., Kooti, F., Sharma, N.K., Korlam, G., Benevenuto, F., Ganguly, N., Gummadi, K.P.: Understanding and combating link farming in the twitter social network. In: Proceedings of the 21st International Conference on World Wide Web, pp. 61–70. ACM, New York (2012)
19. Glasgow, K., Vitak, J., Tausczik, Y., Fink, C.: Grieving in the 21st century: Social media's role in facilitating supportive exchanges following community-level traumatic events. In: Proceedings of the 7th 2016 International Conference on Social Media & Society, p. 4. ACM, New York (2016)
20. Gupta, A., Lamba, H., Kumaraguru, P., Joshi, A.: Faking sandy: characterizing and identifying fake images on twitter during hurricane sandy. In: Proceedings of the 22nd International Conference on World Wide Web. pp. 729–736. ACM, New York (2013)
21. He, X., Lin, Y.R.: Measuring and monitoring collective attention during shocking events. EPJ Data Sci. **6**(1), 30 (2017)
22. Hollerit, B., Kröll, M., Strohmaier, M.: Towards linking buyers and sellers: detecting commercial intent on twitter. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 629–632. ACM, New York (2013)
23. Jansen, B.J., Booth, D.L., Spink, A.: Determining the informational, navigational, and transactional intent of web queries. Inf. Process. Manag. **44**(3), 1251–1266 (2008)
24. Jiang, M., Cui, P., Faloutsos, C.: Suspicious behavior detection: current trends and future directions. IEEE Intell. Syst. **31**(1), 31–39 (2016)
25. Jin, X., Lin, C., Luo, J., Han, J.: A data mining-based spam detection system for social media networks. Proc. VLDB Endowment **4**(12), 1458–1461 (2011)
26. Johri, A., Karbasian, H., Malik, A., Handa, R., Purohit, H.: How diverse users and activities trigger connective action via social media: lessons from the twitter hashtag campaign# ilooklikeanengineer. In: Proceedings of the 51st Hawaii International Conference on System Sciences (2018)
27. Kavanaugh, A.L., Fox, E.A., Sheetz, S.D., Yang, S., Li, L.T., Shoemaker, D.J., Natsev, A., Xie, L.: Social media use by government: from the routine to the critical. Gov. Inf. Q. **29**(4), 480–491 (2012)
28. Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S.A., Sunstein, C.R., Thorson, E.A., Watts, D.J., Zittrain, J.L.: The science of fake news. Science **359**(6380), 1094–1096 (2018). http://science.sciencemag.org/content/359/6380/1094
29. Lee, K., Eoff, B.D., Caverlee, J.: Seven months with the devils: A long-term study of content polluters on twitter. In: ICWSM (2011)
30. Meier, P.: Digital Humanitarians: How Big Data is Changing the Face of Humanitarian Response. CRC Press, Boca Raton (2015)
31. Mondal, M., Silva, L.A., Benevenuto, F.: A measurement study of hate speech in social media. In: Proceedings of the 28th ACM Conference on Hypertext and Social Media, pp. 85–94. ACM, New York (2017)
32. Mukherjee, A., Kumar, A., Liu, B., Wang, J., Hsu, M., Castellanos, M., Ghosh, R.: Spotting opinion spammers using behavioral footprints. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 632–640. ACM, New York (2013)

33. Nazer, T.H., Morstatter, F., Dani, H., Liu, H.: Finding requests in social media for disaster relief. In: 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 1410–1413. IEEE, New York (2016)
34. O'Keeffe, G.S., Clarke-Pearson, K., et al.: The impact of social media on children, adolescents, and families. Pediatrics **127**(4), 800–804 (2011)
35. Pedrood, B., Purohit, H.: Mining help intent on twitter during disasters via transfer learning with sparse coding. In: Proceedings of the 11th International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction and Behavior Representation in Modeling and Simulation (2018, in press). http://ist.gmu.edu/~hpurohit/informatics-lab/papers/sbp18-transferlearning-camera-ready-FINAL.pdf
36. Purohit, H., Castillo, C., Diaz, F., Sheth, A., Meier, P.: Emergency-relief coordination on social media: Automatically matching resource requests and offers. First Monday **19**(1) (2013). http://firstmonday.org/ojs/index.php/fm/issue/view/408
37. Purohit, H., Hampton, A., Bhatt, S., Shalin, V.L., Sheth, A.P., Flach, J.M.: Identifying seekers and suppliers in social media communities to support crisis coordination. Comput. Supported Coop. Work **23**(4–6), 513–545 (2014)
38. Purohit, H., Dong, G., Shalin, V., Thirunarayan, K., Sheth, A.: Intent classification of short-text on social media. In: IEEE International Conference on Smart City/SocialCom/SustainCom (SmartCity), 2015, pp. 222–228. IEEE, New York (2015)
39. Purohit, H., Banerjee, T., Hampton, A., Shalin, V.L., Bhandutia, N., Sheth, A.: Gender-based violence in 140 characters or fewer: a# bigdata case study of twitter. First Monday **21**(1) (2016). http://firstmonday.org/ojs/index.php/fm/issue/view/408
40. Purohit, H., Stabile, B., Grant, A., Pandey, R.: Modeling policy-relevant intent related to gender violence myths on social media using social construction theory. In: International Conference on Computational Social Science (IC2S2) (2018, forthcoming). http://ist.gmu.edu/~hpurohit/informatics-lab/papers/modeling-gbv-policy-intent-ic2s218.pdf
41. Ramanand, J., Bhavsar, K., Pedanekar, N.: Wishful thinking: finding suggestions and 'buy' wishes from product reviews. In: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, pp. 54–61. Association for Computational Linguistics, Los Angeles (2010)
42. Salawu, S., He, Y., Lumsden, J.: Approaches to automated detection of cyberbullying: a survey. IEEE Trans. Affect. Comput. (2017). https://doi.org/10.1109/TAFFC.2017.2761757, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8063898&isnumber=5520654
43. Solon, O.: Tim Berners-Lee: we must regulate tech firms to prevent 'weaponised' web (2018). https://www.theguardian.com/technology/2018/mar/11/tim-berners-lee-tech-companies-regulations?CMP=Share_iOSApp_Other
44. Squicciarini, A., Rajtmajer, S., Liu, Y., Griffin, C.: Identification and characterization of cyberbullying dynamics in an online social network. In: 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 280–285. IEEE, New York (2015)
45. Starbird, K.: Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on twitter. In: ICWSM, pp. 230–239 (2017)
46. Starbird, K., Spiro, E., Edwards, I., Zhou, K., Maddock, J., Narasimhan, S.: Could this be true?: I think so! expressed uncertainty in online rumoring. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 360–371. ACM, New York (2016)
47. Sukthankar, G., Geib, C., Bui, H.H., Pynadath, D., Goldman, R.P.: Plan, activity, and intent recognition: theory and practice. In: Newnes (2014)
48. Tomlinson, M., Bracewell, D., Krug, W.: Capturing cultural differences in expressions of intentions. In: Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, pp. 48–57 (2014)
49. Tyshchuk, Y., Li, H., Ji, H., Wallace, W.A.: Evolution of communities on twitter and the role of their leaders during emergencies. In: Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, pp. 727–733. ACM, New York (2013)

50. Wang, J., Cong, G., Zhao, W.X., Li, X.: Mining user intents in twitter: a semi-supervised approach to inferring intent categories for tweets. In: AAAI, pp. 318–324 (2015)
51. Welser, H.T., Cosley, D., Kossinets, G., Lin, A., Dokshin, F., Gay, G., Smith, M.: Finding social roles in Wikipedia. In: Proceedings of the 2011 iConference, pp. 122–129. ACM, New York (2011)
52. Wu, L., Liu, H.: Tracing fake-news footprints: Characterizing social media messages by how they propagate. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM'18, pp. 637–645. ACM, New York (2018). http://doi.acm.org/10.1145/3159652.3159677
53. Yu, D., Tyshchuk, Y., Ji, H., Wallace, W.: Detecting deceptive groups using conversations and network analysis. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, vol. 1, pp. 857–866 (2015)
54. Zarkov, D., Davis, K.: Ambiguities and dilemmas around #metoo: #forhow long and #whereto? Eur. J. Women's Stud. **25**(1), 3–9 (2018). https://doi.org/10.1177/1350506817749436

Bot-ivistm: Assessing Information Manipulation in Social Media Using Network Analytics

Matthew C. Benigni, Kenneth Joseph, and Kathleen M. Carley

Abstract Social influence bot networks are used to effect discussions in social media. While traditional social network methods have been used in assessing social media data, they are insufficient to identify and characterize social influence bots, the networks in which they reside and their behavior. However, these bots can be identified, their prevalence assessed, and their impact on groups assessed using high dimensional network analytics. This is illustrated using data from three different activist communities on Twitter—the "alt-right," ISIS sympathizers in the Syrian revolution, and activists of the Euromaidan movement. We observe a new kind of behavior that social influence bots engage in—repetitive @mentions of each other. This behavior is used to manipulate complex network metrics, artificially inflating the influence of particular users and specific agendas. We show that this bot behavior can affect network measures by as much as 60% for accounts that are promoted by these bots. This requires a new method to differentiate "promoted accounts" from actual influencers. We present this method. We also present a method to identify social influence bot "sub-communities." We show how an array of sub-communities across our datasets are used to promote different agendas, from more traditional foci (e.g., influence marketing) to more nefarious goals (e.g., promoting particular political ideologies).

M. C. Benigni · K. M. Carley (✉)
Institute for Software Research, Carnegie Mellon University, Pittsburgh, PA, USA
e-mail: kathleen.carley@cs.cmu.edu

K. Joseph
Computer Science and Engineering, SUNY Buffalo, Buffalo, NY, USA
e-mail: kjoseph@buffalo.edu

N. Agarwal et al. (eds.), *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*, Lecture Notes in Social Networks, https://doi.org/10.1007/978-3-319-94105-9_2

1 Introduction

How can individuals and groups be manipulated in social media? What messaging strategies can be used to shape behavior and alter opinions? In general, opinions and behaviors are a function of social influence [1]. That is, whom you know impacts not just what you know, but also your opinions and behavior. Consequently, group members, particularly those in tightly knit groups can come to share the same opinions and behaviors. Herein, we examine this process in social media. We identify a new type of bot that operates through social influence, the social influence bot network (SIBN). We then demonstrate the information strategies employed by such SIBNs to manipulate individuals and groups in social media.

Initially used to spread spam [2] and malware [3], a substantial literature now documents the use of bots on Twitter to influence global politics [4–7]. Many of the original bots discovered were individual bots acting in isolation to effect social goals. These bots or "social bots" [7, 8] have been used to shape discussions during political revolutions [9, 10] and in recruiting and propaganda efforts for terrorist groups [11, 12]. Social bots, and individual bots, were also pervasive and highly active in online conversations during the 2016 US presidential election [13].

More recently, we find evidence of concerted coordinated effort using networks of bots. Thus, we deviate from the existing literature on social bot networks in that we study the creation and use of a new form of social bot—the social influence bot network (SIBN). Most prior work has assessed the network structure of bots on the Twitter follower network [14] or on the directed @mention network [8]. In the latter case, the focus is generally on how bots mention real users in order to gain attention. SIBNs use @mentions in this way and in doing so change the effective influence of those users whom they @mention. Importantly, SIBNs also use @mentions to manipulate the Twitter social network by *@mentioning each other*. In other words, they operate by altering the social network structure, and so impact who has social influence on whom.

An example of the way social influence bots in our data use @mentions is shown in Fig. 1. The tweet in the figure was sent by a bot in our Syrian revolution dataset and contains only a string of mentions to nine other similarly named accounts. Shortly after, the bot sending this tweet was itself mentioned in similarly structured tweets by the other bots mentioned in Fig. 1. The sole purpose of these tweets is to artificially manipulate the reciprocal @mention, or co-mention graph, creating networks of bots with strong ties in the @mention network. More specifically, this kind of behavior produces a mention core of bots—a sub-community of social influence bots displaying "core-like" behavior [15] that have anomalously strong connections in the co-mention network. To distinguish this particular form of social bot network, we will refer to those social bot networks that have a mention core as **social influence bot networks (SIBNs)**.

Social influence bots can impact the Twitter ecosystem on at least three levels. First, they can be used at the "content-level" to rapidly spread specific tweets and/or particular URLs and make them appear artificially popular [8]. Second, they can be