Assessing online media content trustworthiness, relevance and influence: an introductory survey

Eleonora Ciceri, Roman Fedorov, Eric Umuhoza, Marco Brambilla, and Piero Fraternali

Politecnico di Milano. Dipartimento di Elettronica, Informazione e Bioingegneria
Piazza L. Da Vinci 32. I-20133 Milan, Italy
first.last@polimi.it

Abstract. The increasing popularity of social media articles and microblogging systems is changing the way online information is produced: users are both content publishers and content consumers. Since information is produced and shared by common users, who usually have a limited domain knowledge, and due to an exponential growth of the available information, assessing online content trustworthiness is vital. Several works in the state of the art approach this issue and propose different models to estimate online content trustworthiness, content relevance and user influence. In this paper we investigate the most relevant research works in this domain, highlighting their common characteristics and peculiarities in terms of content source type, trust-related content features, trust evaluation methods and performance assessment techniques.

1 Introduction

With the increasing popularity of social media, user-generated content [36] (i.e., content created by users and publicly available on the Web) is reaching an unprecedented mass. The overload of user-generated content makes hard to identify relevant content and to extract trustworthy and high quality information. Assessing trust [24, 31], content relevance [16] and user influence [11] is a critical issue in everyday social activities, where it is vital to filter non-authoritative, low quality and non-verified content to provide users with trusted information and content produced by experts.

As a motivating example, Motutu and Liu [45] report the "Restless Leg Syndrome" case: in 2008, when looking for information about the syndrome on Google, a wrong (and possibly dangerous) treatment promoted by the website WikiHow[1] was returned as top-1 result. This could obviously characterize a serious risk for patients; nevertheless, its rank wrongly suggested it could be trusted as verified and high quality information. Other applications that a misevaluation of user-generated content trustworthiness can affect are: disaster

This work is supported by POR-FESR 2007-2013 PROACTIVE Project and EU FP7-ICT-619172 SmartH2O Project.

[1] http://www.wikihow.com

management [41] (e.g., via false rumors on social networks during emergencies), environmental monitoring [18] (e.g., false reports of environmental phenomena), trend analysis [40] (e.g., via polluting the available information about what users like), detection of news [25] (e.g., via the diffusion of wrong news over the network).

Users often rely on their own knowledge, intuition and analytic capabilities to assess content relevance and trust. However, this becomes unfeasible with the current massive consumption of user-generated content: large volumes of low-quality, non-significant information are produced every day, and valuable content drowns in the large ocean of irrelevant content with little probability of being found by users. Consequently, it is vital to identify an automatic content trust estimation procedure which helps users in discarding unworthy information and focusing on significant content. Three ingredients are necessary to perform trust estimation: *i)* the evaluation of content relevance [15]; *ii)* the identification of influential users and experts [9], which are often focused on a specific topic, and produce mostly valuable content; *iii)* the evaluation of the level of *trust* [26] one can put on the content and people producing it. These ingredients are usually obtained by applying knowledge extraction algorithms and building appropriate trust models on user-generated content.

In recent years, several works in this field have emerged. In particular, several sub-fields significantly overlap between one another [60]: online content quality and relevance estimation [22], user reputation estimation [13] and influencers detection [42] all take part in assessing the quality of information one can find on the Web. This survey overviews the main state-of-the-art methods used in the automatic estimation of content quality, based on either content characteristics (i.e., content trustworthiness, relevance and credibility) or user characteristics (i.e., user trustworthiness and influence), which are strongly intertwined: high quality content often derives from highly experienced and influential users. Specifically, while other survey works go deeper in the details of trust estimation methods and applications [48, 34, 60], we deem our work merges together concepts from all the listed sub-fields and holds a practical relevance for practitioners and researchers who approach these themes for the first time.

The rest of this document is structured as follows: Section 2 introduces the concepts of trust, content relevance and user influence; Section 3 lists content and user profile features used as ingredients to assess content/user trustworthiness; Section 4 discusses methods to aggregate those features and provide a final trust score; Section 5 surveys the different approaches for performance assessment and output validation; finally, Section 6 concludes the work with final considerations and possible future directions in this field.

2 Trust, content relevance and user influence

In this section we introduce the definitions of trust, content relevance and influence, and list the research questions associated with these themes discussed in the state of the art.

2.1 Definitions

The concept of *trust* [39, 44] has been largely studied in the literature, both from a sociological [21] and philosophical [6] point of view. However, with the advent of social media [32, 17], studies on trust have recently shifted towards the construction of a trustworthiness model for digital content [45]. Siegrist and Cvetkovich [56] define trust as a tool that reduces social complexity: users that trust other users believe in their opinions, without making rational judgments. Sztompka [59] defines trust as "*the gambling of the belief of other people's possible future behavior*".

The concept of *relevance* (or *pertinence*) is crucial in the ability of an information retrieval system to find *relevant* content [43]. Many research works study the definition of relevance and its subjectivity in terms of system-oriented relevance [54], user relevance judgment [14], situation relevance [65], etc. Content relevance and popularity [10, 38, 19] are often connected: topic-related high quality content becomes often viral.

Social influence [61, 62] is defined as the power exerted by a minority of people, called *opinion leaders*, who act as intermediaries between the society and the mass media [33]. An opinion leader is a subject which is very informed about a topic, well-connected with other people in the society and well-respected.

The concepts of trust, content relevance and social influence are strongly intertwined: *i)* influential users (i.e., opinion leaders) are often experts in a specific field; *ii)* domain experts produce trustworthy content; *iii)* trustworthy, topic-related content has high relevance to the selected field. Moreover, popularity plays its role too: viral content is transmitted through the network in the same way a disease spreads among the population, and the more influential are the users sharing it, the larger is its popularity [47].

In this work we talk indistinctly about *trust*, *relevance* and *influence*, since they all represent quality measures for the object in question (i.e., either users or content). For the ease of the reader, henceforth, *trust* refers also to other discussed qualities, namely *relevance* and *influence*.

2.2 Research questions

A model of trust is defined as a function that extracts a set of *features* from a content object and aggregates them into a trustworthiness index. The construction of such model raises three research questions:

1. Which features better define the concept of trust and content quality?
2. How do we aggregate such features into a trustworthiness index?
3. How do we assess the quality of the trustworthiness index?

In the next sections these questions are addressed separately.

3 Trust model: features selection

In this section we describe features frequently used in the literature to assess the trustworthiness of Web content.

3.1 Source-based features

User-generated content is retrieved from a Web publishing source. Thus, the features one can extract from content to assess its quality depend on what can be extracted from the Web source. Although each source has its own characteristics and differences, we can classify them into two main categories:

– *Article-based sources* focus on the content itself, published in the form of articles. Content is usually long, and sometimes authors are encouraged to review, edit, rate and discuss it, thus creating high quality, multi-authored information. The author of the content may be thus unknown. Examples of these kind of sources are blogs, online encyclopedias (e.g., Wikipedia[2]) and question-answering communities (e.g., Stackoverflow[3]). Several works apply trust estimation techniques on these sources (e.g., [1, 3, 46]).
– *Social media* promote users as content authors: common people produce content which could become viral in short time. Users' authority becomes a key factor in the evaluation of content trustworthiness: non-expert authors often generate low quality, untrusted content. Examples of these kind of sources include Facebook[4], Twitter[5] and LinkedIn[6]. Several works apply trust estimation techniques on these sources: some examples can be found in [9, 41, 63].

Trust assessment studies performed on article-based sources tend to use content-based features (e.g., article length), since often the author is unknown, while works performed on social media focus both on author properties (e.g., number of connection with others) and content characteristics.

3.2 Content and author-based features

Moturu and Liu [45] propose a classification of features which takes inspiration from what people use to assess the trustworthiness of a person or a content in the real world. To evaluate user and content trustworthiness, we base our analysis on user's past actions (i.e., *reputation*), user/content present status (i.e., *performance*) and user/content perceived qualities (i.e., *appearance*). In the following, we describe each category separately. For a more complete overview see [7, 66, 49].

Reputation User reputation suggests how much one should trust their content [59]. The reputation depends on which actions users perform on social media, such as: *i)* content creation or consumption, *ii)* answers to others' content, *iii)* interactions with others, and *iv)* social networking. Reputation can be further split in the following feature categories:

[2] `http://en.wikipedia.org`
[3] `http://stackoverflow.com`
[4] `http://www.facebook.com`
[5] `http://twitter.com`
[6] `http://www.linkedin.com`

- *Connectedness.* The more a user is connected with others, the higher is his reputation in the network. Connectedness features are related to connections between users, and comprise simple features such as author registration status [45], number of followers/friends [51, 63, 4], number of accounts in different social media [29]. Furthermore, more complex features can be defined in this context, such as author centrality in graph of co-author network [50], social connectedness [45], number of reading lists the author is listed in [51], H-index [52] and IP-influence (i.e., influence vs. passivity) [52]. The identification of highly connected people is vital in case the objective is to spread content virally [55, 58].
- *Actions on the content.* The more acknowledged is the content one produces, the higher is his reputation on the network. Features in this category include the quantity/frequency of contributions to articles [45, 29], the amount of content sharing on social media [45, 3], the number of upvotes/likes [29], the number of answers to others' content [29], the number of retweets and retweeting rate [29, 52] and the Klout influence score [29].

Performance User performance describes the behavior of that user and his actions [45], and can be used to estimate his trustworthiness [59]. On the other hand, content performance can be determined from user's actions towards it and from the interest it generates. Performance-related features vary significantly depending on which social media platform we consider in our analysis. Example of such features include:

- Number of content edits [45].
- Direct actions on the content (e.g., number of responses/comments to a blog post [2] and retweets [29]).
- Characteristics of content update procedures (e.g., edit longevity [50], median time between edits, median edit length, proportion of reverted edits [45]).
- References to content by external sources (e.g., number of internal links [45, 2], incoming links [2], references by other posts [2], weighted reference score [45], publication date and place [29], variance on received ratings [29]).

Appearance External characteristics that represent the individual's appearance, personality, status and identity can be used to assess his trustworthiness. Similarly, the characteristics of content, such as style, size and structure, are useful in judging its quality. The most used features of this category include:

- Measure of the author reliability based on the structure of the content (e.g., length of blog posts, number of sections and paragraphs [45]).
- Language style (e.g., punctuation and typos [8], syntactic and semantic complexity and grammatical quality [3], frequency of terms belonging to a specific category [51], keywords in a tweet [8]).
- Originality of the content (e.g., presence of reused content [29], patterns of content replication over the network [8]).

4 Trust model: features aggregation

In Section 3 we present various feature categories used to assess the trustworthiness of online media content and users. Those features are transformed in a *trust/quality* index (usually scalar) through trustworthiness estimation algorithms.

Although it is common to find naive feature aggregation methods [66, 29, 50], the literature proposes a variety of more complex methods used to compute the final trust score. The categorization of such methods is not trivial, due to a fuzzy separation between *feature definition* and *feature aggregation* methods.

- *Statistical approaches.* It is common for features to be aggregated through cluster rank scores [45, 29, 12] or maximum feature values [2]. Several works use more refined approaches, such as cumulative distribution-based ranking, [66], K-nearest neighbors and Naive-Bayes classification [8], regression trees [4], mixture models [5], Gaussian Mixture Model and Gaussian ranking [49].
- *Graph-based algorithms.* Social connections play an important role in assessing the level of trust of an user and his content: the more connected is the user, the more others are interested in what he produces. Thus, several algorithms use graph-based methods, e.g., PageRank [52, 53, 27, 37, 64, 50, 63] and its variants [28], HITS [35], impact of a user on the social connections graph entropy [55], graph centrality measures [63, 27], indegree vs. outdegree [64] and other custom metrics based on information exchange over graphs [58, 8]. In some cases, trust is computed based on characteristics of a specific content source, e.g., number of followers vs. friends in the Twitter graph [28, 37].
- *Feature correlation.* Several works do not define an aggregation method, and simply study the correlation between features [52, 51].
- *Correlation between user influence and content relevance.* Some works use influencers retrieval techniques to identify influential users from a social network, and then navigate through the content they produce to collect the most relevant one [57].

Generally, the lack of uniformity in the proposed evaluation metrics and the heavy dependence on the type of content source (see Section 3.1) make it difficult to compare such metrics and state which one is most suited for a specific context. We believe that a further standardization of features would encourage the development of more sophisticated aggregation methods, e.g., based on supervised machine learning regressors and classifiers, as already proposed by Agichtein et al. [3] and by Castillo et al. [7].

5 Trust model: evaluation techniques

In this section we describe the experimental evaluation techniques that are used to assess the performance of the proposed trustworthiness estimation methods.

We state that the discussed research fields suffer from the absence of standardized requirements for the expected output. Thus, it is often difficult for the authors to compare their methods with respect to other state-of-the-art approaches.

5.1 Datasets

Due to the high variance of the type of data one can retrieve from each content source type, there exists a large collection of datasets in the state of the art, rarely made publicly available.

- *Custom datasets.* Almost all works create their own dataset by crawling data from the selected content publishing platforms. Several works (e.g., [7, 51, 52, 63]) base their analysis on Twitter, for several reasons: *i)* high volume of publicly available user-generated content; *ii)* presence of both textual and multimedia data; *iii)* access to public user profiles and their connections with other users; *iv)* easy storage of content (for further analysis), due to the limited length of posts. However, sometimes also article based platforms are taken into account (e.g., Wikipedia in Qin et al. [50], or question-answer platforms in Agichtein et al. [3]).
- *Use of standard datasets.* Sometimes, more standard datasets are used, e.g., the Enron Email Database[7] analyzed by Shetty and Adibi [55] or the WikiProject History[8] in [50], in which articles have been assigned class labels according to the Wikipedia Editorial Teams quality grading scheme.
- *Building a gold standard.* To assess the performance of a trust computation technique, it is often necessary to build a *gold standard* (or *ground truth*), i.e., a set of manually annotated data in which annotators are asked to state whether the content can be trusted, and labels are supposed to be error-free. In several contexts, labeling content is usually performed by a group of people (either part of an internal crowd or workers in some crowdsourcing platform [20, 67]), which manually annotate content. Then, the output of the proposed algorithm is compared with the ground truth, to assess the precision and recall of the retrieved set of trusted content/users [45, 66, 49]. However, trustworthiness, content quality and relevance are highly subjective characteristics, and thus the ground truth one builds is based on each annotator's perception of what being trustworthy means, which makes it biased and not reliable.

5.2 Performance assessment

State-of-the-art trust and influence metrics are all different and sometimes difficult to compare. Several works, thus, evaluate their performance with respect to similar algorithms applied to the same content sources. For this reason, the range of the metrics considered in this document is wide.

---

[7] https://www.cs.cmu.edu/~./enron/

[8] https://en.wikipedia.org/wiki/Wikipedia:WikiProject_History

- *Manual validation.* Many works tend to evaluate and discuss the results through manual inspection, where an internal crowd [55, 49, 9, 35, 27, 66] or anonymous users via user studies [23, 28, 12] assess the quality of the retrieved set of users/content.
- *Classification performance.* In some works, the authors manage to cast the trust evaluation problem as a classification problem, in which users are classified as influential/non-influential and content is labeled as trusted/non-trusted. These works are likely to present standard classification performance metrics: precision, TP-rate, FP-rate, accuracy [7] and ROC curves [3].
- *Evaluation of rankings.* In other cases, the output of the algorithm is a ranked list of authoritative content/users, and thus ranking correlation indexes (i.e., Pearson correlation [49] or generalized Kendall-Tau metrics [37]) are used to assess the performance of the proposed algorithm. In the same perspective, NDCG [30] (originally designed to test the ability of a document retrieval query to rank documents by relevance) is used to evaluate quality, trustworthiness and influence estimations, both in article-based content sources [45, 50] and microblogging platforms [66].
- *Comparison with known rankings.* Some works compare the output ranking of content/user with some rankings one can found on the Web, e.g., Digg [2], Google Trend and CNN Headlines [37].
- *Characteristics of users.* In some cases, one takes into account some characteristics of the involved users (e.g., activity [64] or validation of profile on Twitter [5]) to assess the performance of the algorithm. A high-performance result, in this sense, is the one maximizing the overlap between the set of active (validated) users and the users retrieved by the proposed algorithm.
- *Custom metrics.* Finally, some works build their own performance metrics, since in such cases it is difficult to compare the proposed algorithm with the ones available in the state of the art [51].

6 Conclusions and open challenges

In this survey we presented an overview of major recent works in the field of automatic estimation of trustworthiness, relevance and influence of online content. As discussed, trust estimation is important in Web search, and can be performed by capturing multiple signals deriving from both user profiles and content characteristics: authoritative (or influential) users produce mainly high quality content, and high quality content is largely trusted on the network of users. We thus reviewed several algorithms, listing their common characteristics and peculiarities in terms of content type, trust evaluation features and algorithms and performance assessment metrics.

We believe that these recent research topics are of great interest and practical importance in several domains such as automatic content retrieval and analysis, viral marketing, trend analysis, sales prediction and personal security. Nevertheless, in our opinion, there is enough space and need for future works that aim at building a concrete base of gold standards common to all discussed topics, and solidly integrating the proposed techniques to merge the efforts and converge towards a unified approach for user trust and content relevance estimation.

Current research works by the authors include methods for multi-platform and multimedia collective intelligence extraction from user-generated content, e.g., to perform trend analysis on the preference of Twitter users and to estimate environmental characteristics such as the presence of snow on mountains. Extracting relevant information from user-generated content implies: *i)* the identification of the influential users; *ii)* the estimation of content relevance; *iii)* the estimation of content trustworthiness. We believe that a strong cooperation of methods operating on multiple platforms and multiple content types (e.g., text, images, videos) is fundamental to define new standards this field lacks of.

References

1. Adler, B.T., Chatterjee, K., De Alfaro, L., Faella, M., Pye, I., Raman, V.: Assigning trust to wikipedia content. In: Proceedings of the 4th International Symposium on Wikis. p. 26. ACM (2008)
2. Agarwal, N., Liu, H., Tang, L., Yu, P.S.: Identifying the influential bloggers in a community. In: Proceedings of the 2008 international conference on web search and data mining. pp. 207–218. ACM (2008)
3. Agichtein, E., Castillo, C., Donato, D., Gionis, A., Mishne, G.: Finding high-quality content in social media. In: Proceedings of the 2008 International Conference on Web Search and Data Mining. pp. 183–194. ACM (2008)
4. Bakshy, E., Hofman, J.M., Mason, W.A., Watts, D.J.: Everyone's an influencer: quantifying influence on twitter. In: Proceedings of the fourth ACM international conference on Web search and data mining. pp. 65–74. ACM (2011)
5. Bi, B., Tian, Y., Sismanis, Y., Balmin, A., Cho, J.: Scalable topic-specific influence analysis on microblogs. In: Proceedings of the 7th ACM international conference on Web search and data mining. pp. 513–522. ACM (2014)
6. Blomqvist, K.: The many faces of trust. Scandinavian journal of management 13(3), 271–286 (1997)
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 (2011)
8. Cataldi, M., Aufaure, M.A.: The 10 million follower fallacy: audience size does not prove domain-influence on twitter. Knowledge and Information Systems pp. 1–22 (2014)
9. Cha, M., Haddadi, H., Benevenuto, F., Gummadi, P.K.: Measuring user influence in twitter: The million follower fallacy. ICWSM 10(10-17), 30 (2010)
10. Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.Y., Moon, S.: Analyzing the video popularity characteristics of large-scale user generated content systems. IEEE/ACM Transactions on Networking (TON) 17(5), 1357–1370 (2009)
11. Chan, K.K., Misra, S.: Characteristics of the opinion leader: A new dimension. Journal of advertising 19(3), 53–60 (1990)
12. Chen, C., Gao, D., Li, W., Hou, Y.: Inferring topic-dependent influence roles of twitter users. In: Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval. pp. 1203–1206. ACM (2014)

13. Cook, K.S., Yamagishi, T., Cheshire, C., Cooper, R., Matsuda, M., Mashima, R.: Trust building via risk taking: A cross-societal experiment. Social Psychology Quarterly 68(2), 121–142 (2005)
14. Cuadra, C.A., Katter, R.V.: Opening the black box of'relevance'. Journal of Documentation 23(4), 291–303 (1967)
15. De Choudhury, M., Counts, S., Czerwinski, M.: Find me the right content! diversity-based sampling of social media spaces for topic-centric search. In: ICWSM (2011)
16. De Choudhury, M., Counts, S., Czerwinski, M.: Identifying relevant social media content: leveraging information diversity and user cognition. In: Proceedings of the 22nd ACM conference on Hypertext and hypermedia. pp. 161–170. ACM (2011)
17. Ellison, N.B., et al.: Social network sites: Definition, history, and scholarship. Journal of Computer-Mediated Communication 13(1), 210–230 (2007)
18. Fedorov, R., Fraternali, P., Tagliasacchi, M.: Snow phenomena modeling through online public media. In: Image Processing (ICIP), 2014 IEEE International Conference on. pp. 2174–2176. IEEE (2014)
19. Figueiredo, F., Benevenuto, F., Almeida, J.M.: The tube over time: characterizing popularity growth of youtube videos. In: Proceedings of the fourth ACM international conference on Web search and data mining. pp. 745–754. ACM (2011)
20. Finin, T., Murnane, W., Karandikar, A., Keller, N., Martineau, J., Dredze, M.: Annotating named entities in twitter data with crowdsourcing. In: Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk. pp. 80–88 (2010)
21. Golbeck, J.: Combining provenance with trust in social networks for semantic web content filtering. In: Provenance and Annotation of Data, pp. 101–108 (2006)
22. Grady, C., Lease, M.: Crowdsourcing document relevance assessment with mechanical turk. In: Proceedings of the NAACL HLT 2010 workshop on creating speech and language data with Amazon's mechanical turk. pp. 172–179. Association for Computational Linguistics (2010)
23. Hannon, J., Bennett, M., Smyth, B.: Recommending twitter users to follow using content and collaborative filtering approaches. In: Proceedings of the fourth ACM conference on Recommender systems. pp. 199–206. ACM (2010)
24. Hsieh, H.F., Shannon, S.E.: Three approaches to qualitative content analysis. Qualitative health research 15(9), 1277–1288 (2005)
25. Hu, M., Liu, S., Wei, F., Wu, Y., Stasko, J., Ma, K.L.: Breaking news on twitter. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 2751–2754. ACM (2012)
26. Huang, F.: Building social trust: A human-capital approach. Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft pp. 552–573 (2007)
27. Huang, P.Y., Liu, H.Y., Chen, C.H., Cheng, P.J.: The impact of social diversity and dynamic influence propagation for identifying influencers in social networks. In: Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on. vol. 1, pp. 410–416. IEEE (2013)
28. Jabeur, L.B., Tamine, L., Boughanem, M.: Active microbloggers: identifying influencers, leaders and discussers in microblogging networks. In: String Processing and Information Retrieval. pp. 111–117. Springer (2012)
29. Jaho, E., Tzoannos, E., Papadopoulos, A., Sarris, N.: Alethiometer: a framework for assessing trustworthiness and content validity in social media. In: Proceedings

of the 23th International Conference on World Wide Web Companion. pp. 749–752. International World Wide Web Conferences Steering Committee (2014)
30. Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of ir techniques. ACM Transactions on Information Systems (TOIS) 20(4), 422–446 (2002)
31. Jøsang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision support systems 43(2), 618–644 (2007)
32. Kaplan, A.M., Haenlein, M.: Users of the world, unite! the challenges and opportunities of social media. Business horizons 53(1), 59–68 (2010)
33. Katz, E., Lazarsfeld, P.F.: Personal Influence, The part played by people in the flow of mass communications. Transaction Publishers (1955)
34. Kelton, K., Fleischmann, K.R., Wallace, W.A.: Trust in digital information. Journal of the American Society for Information Science and Technology 59(3), 363–374 (2008)
35. Kong, S., Feng, L.: A tweet-centric approach for topic-specific author ranking in micro-blog. In: Advanced Data Mining and Applications, pp. 138–151 (2011)
36. Krumm, J., Davies, N., Narayanaswami, C.: User-generated content. IEEE Pervasive Computing (4), 10–11 (2008)
37. Kwak, H., Lee, C., Park, H., Moon, S.: What is twitter, a social network or a news media? In: Proceedings of the 19th international conference on World wide web. pp. 591–600. ACM (2010)
38. Lerman, K., Hogg, T.: Using a model of social dynamics to predict popularity of news. In: Proceedings of the 19th international conference on World wide web. pp. 621–630. ACM (2010)
39. Maheswaran, M., Tang, H.C., Ghunaim, A.: Towards a gravity-based trust model for social networking systems. In: Distributed Computing Systems Workshops, 2007. ICDCSW'07. 27th International Conference on. pp. 24–24. IEEE (2007)
40. Mathioudakis, M., Koudas, N.: Twittermonitor: trend detection over the twitter stream. In: Proceedings of the 2010 ACM SIGMOD International Conference on Management of data. pp. 1155–1158. ACM (2010)
41. Mendoza, M., Poblete, B., Castillo, C.: Twitter under crisis: Can we trust what we rt? In: Proceedings of the first workshop on social media analytics. pp. 71–79. ACM (2010)
42. Merton, R.K.: Patterns of influence: Local and cosmopolitan influentials. Social theory and social structure 2, 387–420 (1957)
43. Mizzaro, S.: How many relevances in information retrieval? Interacting with computers 10(3), 303–320 (1998)
44. Molm, L.D., Takahashi, N., Peterson, G.: Risk and trust in social exchange: An experimental test of a classical proposition. American Journal of Sociology pp. 1396–1427 (2000)
45. Moturu, S.T., Liu, H.: Quantifying the trustworthiness of social media content. Distributed and Parallel Databases 29(3), 239–260 (2011)
46. Nam, K.K., Ackerman, M.S., Adamic, L.A.: Questions in, knowledge in?: a study of naver's question answering community. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 779–788. ACM (2009)
47. Ni, M., Chan, B., Leung, G., Lau, E., Pang, H.: Data from: Transmissibility of the ice bucket challenge among globally influential celebrities: retrospective cohort study (2014), `http://dx.doi.org/10.5061/dryad.n4sc4`
48. Nurse, J.R., Rahman, S.S., Creese, S., Goldsmith, M., Lamberts, K.: Information quality and trustworthiness: A topical state-of-the-art review. In: Proceedings of the International Conference on Computer Applications and Network Security (ICCANS) (2011)

49. Pal, A., Counts, S.: Identifying topical authorities in microblogs. In: Proceedings of the fourth ACM international conference on Web search and data mining. pp. 45–54. ACM (2011)
50. Qin, X., Cunningham, P.: Assessing the quality of wikipedia pages using edit longevity and contributor centrality. arXiv preprint arXiv:1206.2517 (2012)
51. Quercia, D., Ellis, J., Capra, L., Crowcroft, J.: In the mood for being influential on twitter. In: Privacy, Security, Risk and Trust (PASSAT), 2011 IEEE Third International Conference on. pp. 307–314. IEEE (2011)
52. Romero, D.M., Galuba, W., Asur, S., Huberman, B.A.: Influence and passivity in social media. In: Machine learning and knowledge discovery in databases, pp. 18–33. Springer (2011)
53. Saez-Trumper, D., Comarela, G., Almeida, V., Baeza-Yates, R., Benevenuto, F.: Finding trendsetters in information networks. In: Proceedings of the 18th ACM SIGKDD. pp. 1014–1022. ACM (2012)
54. Schamber, L., Eisenberg, M.B., Nilan, M.S.: A re-examination of relevance: toward a dynamic, situational definition. Information processing & management 26(6), 755–776 (1990)
55. Shetty, J., Adibi, J.: Discovering important nodes through graph entropy the case of enron email database. In: Proceedings of the 3rd international workshop on Link discovery. pp. 74–81. ACM (2005)
56. Siegrist, M., Cvetkovich, G.: Perception of hazards: The role of social trust and knowledge. Risk analysis 20(5), 713–720 (2000)
57. Silva, A., Guimarães, S., Meira Jr, W., Zaki, M.: Profilerank: finding relevant content and influential users based on information diffusion. In: Proceedings of the 7th Workshop on Social Network Mining and Analysis. p. 2. ACM (2013)
58. Sun, B., Ng, V.T.: Identifying influential users by their postings in social networks. Springer (2013)
59. Sztompka, P.: Trust: A sociological theory. Cambridge University Press (1999)
60. Thirunarayan, K., Anantharam, P., Henson, C., Sheth, A.: Comparative trust management with applications: Bayesian approaches emphasis. Future Generation Computer Systems 31, 182–199 (2014)
61. Watts, D.J., Dodds, P.S.: Influentials, networks, and public opinion formation. Journal of consumer research 34(4), 441–458 (2007)
62. Weimann, G.: The influentials: People who influence people. SUNY Press (1994)
63. Weitzel, L., Quaresma, P., de Oliveira, J.P.M.: Measuring node importance on twitter microblogging. In: Proceedings of the 2nd International Conference on Web Intelligence, Mining and Semantics. p. 11. ACM (2012)
64. Weng, J., Lim, E.P., Jiang, J., He, Q.: Twitterrank: finding topic-sensitive influential twitterers. In: Proceedings of the third ACM international conference on Web search and data mining. pp. 261–270. ACM (2010)
65. Wilson, P.: Situational relevance. Information storage and retrieval 9(8), 457–471 (1973)
66. Zhai, Y., Li, X., Chen, J., Fan, X., Cheung, W.K.: A novel topical authority-based microblog ranking. In: Web Technologies and Applications, pp. 105–116. Springer (2014)
67. Zubiaga, A., Liakata, M., Procter, R., Bontcheva, K., Tolmie, P.: Crowdsourcing the annotation of rumourous conversations in social media. In: Proceedings of the 24th International Conference on World Wide Web Companion. pp. 347–353 (2015)