

# Visual Exploration and Debugging of Machine Learning Classification over Social Media Data



Mayank Kejriwal and Peilin Zhou

**Abstract** Humanitarian and geopolitical crises (such as COVID-19) are frequently extra-national in scope. Technology, including applications of natural language processing and machine learning, can play a vital role in mitigating this burden, especially with availability and real-time analyses of social media. One such application is *situation labeling*, intuitively defined as the semi-automatic assignment of one or more *actionable* labels (such as food, medicine or water) from a controlled vocabulary to tweets or documents that become available in the aftermath of a crisis, such as an earthquake. Despite multiple advances, users of current situation labeling systems are often unwilling to trust these (and other machine learning) outputs without some provenance and visualization of results. This article describes an interactive visualization approach called SAVIZ that allows non-technical users to intuitively and interactively explore outputs of situation labeling systems. We illustrate the potential of SAVIZ with two real-world crisis datasets from Twitter. Our platform is completely built using open-source tools, can be rendered on a web browser and is backward-compatible with several pre-existing crisis intelligence platforms.

**Keywords** Crisis informatics · Situation labeling · Visualization · Machine learning · Text classification · Situational awareness · Natural language processing · Embeddings · Representation learning

---

M. Kejriwal (✉) · P. Zhou  
University of Southern California, Los Angeles, CA, USA  
e-mail: [kejriwal@isi.edu](mailto:kejriwal@isi.edu)

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2021  
M. Çakırtaş, M. K. Ozdemir (eds.), *Big Data and Social Media Analytics*, Lecture Notes in Social Networks, [https://doi.org/10.1007/978-3-030-67044-3\\_8](https://doi.org/10.1007/978-3-030-67044-3_8)

## 1 Introduction

In its global humanitarian overview in 2020, the United Nations Office for the Coordination of Human Affairs (OCHA) reported<sup>1</sup> forecast that in 2020, more than 167 million people will need humanitarian assistance. In 2019, many more needed humanitarian assistance than was originally forecast, primarily because of the effects of conflicts and ‘extreme climate’ events. According to the overview, not only is global climate change increasing people’s vulnerability to humanitarian disasters, with the top eight worst food crises all linked to conflict and climate change effects, but infectious diseases are becoming harder to control and are becoming more widespread. This is in addition to persistent problems like malnutrition. At the start of 2019, some 821 million people were under-nourished (according to the same overview), with more than a 113 million suffering from acute hunger.<sup>2</sup>

While donors do step up to the challenge, the massive scale of these problems typically leads to a funding shortfall, and less effectiveness in tackling these issues. One solution that has been touted is *crisis informatics*, which has emerged as an important interdisciplinary area [32], with contributions from both social and computational sciences, including machine learning, information retrieval, natural language processing, social networks and visualization [19, 28, 29, 34, 38]. The key idea behind crisis informatics is to use technological solutions, particularly those powered by information science or even ‘Big Data’, to help predict disasters and mobilize resources more effectively [9, 26].

To realize this vision in more specific ways, multiple government and private programs have been instituted, some with direct, and others with indirect, salience to crisis informatics. An example of the latter is the DARPA LORELEI program<sup>3</sup> that was established with the explicit agenda of providing *situational awareness for emergent incidents*, under the assumption that the emergent incident occurs in a region of the world where the predominant language is computationally *low-resource* [16]. Emergent incidents do not have to be limited to natural disasters, though that they were considered a critical use case. An example of a computationally low-resource language is Uighyur, a Turkic language spoken by about 10–25 million people in Western China, for which few automated technology capabilities currently exist [2]. LORELEI situational awareness systems like THOR [20] must first translate tweets and messages into English, using automated machine translation algorithms, and to provide further analytical capabilities, must execute additional Natural Language Processing (NLP) and Artificial Intelligence (AI)

---

<sup>1</sup>[https://www.unocha.org/sites/unocha/files/GHO-2020\\_v9.1.pdf](https://www.unocha.org/sites/unocha/files/GHO-2020_v9.1.pdf)

<sup>2</sup>This chapter is an extended version of [23], which was a 4-page demonstration submission at the 2019 IEEE/ACM ASONAM conference in Vancouver, Canada. In this article, we describe the system in depth, including a description of its workflow, full description of its use-cases on real data, links to the system (which has never been released prior) and an updated version that includes time as a control variable.

<sup>3</sup><http://www.darpa.mil/program/low-resource-languages-for-emergent-incidents>

algorithms like named entity recognition, automatic detection of *need types* (e.g., does the tweet express a *food* need or a *medical* need?) and sentiment analysis.

Despite advances in NLP and AI, such algorithms continue to be imperfect. For example, we executed a state-of-the-art crisis informatics NLP system called ELISA [4] on an Ebola<sup>4</sup> dataset collected over Twitter. Among other things, ELISA ingests a tweet as input and uses a pre-trained machine learning module developed over the course of the ongoing DARPA LORELEI program to output categorical *situation labels* such as food, medicine, water and infrastructure that allow humanitarian responders to quickly decide where to focus their attention and resources (as opposed to reading every single tweet in the corpus). While for some (pre-processed) tweets such as ‘ebola in sierra leone’, ELISA correctly outputs the label ‘med’, it also erroneously outputs labels like ‘med’ for other tweets like ‘vivian dou yemm moh’, which have become meaningless due to mangled machine translation or heavy dependence on emoticons and symbols (that get removed during pre-processing). While the modules get better over time, performance is still well below 70% F-measure due to data sparsity and noise. Performance is even worse when the modules are trained on one type of disaster or locale, but have to be applied in another. This is a pervasive problem in crisis informatics, since every crisis is different, making generalization difficult. It also makes humanitarian responders question the veracity of such a system, making transition and uptake of advanced AI technology a social challenge.

In this chapter, we present a highly lightweight, interactive visualization platform called SAVIZ that can be deployed on a web browser in less than 30 seconds for thousands of tweets, and is designed for short, crisis-specific messages collected over social media like Twitter, and processed by systems like ELISA. SAVIZ relies on established, pre-existing and open-source technologies from the representation learning, visualization and data processing communities. SAVIZ is backward-compatible with crisis informatics sub-systems recently released under the LORELEI program, and has been applied on real-world datasets collected from the Twitter API.

## 2 Related Work

Visualization is an important part of any human-centric system that is attempting to make sense of a large amount of information. Several good crisis informatics platforms that provide visualizations include Twitris [17], CrisisTracker [35], Twitcident [1], TweetTracker [24], AIDR [14], and several others. Some research efforts focus on improving accuracy on a narrow but difficult and important problem, such as extracting information from micro-blogs or determining if a particular

---

<sup>4</sup>Ebola is a rare and deadly disease that currently has no approved vaccine or cure: <https://www.cdc.gov/vhf/ebola/index.html>

message is relevant to the disaster in question (from a much bigger stream of messages). Specific examples include work on extracting parcels of information from disaster-related social media messages [15], work on semi-automatic detection of informative tweets during emerging disasters [44], the Twitter-specific case study by Thom et al. [40], among several others [5, 12, 36, 37]. In the last few years, novel advances in AI, including deep learning, have also been explored for crisis response applications. The preprint by Nguyen et al. [30] is a good reference. Finally, work such as [21] address the critical issue of how to acquire Twitter data efficiently in the aftermath of a crisis, especially without paid subscriptions or unlimited API calls.

Another line of work is algorithmic, rather than applied, but is relevant because improvements in some of these algorithms has a direct effect on the functioning and performance of downstream interactive tools and applications. For example, work on discovering geographical topics in the Twitter stream [13] has a direct effect on information extraction and relevance detection. Algorithmic innovations in AI areas such as entity linking [6, 27], event detection [3, 11], crowdsourcing [10, 41], representation learning [18, 25], and sentiment analysis [8] also have consequential effects.

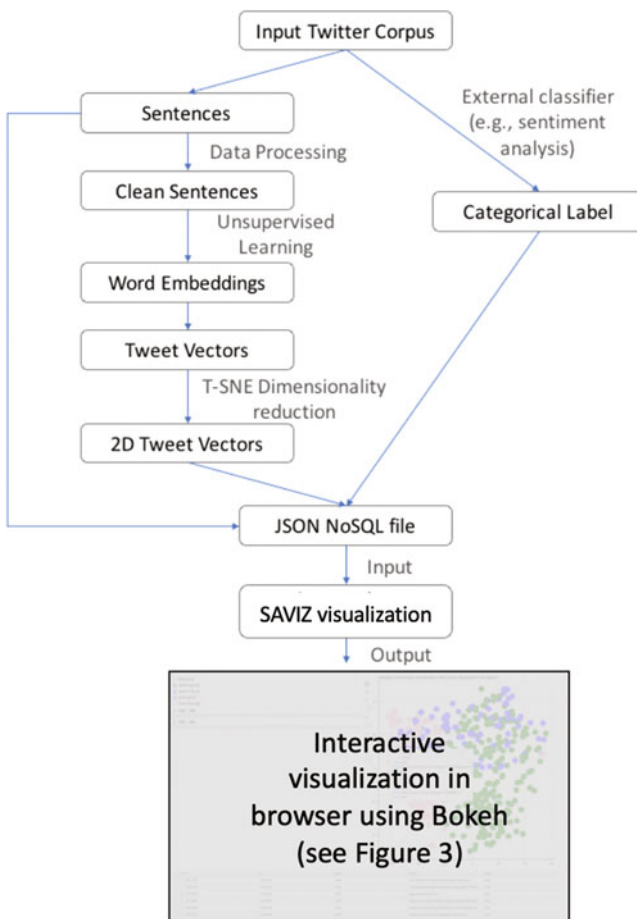
A more sophisticated interactive system, THOR (Text-enabled Humanitarian Operations in Real-time) [20], also provides sophisticated situational awareness, but is designed for computationally low-resource languages like Uighyur and Bengali, and consequently has a stronger focus on NLP tasks like machine translation.

Several aspects of SAVIZ distinguish it from the systems referenced above. The most important difference is that, unlike the systems above, SAVIZ ingests not just the raw social media data stream itself, but the categorical outputs of NLP and machine learning systems like situation labeling and sentiment analysis [33]. Thus, SAVIZ allows the user to jointly explore both the social media data and the labels, which serves two purposes: to understand the noise in the classification system, and to understand the social media stream in aggregate. For example, consider again Fig. 3, which expresses the (initially non-intuitive) finding that ‘water’ (green) is as big an issue in the context of the collected data, as ‘med’ (pink), which is what one would expect in a dataset collected from the Twitter API specifically using Ebola-related keywords. Other differences between SAVIZ and systems like CrisisTracker [35] is its use of embeddings and 2D projection (using t-SNE [25]) as an interactive visual aid. As more advanced embeddings (including network and knowledge graph embeddings [39, 42, 43]) continue to be developed and released as open-source, SAVIZ will be well-positioned to use these to provide an alternate ‘view’ of the data. The current version of SAVIZ is already capable of treating embeddings as a black-box, by directly ingesting the high-dimensional vectors as its input. This allows the system to be lightweight, simple and customizable.

### 3 SAVIZ: Brief Overview

SAVIZ has a simple processing workflow that is illustrated schematically in Fig. 1. As a first step, the system ingests an input Twitter corpus that has been collected in the aftermath of a crisis. A good system that is capable of such focused data collection is CrisisLex [31]; also, the recent pipeline approach by Gu and Kejriwal is an alternative way of collecting relevant tweets using methods like active learning [21].

Once the corpus has been acquired, an NLP-based situational awareness system like ELISA [4] is typically executed over it. In addition to machine translation and named entity recognition, ELISA also does *situation labeling* on each message.



**Fig. 1** The SAVIZ workflow over a corpus of Twitter data collected in the aftermath of a crisis. A detailed version of the interactive visualization at the bottom is demonstrated in Fig. 3

Situation labeling is a multi-label problem wherein one or more *situation types*, such as food, medicine, water, evacuation etc. from a small set of pre-defined types, are assigned to each message. However, generic black-box algorithms from the sentiment analysis and classification literature could also be used in this phase. The result of such analysis is one or more *categorical* labels per tweet. For example, for a given message ‘Tragic loss of life when dam collapsed. Many more trapped, awaiting rescue’, the situation labeling may output ‘evacuation’ and ‘infrastructure’ (depending on how the system was trained, and how broad its outputs should be), while sentiment analysis may output ‘negative’. Other kinds of classification have also been explored in the literature including *urgency detection* [22]. For example, the message above is arguably high on an urgency scale, since many people could lose lives if the response is not expedient enough. Other messages that may be discussing a temporary power outage or a minor flood warning may not be deemed as urgent. By using these systems, first responders can plan efficient resource and personnel deployment.


The next few steps are unsupervised. The tweets are first preprocessed by converting them to lower-case and removing special symbols and characters (like #, @ etc.), along with tabs and newlines. For example, the tweet ‘massive earthquake in NEPAL ————— Bhimsen Tower aka Dharahara In Nepal... <http://t.co/4tUDQDWvC4>’ would, after preprocessing, become ‘massive earthquake in nepal bhimsen tower aka dharahara in nepal <http://t.co/4tUDQDWvC4>’ (Last accessed January 2021).

Next, we use the ‘bag-of-tricks’ *word embedding* package (called fastText [18]) released by Facebook Research to embed the words and sentences into a dense, continuous and low-dimensional (specifically, 100-dimensional) vector space. Word embedding algorithms, which are a more specific sub-field of the more general area of *representation learning*, have been extremely influential in NLP for the last ten years, leading to improvements in performance across multiple NLP problems without necessitating domain-specific feature engineering [7]. More recently, the impact of word embeddings and other kinds of embeddings have percolated into multiple communities that rely on text and image analytics. Social media analytics and crisis informatics are good examples, including systems such as SAVIZ. In previous work, for instance, we used these word embeddings to derive vectors for hashtags, and found that simply exploring hashtags in a 2D space (described further below) can yield important insights about disasters (Fig. 2), such as the tragic mass shooting in Las Vegas towards the end of 2017.

To enable visualization, we use t-SNE<sup>5</sup> [25] to project the sentence vectors into a 2D space. The t-SNE algorithm has achieved tremendous impact as a standard embedding-visualization tool in the machine learning community (and also beyond). In principle, other dimensionality reduction tools existed before t-SNE was proposed, but by using a neural network for its optimization, t-SNE is able to achieve experimentally superior visualization wherein points that are close in the

---

<sup>5</sup>Stands for *t-Distributed Stochastic Neighbor Embedding*.



**Fig. 2** Embedding of hashtags and visualization in 2D space using the t-SNE algorithm. The dimensions do not have any intrinsic meaning and are only useful as a coordinate system for visualization purposes

original vector space tend to be close in the reduced space also. This means that the visualization more accurately reflects what is happening in the higher dimensional space. In the context of this chapter, using t-SNE allows SAVIZ to present more compact and accurate visualizations (subsequently described) to a user.

SAVIZ uses all the information sets mentioned above, including the categorical labels output by systems like ELISA and the 2D points output by t-SNE, to compile a NoSQL file that is input to the SAVIZ visualization module. This module is based on Bokeh,<sup>6</sup> an interactive, well-documented visualization library that targets modern web browsers for presentation. Bokeh aims to provide elegant, concise construction of versatile graphics, and to extend this capability with high-performance interactivity over very large or streaming datasets. Because it uses Bokeh, SAVIZ requires no extensive set up, since the visualization itself is rendered on a web browser, making the system portable.

<sup>6</sup>[https://bokeh.pydata.org/en/latest/docs/dev\\_guide.html](https://bokeh.pydata.org/en/latest/docs/dev_guide.html)

### 3.1 User Experience

We demonstrate the simplicity of using, and the key features of, the system for a complex disaster (the Ebola crisis) that continues to unfold in Africa. Figure 3 illustrates the SAVIZ interface for an Ebola dataset that was collected from Twitter. The full corpus that we collected, using Ebola-specific keywords, comprises 18,224 tweets, with timestamps ranging from 2014-Aug-01 00:03 to 2014-Sep-24 23:16. ELISA [4] was executed on this corpus, yielding zero or more situation labels per tweet from a vocabulary of eleven types: *food, infrastructure, water, utilities, regime-change, terrorism, medicine, evacuation, shelter, search, and crime/violence*. These labels could be noisy; no ground truth was available against which the accuracy of ELISA on the situation labeling task could be ascertained.<sup>7</sup> For visualization purposes (Fig. 3), we sampled 720 points from this corpus with timestamps ranging from 2014-Aug-01 00:03 to 2014-Sep-24 23:16, and over five common types (infrastructure, water, search, medicine, and food).

To validate user experience, we recently demonstrated SAVIZ at the ACM/IE-EE ASONAM conference in 2019 and allowed the user to play with the Ebola dataset and interface, including facet selection and de-selection, and interaction with points (including drawing of bounding boxes around points). Furthermore, as evidence that the system works for arbitrary disasters, we also considered a second disaster, namely the devastating Gorkha earthquake in Nepal in 2015. This corpus was also collected over Twitter and consists of 29,946 points, with timestamps ranging from 2015-Apr-25 01:00 to 2015-May-06 09:42. Once again, ELISA was executed over this corpus to yield zero or more situation labels from a vocabulary of seven<sup>8</sup> labels (utilities, water, food, medicine, shelter, search, and infrastructure). We present a visualization of the system for the Nepal disaster in Fig. 4. For visualization purposes, we considered a sample of 1,810 tweets, with timestamps ranging from 2015-Apr-25 01:00 to 2015-May-02 06:56, and over five common types (food, medicine, shelter, infrastructure and search).

**Availability and Future Development.** More recently, we have made SAVIZ available as both a Docker container<sup>9</sup> and a GitHub project to enable open download and experimentation. These projects are respectively available at the following links.<sup>10</sup>

Finally, SAVIZ continues to undergo development and feature additions that will allow users to gain more situational insights in crisis data. An important recent addition has been the ability for users to facet and filter based on *time*. Consider, for example, Fig. 5, which now includes a slider for time to enable a user to limit

---


<sup>7</sup>The LORELEI program conducts regular evaluations to compute such numbers on specific datasets gathered and curated by the Linguistic Data Consortium.

<sup>8</sup>Slightly different versions of ELISA were available at the times of data collection; hence the vocabularies are different between Nepal and Ebola.

<sup>9</sup>To access and download the Docker package, users may have to first create a free Docker account.

<sup>10</sup><https://hub.docker.com/repository/docker/ppplinday/situation-awareness-visualization?ref=login>; <https://github.com/ppplinday/Situation-Awareness-Visualization>





**Fig. 3** The SAVIZ interface over an Ebola dataset collected over Twitter. The tabular view can be re-generated, and made more focused, by drawing a bounding box around any portion of the screen




Fig. 4 The SAVIZ interface over the Nepal earthquake dataset collected over Twitter




Fig. 5 SAVIZ with a time-filtering slider (below the x-y sliders) that allows the user more control over the visualization

what they see on the screen to a particular time frame. This facility is expected to be particularly useful for large datasets that have hundreds of thousands of points over a long (or dense) time period.

## 4 Conclusion

Despite many advances in the Artificial Intelligence and NLP communities, outputs from NLP and text classification algorithms still tend to be imperfect. In humanitarian domains, such as crisis response, first responders and other stakeholders who have to make decisions in the aftermath of crises, are not likely to trust such systems blindly. In this chapter, we presented a highly lightweight, interactive visualization platform called SAVIZ that can be deployed on a web browser in less than 30 s for thousands of tweets, and is designed for short, crisis-specific messages collected over social media like Twitter, and processed by NLP systems like ELISA. SAVIZ relies on established, pre-existing and open-source technologies from the representation learning, visualization and data processing communities. SAVIZ is backward-compatible with crisis informatics sub-systems recently released under the DARPA LORELEI program, and has been applied on real-world datasets collected from the Twitter API.

SAVIZ is intended to provide non-technical first responders with interactive situational awareness capabilities in support of crisis informatics. In the future, we are looking to extend its capabilities to help users correct existing, and provide new, annotations directly using the interface. Currently, data annotation services are expensive and require sharing of data, in addition to not being real time. SAVIZ places control firmly in the hands of the humanitarian and field users, who are best equipped to be both labeling and exploring the data. This is especially true for global crises such as COVID-19, which affect different communities and countries in different ways, and during which local decision-makers have a key role to play in mitigating the damaging impacts of the crisis.

**Acknowledgments** The authors were supported under the DARPA LORE-LEI program.

## References

1. Abel, F., Hauff, C., Houben, G. J., Stronkman, R., & Tao, K. (2012). Twitcident: fighting fire with information from social web streams. In *Proceedings of the 21st International Conference on World Wide Web* (pp. 305–308). ACM.
2. Aibaidulla, Y., & Lua, K. T. (2003). The development of tagged uyghur corpus. In *Proceedings of PACLIC17* (pp. 1–3).
3. Atefeh, F., & Khreich, W. (2015). A survey of techniques for event detection in Twitter. *Computational Intelligence*, 31(1), 132–164.
4. Cheung, L., Gowda, T., Hermjakob, U., Liu, N., May, J., Mayn, A., Pourdaghani, N., Pust, M., Knight, K., Malandrakis, N., et al. Elisa system description for lorehlt (2017).

5. Choi, S., & Bae, B. (2015). The real-time monitoring system of social big data for disaster management. In *Computer science and its applications* (pp. 809–815). Springer.
6. Chong, W. H., Lim, E. P., & Cohen, W. (2017). Collective entity linking in tweets over space and time. In *European Conference on Information Retrieval* (pp. 82–94). Springer.
7. Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011). Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12, 2493–2537.
8. Dos Santos, C. N., & Gatti, M. (2014). Deep convolutional neural networks for sentiment analysis of short texts. In *Proceedings of International Conference on Computational Linguistics (COLING)*. Dublin, Ireland.
9. Ford, J. D., Tilleard, S. E., Berrang-Ford, L., Araos, M., Biesbroek, R., Lesnikowski, A. C., MacDonald, G. K., Hsu, A., Chen, C., & Bizikova, L. (2016). Opinion: Big data has big potential for applications to climate change adaptation. *Proceedings of the National Academy of Sciences*, 113(39), 10729–10732.
10. Gao, H., Barbier, G., & Goolsby, R. (2011). Harnessing the crowdsourcing power of social media for disaster relief. *IEEE Intelligent Systems*, 26(3), 10–14.
11. Ghaeini, R., Fern, X. Z., Huang, L., & Tadepalli, P. (2016). Event nugget detection with forward-backward recurrent neural networks. In *The 54th Annual Meeting of the Association for Computational Linguistics* (p. 369).
12. He, X., Lu, D., Margolin, D., Wang, M., Idrissi, S. E., & Lin, Y. R. (2017). The signals and noise: Actionable information in improvised social media channels during a disaster. In *Proceedings of the 2017 ACM on Web Science Conference* (pp. 33–42). ACM.
13. Hong, L., Ahmed, A., Gurumurthy, S., Smola, A. J., & Tsioutsouliklis, K. (2012). Discovering geographical topics in the twitter stream. In *Proceedings of the 21st International Conference on World Wide Web* (pp. 769–778). ACM.
14. Imran, M., Castillo, C., Lucas, J., Meier, P., & Vieweg, S. (2014). Aidr: Artificial intelligence for disaster response. In *Proceedings of the 23rd International Conference on World Wide Web* (pp. 159–162). ACM.
15. Imran, M., Elbassuoni, S., Castillo, C., Diaz, F., & Meier, P. (2013). Extracting information nuggets from disaster-related messages in social media. In *ISCRAM*.
16. Irvine, A., & Klementiev, A. (2010). Using mechanical turk to annotate lexicons for less commonly used languages. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk* (pp. 108–113). Association for Computational Linguistics.
17. Jadhav, A. S., Purohit, H., Kapanipathi, P., Anantharam, P., Ranabahu, A. H., Nguyen, V., Mendes, P. N., Smith, A. G., Cooney, M., & Sheth, A. P. (2010). Twitris 2.0: Semantically empowered system for understanding perceptions from social data.
18. Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2016). Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759.
19. Kaufhold, M. A., Rupp, N., Reuter, C., & Habdank, M. (2019). Mitigating information overload in social media during conflicts and crises: design and evaluation of a cross-platform alerting system. *Behaviour & Information Technology*, 39(3) 1–24.
20. Kejriwal, M., Gilley, D., Szekely, P., & Crisman, J. (2018). Thor: Text-enabled analytics for humanitarian operations. In *Companion of the The Web Conference 2018 on The Web Conference 2018* (pp. 147–150). International World Wide Web Conferences Steering Committee.
21. Kejriwal, M., & Gu, Y. (2018). A pipeline for post-crisis twitter data acquisition. arXiv preprint arXiv:1801.05881.
22. Kejriwal, M., & Zhou, P. (2019). Low-supervision urgency detection and transfer in short crisis messages. arXiv preprint arXiv:1907.06745.
23. Kejriwal, M., & Zhou, P. (2019). Saviz: Interactive exploration and visualization of situation labeling classifiers over crisis social media data.
24. Kumar, S., Barbier, G., Abbasi, M. A., & Liu, H.: Tweettracker: An analysis tool for humanitarian and disaster relief. In *ICWSM* (2011).

25. Maaten, L. V. D., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research* 9, 2579–2605.
26. Meier, P. (2015). *Digital humanitarians: how big data is changing the face of humanitarian response*. Taylor & Francis Press.
27. Moro, A., Raganato, A., & Navigli, R.: Entity linking meets word sense disambiguation: a unified approach. *Transactions of the Association for Computational Linguistics*, 2, 231–244 (2014).
28. Morss, R. E., Demuth, J. L., Lazrus, H., Palen, L., Barton, C. M., Davis, C. A., Snyder, C., Wilhelmi, O. V., Anderson, K. M., Ahijevych, D. A., et al. (2017). Hazardous weather prediction and communication in the modern information environment. *Bulletin of the American Meteorological Society*, 98(12), 2653–2674.
29. Nazer, T. H., Xue, G., Ji, Y., & Liu, H. (2017). Intelligent disaster response via social media analysis a survey. *ACM SIGKDD Explorations Newsletter*, 19(1), 46–59.
30. Nguyen, D. T., Joty, S., Imran, M., Sajjad, H., & Mitra, P. (2016). Applications of online deep learning for crisis response using social media information. arXiv preprint arXiv:1610.01030.
31. Olteanu, A., Castillo, C., Diaz, F., & Vieweg, S. (2014). CrisisLex: A lexicon for collecting and filtering microblogged communications in crises. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)*. Oxford.
32. Palen, L., & Anderson, K. M. (2016). Crisis informaticsnew data for extraordinary times. *Science*, 353(6296), 224–225.
33. Pang, B., Lee, L., et al. (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, 2(1–2), 1–135.
34. Reuter, C., & Kaufhold, M. A. (2018). Fifteen years of social media in emergencies: a retrospective review and future directions for crisis informatics. *Journal of Contingencies and Crisis Management*, 26(1), 41–57.
35. Rogstadius, J., Vukovic, M., Teixeira, C., Kostakos, V., Karapanos, E., & Laredo, J. A. (2013). Crisistracker: Crowdsourced social media curation for disaster awareness. *IBM Journal of Research and Development*, 57(5), 4–1.
36. Schulz, A., Ristoski, P., & Paulheim, H. (2013). I see a car crash: Real-time detection of small scale incidents in microblogs. In *Extended Semantic Web Conference* (pp. 22–33). Springer.
37. Simon, T., Goldberg, A., & Adini, B. (2015). Socializing in emergenciesa review of the use of social media in emergency situations. *International Journal of Information Management*, 35(5), 609–619.
38. Stowe, K., Palmer, M., Anderson, J., Kogan, M., Palen, L., Anderson, K. M., Morss, R., Demuth, J., & Lazrus, H. (2018). Developing and evaluating annotation procedures for twitter data during hazard events. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)* (pp. 133–143).
39. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., & Mei, Q. (2015). Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 1067–1077). International World Wide Web Conferences Steering Committee.
40. Thom, D., Krüger, R., Ertl, T., Bechstedt, U., Platz, A., Zisgen, J., & Volland, B. (2015). Can twitter really save your life? A case study of visual social media analytics for situation awareness. In *Visualization Symposium (PacificVis), 2015 IEEE Pacific* (pp. 183–190). IEEE.
41. Tierney, T. F. (2014). Crowdsourcing disaster response: Mobilizing social media for urban resilience. *The European Business Review*, 80(9), 1854–1867.
42. Wang, D., Cui, P., & Zhu, W. (2016). Structural deep network embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1225–1234). ACM.
43. Wang, Q., Mao, Z., Wang, B., & Guo, L. (2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12), 2724–2743.
44. Zhang, S., & Vucetic, S. (2016). Semi-supervised discovery of informative tweets during the emerging disasters. arXiv preprint arXiv:1610.03750.