

Review article: Detection of actionable tweets in crisis events

Anna Kruspe¹, Jens Kersten², and Friederike Klan²

¹Technical University of Munich, Munich, Germany

²German Aerospace Center (DLR), Jena, Germany

Correspondence: Anna Kruspe (anna.kruspe@tum.de)

Abstract. Messages on social media can be an important source of information during crisis situations. They can frequently provide details about developments much faster than traditional sources (e.g. official news) and can offer personal perspectives on events, such as opinions or specific needs. In the future, these messages can also serve to assess disaster risks.

One challenge for utilizing social media in crisis situations is the reliable detection of relevant messages in a flood of data.

5 Researchers have started to look into this problem in recent years, beginning with crowd-sourced methods. Lately, approaches have shifted towards an automatic analysis of messages. A major stumbling block here is the question of exactly what messages are considered relevant or informative, as this is dependent on the specific usage scenario and the role of the user in this scenario. In this review article, we present methods for the automatic detection of crisis-related messages (tweets) on Twitter. We start by showing the varying definitions of importance and relevance relating to disasters, leading into the concept of use case-
10 dependent actionability that has recently become more popular, and is the focal point of the review paper. This is followed by an overview of existing crisis-related social media data sets for evaluation and training purposes. We then compare approaches for solving the detection problem based (1) on filtering by characteristics like keywords and location, (2) on crowdsourcing, and (3) on machine learning technique. We analyze their suitability and limitations of the approaches with regards to actionability. We then point out particular challenges, such as the linguistic issues concerning social media data. Finally, we suggest future
15 avenues of research, and show connections to related tasks, such as the subsequent semantic classification of tweets.

Copyright statement. TEXT

1 Introduction

During a crisis situation, quickly gaining as much information as possible about the tide of events is of crucial importance. Having access to information is necessary for developing situational awareness, and can mean the difference between life and
20 death. One source of such information that has started gaining interest in the last couple of years is social media. Twitter users, as an example, write about disaster preparations, developments, recovery, and a host of other topics (Niles et al., 2019). Retrieving this information could lead to significant improvements in disaster management strategies. In contrast to most other information sources, social media posts show up nearly immediately whenever there is a new occurrence (as long as telecommunication infrastructure is still intact), and as such can deliver information very quickly (Sakaki et al., 2010). Such

25 messages can also provide new perspectives that would not be available any other way at this speed, e.g. ground photos. In addition to factual information, social media can offer personal insights into the occurrences, as well as a back-channel to users for relief providers, government agencies, and other official institutions as well as the media. From a user perspective, 69% of Americans think that emergency response agencies should respond to calls for help sent through social media channels according to a 2010 Red Cross study (American Red Cross, 2010). A very comprehensive overview of social media usage in crisis situations is given in (Reuter and Kaufhold, 2018). Researchers have begun integrating this data source into large cross-national platforms for emergency management, e.g. in the *I-REACT*¹ (Rossi, 2019) and *E2MC*² (Havas et al., 2017) projects. The crux of social media analysis in disasters lies in the reliable retrieval and further analysis, for instance classification, of relevant messages. Twitter users worldwide generate 5,800 tweets per second on average³. In any given event, the majority of these posts will not be relevant to the event, or useful to service providers. The question is thus: What messages should be detected during a crisis event, and how can such a detection be implemented? This review article will provide an overview over existing approaches to this problem. We will focus on Twitter data as most other social media sources do not offer a possibility to obtain large amounts of their data to outside researchers, or are not commonly used in a way that facilitates gaining information quickly during a disaster.

In this context, models are commonly trained only once on a fixed set of data, making them inflexible and known to have limited generalization capability in case of new incidents. In contrast, thorough studies conducted by Stieglitz et al. (2018) and Fathi et al. (2020) revealed that interactivity and a customization of social media filtering and analysis algorithms are essential to support responses in various specific crisis situations. In order to take into account this important user-centric perspective, we focus our review not just on pre-trained general-purpose models, but also on adjustable and flexible methods that allow for more interactive data filtering and preparation for further processing.

45 In the next section, we will examine the problem definition more closely and show why the conventional concepts of “related”, “informative”, or “relevant” are problematic. Section 3 introduces social media data sets useful for analyzing the task of retrieving tweets of interest, and for training and as testing modeling approaches. In section 4, we will then show how such approaches have been implemented so far, grouped into filtering, crowdsourcing, and machine learning methods. Furthermore, aspects of adapting machine learning methods to the concept of actionability are discussed. Section 5 then goes into detail about the challenges these approaches frequently face, while section 6 briefly describes some related problems. We finish with suggestions for new developments in section 7, and a conclusion in section 8.

2 Problem definition

The task of finding social media posts in a crisis may appear clearly defined at first, but quickly becomes more convoluted when attempting an exact definition. Existing publications have gone about defining their problem statement in a variety of ways. An overview is provided in table 1.

¹<https://www.i-react.eu/>

²<https://www.e2mc-project.eu/>

³<https://www.omnicoreagency.com/twitter-statistics/>

What emerges from this table is a trichotomy between the concepts “related”, “relevant”, and “informative”. Several overlaps between these definitions can be observed. For instance, the class *not related or irrelevant* in (Nguyen et al., 2017a) contains *unrelated* tweets (like in (Burel and Alani, 2018)), but also *related but irrelevant* ones (like class *personal* in (Imran et al., 2013)). Compared to rather subjective classes, like *informative*, *personal* or *useful*, the relatedness to an event is a more
60 objective criterion. As a tentative definition, we subsume that “related” encompasses all messages that make implicit or explicit mention of the event in question. The “relevant” concept is a subset of the “related” concept, comprised of messages that contain actual information pertaining to the event. “Informative” messages, finally, offer information useful to the user of the system, and can be seen as a subset of “relevant” in turn. Not all publications necessarily follow this pattern, and lines between these concepts are blurry. In reality, many border cases arise, such as jokes, sarcasm, and speculation. In addition, the question of
65 what makes a tweet informative, or even relevant, is highly dependent on who is asking this question, i.e. who the user of this system is. Such users are often assumed to be relief providers, but could also be members of the government, the media, affected citizens, their family members, and many others. Building on top of this, each of these users may be interested in a different use case of the system, and the employed categorization may be too coarse for their purposes. For instance, humanitarian and governmental emergency management organizations are interested in understanding “the big picture”, whereas local police
70 forces and firefighters desire to find “implicit and explicit requests related to emergency needs that should be fulfilled or serviced as soon as possible” (Imran et al., 2018). These requirements also strongly depend on the availability of information from other sources, e.g. government agencies or news outlets.

In recent years, researchers have begun to address these challenges by introducing the concept of “actionability” to describe information relevance from the end user perspective of emergency responders (He et al., 2017) as opposed to generalized
75 situational awareness. Zade et al. (2018) loosely define actionability as “information containing a request or a suggestion that a person should act on and an assumption that a message actionable to some responders may be irrelevant to others”, while McCreadie et al. (2020) specify it implicitly via certain topical classes. The concept “serviceability” as introduced in (Purohit et al., 2018) is similar, but only applies to messages directly addressed to relief providers and is defined more narrowly. Similarly, according to (Kropczynski et al., 2018), a “golden tweet” – a post on Twitter containing actionable information for
80 emergency dispatch and supporting the immediate situational awareness needs of first responders – should contain information that addresses the well known five W’s (where, what, when, who, why) as well as information on weapons.

In this paper, we define an actionable tweet in a crisis event as one that is relevant and informative in a certain use case or to a certain user. Naturally, focusing on user-centric actionability adds complexity to the corresponding methodological and technical systems, and many of the presented methods do not yet offer this flexibility. However, we believe that this is a viable
85 path forward to make such systems more useful in real-life situations. For the remainder of the paper, we will point out how existing data sets and methods can be enhanced in the future to make systems adaptable to individual requirements by different users. We deliberately do not focus on specific use cases, but rather on approaches to guarantee this adaptability.

An aspect that is often neglected in social media-based crisis analytics is the existence of mature and well-established workflows for disaster response activities that have so far been mainly based on geo-data and remote sensing (Voigt et al., 2016; Lang
90 et al., 2020). Information from social media channels should therefore not be seen as solitary but rather as an additional,

Table 1. Overview of class definitions for filtering crisis-related tweets

| <i>Article</i> | <i>Class</i> | <i>Definition</i> |
|-------------------------------|----------------------------------|---|
| (Imran et al., 2013) | <i>Personal</i> | A message only of interest to its author and her immediate circle of family/friends - does not convey any useful information to people who do not know its author |
| | <i>Informative</i> | Messages of interest to other people beyond the author's immediate circle |
| | <i>Other</i> | Not related to the disaster |
| (Parilla-Ferrer et al., 2014) | <i>Informative</i> | A tweet provides useful information to the public and is relevant to the event |
| | <i>Uninformative</i> | Tweets that are not relevant to the disaster and these do not convey enough information or are personal in nature and may only be beneficial to the family or friends of the sender |
| (Caragea et al., 2016) | <i>Informative</i> | Useful information |
| | <i>Not informative</i> | Not relevant to the event and no useful information |
| (Win and Aung, 2017) | <i>Informative</i> | Useful information |
| | <i>Not informative</i> | Not relevant to the event and no useful information |
| | <i>Other information</i> | Messages related to the event but without useful information |
| (Nguyen et al., 2017a) | <i>Useful/Relevant</i> | Information that is useful to others |
| | <i>Not related or irrelevant</i> | Not related to the event or does not contain useful information for others |
| (Burel and Alani, 2018) | <i>Crisis related</i> | Message related to a crisis situation in general without taking into account informativeness or usefulness |
| | <i>Non-crisis related</i> | Message that is not related to a crisis situation |
| (Stowe et al., 2018) | <i>Relevant</i> | Any information that is relevant to disaster events, including useful information but also jokes, retweets, and speculation |
| | <i>Irrelevant</i> | Not related to a disaster event |

complementary source of information. In this context, further interesting use-cases, corresponding questions and problem definitions arise in which social media may fill temporal gaps between satellite data acquisitions, could be used to identify areas that need assistance, and to trigger local surveys.

3 Data sets

95 Collections of social media data created during crises are necessary to study what users write about, how this develops over time, and to create models for automatic detection and other tasks. For these reasons, several such data sets have already been

created. As mentioned above, Twitter is the most salient source of data for this use case; therefore, available data sets are mainly composed of Twitter data.

100 Table 2 lists an overview of available Twitter data sets collected during disaster events. These mainly focus on the text content of tweets, except for *CrisisMMD* which provides tweets with both text and images. Some of these data sets only contain data for one event, while others aggregate multiple ones. Based on various existing data sets, Wiegmann et al. (2020a) recently proposed a balanced compilation of labeled Tweets from 48 different events covering the ten most common disaster types. A distinction can also be made for corpora focusing on natural disasters and those also including man-made disasters. *Events2012* goes even further, containing around 500 events of all types, including disasters.

105 Annotations vary between these data sets. Some of them do not contain any labels beyond the type of event itself, while others are labeled according to content type (e.g. “Search and rescue” or “Donations”), information source (first-party observers, media, etc.), and priority or importance of each tweet (*CrisisLexT26* and *TREC-IS 2019B*).

A general issue with these data sets lies in the fact that researchers cannot release the full tweet content due to Twitter’s redistribution policy⁴. Instead, these data sets are usually provided as lists of tweet ID’s, which must then be expanded to the full information (“hydrated”). This frequently leads to data sets becoming smaller over time as users may choose to delete their tweets or make them private. For instance, as of September 2020, only ~30 % of all labeled Tweets from the *Events2012* data set are available. Additionally, the teams creating these corpora have mainly focused on English- and occasionally Spanish-language tweets to facilitate their wider usage for study. More insights would be possible if tweets in the language(s) of the affected area were available. However, Twitter usage also varies across countries. Another factor here is that less than 1% of all tweets contain geolocations (Sloan et al., 2013), which are often necessary for analysis. The following sections provide descriptions of the data sets in more detail:

110
115

Events2012 This data set was acquired between October 9 and November 7 in 2012 and contains 120 million tweets, of which around 150,000 were labeled to belong to one of 506 events (which are not necessarily disaster events) (McMinn et al., 2013). The event types are categorized into eight groups, such as “Business & Economic” “Arts, Culture & Entertainment”, “Disasters & Accidents”, or “Sports”.

120

CrisisLexT6 and T26 *CrisisLexT6* was first published by Olteanu et al. (2014) and expanded later to *CrisisLexT26* (Olteanu et al., 2015). The sets contain tweets collected during 6 and 26 crises, respectively, mainly natural disasters like earthquakes, wildfires and floods, but also human-induced disasters like shootings and a train crash. Amounts of these tweets per disaster range between 1,100 and 157,500. In total, around 285,000 tweets were collected. They were then annotated by paid workers on the *CrowdFlower* crowdsourcing platform⁵ according to three concepts: Informativeness, information type, and tweet source.

125

⁴<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

⁵Later named *Figure Eight*, <https://www.figure-eight.com/>; acquired in 2019 by *Appen*, <https://appen.com>

Table 2. Overview of crisis-related Twitter data sets

| <i>Name</i> | <i># Labeled tweets</i> | <i># Total tweets</i> | <i>Labeled concepts (#classes)</i> | <i>Covered event types</i> |
|---|-------------------------|-----------------------|---|---|
| Events2012 (McMinn et al., 2013) | ~150,000 | 120 mio. | 506 Events (8) | Disasters and accidents, other events in sports, arts, culture and entertainment |
| CrisisLexT6 (Olteanu et al., 2014) | ~6,000 | ~6,000 | Relatedness (2) | Hurricane, flood, bombing, tornado, explosion |
| CrisisLexT26 (Olteanu et al., 2015) | 26,000 | 285,000 | Informativeness (2), information type (6), tweet source (6) | Earthquake, flood, wildfire, meteor, typhoon, flood, explosion, bombing, train crash, building collapse |
| Disasters on Social Media (DSM) (Crowdflower, 2015) | ~10,000 | ~10,000 | Relevance (4) | Not provided |
| Incident-related Twitter Data (IRTD) (Schulz and Guckelsberger, 2016) | ~21,000 | ~21,000 | Relatedness (2), incident type (4) | Crash, fire, shooting |
| CrisisNLP (Imran et al., 2016b) | 23,000 | 53 mio. | Information type (9) | Earthquake, hurricane, flood, typhoon, cyclone, ebola, MERS |
| CrisisMMD (Alam et al., 2018b) | 16,000 | 16,000 | Informativeness (2), information type (8), 3 damage severity (3) | Hurricane, earthquake, wildfire, flood |
| Epic (Stowe et al., 2018) | ~25,000 | 25,000 | Relevance (2), information type (17), sentiment (3) | Hurricane |
| Disaster Tweet Corpus 2020 (DTC) (Wiegmann et al., 2020b, a) | ~150,000 | ~5.1 mio. | Relatedness (2) | Biological, earthquake, tornado, hurricane, flood, industrial, societal, transportation, wildfire |
| TREC-IS 2019B (McCreadie et al., 2019, 2020) | ~38,000 | ~45,000 | Information type (25), priority (4), actionability (2) | Bombing, earthquake, flood, typhoon/hurricane, wildfire, shooting |
| Appen Disaster Response Messages (Appen Ltd., 2020) | ~30,000 | ~30,000 | Information type (36) | Earthquake, flood, hurricane |
| Storm-related Social Media (SSM) (Grace, 2020) | ~22,000 | 22,000 | Relatedness (2), information type (19), aggregated information type (6), 3 toponym concepts (2/2/3) | Tornado |

Disasters on Social Media (DSM) This resource is available on *CrowdFlower*⁶ and contains around 10,000 tweets that were identified via keyword-based filtering (for example “ablaze”, “quarantine”, and “pandemonium”). At its finest granularity, four different classes are distinguished: (1) Relevant (65.52 %), (2) Not Relevant (27.59 %), (3) Relevant Can’t Decide (4.6 %), and (4) Not Relevant Can’t Decide (2.3 %). No information regarding the covered event types is available, but a cursory review of the data reveals that a multitude of events is found with the keywords, e.g. floods, (wild)fires, car crashes, earthquakes, typhoons, heat waves, plane crashes, terrorist attacks, etc.

Incident-related Twitter Data (IRTD) Within three time periods in 2012–2014, around 15 million tweets in a 15 km radius around the city centers of Boston (USA), Brisbane (AUS), Chicago (USA), Dublin (IRE), London (UK), Memphis (USA), New York City (USA), San Francisco (USA), Seattle (USA) and Sidney (AUS), were collected. After filtering by means of incident-related keywords, redundant tweets and missing textual content, the remaining set of around ~21,000 tweets was manually labeled by five annotators using the *CrowdFlower* platform. The annotators labeled according to two different concepts: (1) 2 classes: “incident related” and “not incident related”, and (2) 4 classes: “crash”, “fire”, “shooting”, and a neutral class “not incident related”. Manual labels for which the annotator agreement was below 75 % were discarded (Schulz and Guckelsberger, 2016).

CrisisNLP The team behind *CrisisNLP* collected tweets during 19 natural and health-related disasters between 2013 and 2015 on the *AIDR* platform (see section 4.2) using different strategies (Imran et al., 2016b). Collected tweets range between 17,000 and 28 million per event, making up around 53 million in total. Out of these, around 50,000 were annotated both by volunteers and by paid workers on *CrowdFlower* with regard to nine information types.

CrisisMMD *CrisisMMD* is an interesting special case because it only contains tweets with both text and image content. 16,000 tweets were collected for seven events that took place in 2017 in five countries. Annotation was performed by *Figure Eight* for text and images separately. The three annotated concepts are: Informative/Non-informative, eight semantic categories (like “Rescue and volunteering” or “Affected individuals”), and damage severity (only applied to images) (Alam et al., 2018b).

Epic This data set with a focus on Hurricane Sandy was collected in a somewhat different manner than most others. The team first assembled tweets containing hashtags associated with the hurricane, and then aggregated them by user. Out of these users, they selected those who had geotagged tweets in the area of impact, suggesting that these users would have been affected by the hurricane. Then, 105 of these users were selected randomly, and their tweets from a week before landfall to a week after were assembled. This leads to a data set that in all probability contains both related and unrelated tweets by the same users. Tweets were annotated according to their relevance as well as 17 semantic categories (such as “Seeking info” or “Planning”) and sentiment (Stowe et al., 2018).

Disaster Tweet Corpus 2020 (DTC) This data set contains tweets collected, annotated, and published in several other works (Imran et al., 2014; Olteanu et al., 2014, 2015; Imran et al., 2016c; Alam et al., 2018b; Stowe et al., 2018; McMinn et al.,

⁶<https://data.world/crowdflower/disasters-on-social-media>

160 2013), and covers 48 disasters over 10 common disaster types. This balanced collection is intended as a benchmarking data set for filtering algorithms in general (Wiegmann et al., 2020b, a). Additionally, a set of 5 million unrelated tweets, collected during a tranquil period, i.e., where no disasters happened, is provided. This is intended to test filtering models in terms of false positive rates.

165 **TREC-IS 2019B** A crisis classification task named “Incident Streams” has been a part of the Text REtrieval Conference (TREC) organized by NIST since 2018 (McCreadie et al., 2019). In the first iteration, tweets for six events were first collected automatically using a pre-defined list of keywords, and then annotated with one of 25 information type categories. Further iterations were conducted twice in 2019, for which the data set was expanded each time through a sophisticated process of crawling Twitter and then downsampling the results. The format was also changed to allow multiple labels per tweet. There are several subsets that have been flexibly used for training and testing in the task, partially comprised of *CrisisNLP* and *CrisisLex*. We show the 2019B iteration here, but each iteration has been composed of somewhat different data, comprising 48 crisis events, 50,000 tweets, and 125,000 labels in total. In the 2020 iterations, only events that took place in 2019 were included (McCreadie et al., 2020). *TREC-IS* also contains a concept of actionability defined by a selection of the semantic classes.

175 **Appen Disaster Response Messages** This data set was published in an open-source format originally by *Figure Eight*, now part of private company *Appen* (Appen Ltd., 2020). It contains 30,000 messages split into training, test, and validation sets collected during various disaster events between 2010 and 2012. These tweets are annotated according to 36 content categories, such as “Search and rescue”, “Medical help”, or “Military”, as well as with a “Related” flag. These messages contain multiple languages plus English translations. The data set also includes news articles related to disasters. The data set is used in a Udacity course⁷ as well as a Kaggle challenge⁸.

180 **Storm-related Social Media (SSM)** Presented in (Grace, 2020), this data set was collected during a 2017 tornado in Pennsylvania using three methods: Filtering by Twitter-provided geolocation in the affected area; keyword filtering by place names in the affected area; and filtering by networks of users located in the affected county. For the last approach, user IDs are available in a supplementary data set. Tweets were then labeled according to six concepts: Relatedness to the storm; semantic information type (subsumed from other publications, e.g. (Olteanu et al., 2015)); an aggregated set of the semantic information types (e.g. disruptions, experiences, forecasts); and three toponym-related concepts. Labeling was done by three assessors for part of the data set, then split between them for the rest, after consolidating discrepancies. 185 The data is available as supplementary material for (Grace, 2020)⁹.

All presented data sets offer advantages and disadvantages, depending on the use case. Almost all of them contain information type annotations, but there is no universal agreement on an ontology here. Many of the used information type definitions are compatible across data sets, but this requires manual work. In addition, interpretation that may lead to errors is required, on

⁷<https://www.udacity.com/course/data-scientist-nanodegree--nd025>

⁸<https://www.kaggle.com/jannesklaas/disasters-on-social-media>

⁹<https://www.sciencedirect.com/science/article/pii/S2352340920304893>

190 the one hand because the classes are often not clearly defined, and on the other because even the meanings of classes with the same name can vary between data sets. The information type ontology provided in *TREC-IS 2019B* was developed and refined in collaboration with help providers, and could therefore be a valuable basis for future annotations.

In published works, *CrisisNLP* and *CrisisLexT26* are used most frequently to demonstrate novel approaches because they are relatively large and cover a wide range of event types. As mentioned above, the *Appen* material is used in Udacity courses
195 and on Kaggle, and may therefore also be a useful starting point for new researchers. For detection of disaster-related tweets, *Events2012* is also very interesting because it contains both disaster events as well as other events, and is much larger than the two others. It does not contain information type annotations, however.

All four of these data sets contain tweets created before 2017, which is relevant because the character limit for tweets was increased from 140 to 280 in 2017. For a large data set of newer tweets, the latest iteration of the *TREC-IS* set is very inter-
200 esting. In addition, existing approaches for this data set can be recreated from the TREC challenge. *CrisisMMD* has not been used as frequently so far, but is interesting because of the added image content. This data set as well as *Epic* and *SSM* does not cover as many different events, but in exchange, they have a much wider selection of labeled concepts that have not received as much attention so far. *DTC* is interesting due to its aggregation of several data sets and resulting large size and wide coverage, making it usable for benchmarks.

205 All of these data sets operate under the notion of “related”/“informative”/“relevant” tweets, either by providing explicit labels for these concepts, or by assuming that all contained tweets belong to these concepts. As described in section 2, these conventional annotations are too rigid to implement a detection of actionability for different use cases. We suggest two solutions for future systems:

1. Explicitly annotating tweets with use case-dependent actionability labels. This is, of course, a costly option, but would
210 be highly interesting as a starting point for developing adaptable systems.
2. Defining actionability in a use case-specific way as a composite of other (basic) concepts. A data set labeled with those basic concepts could then be used for different use cases. This is, for example, done in the *TREC-IS 2019B* data set through a selection of information type classes, primarily request and report classes. With the refined ontologies of information types and other concepts contained in the presented data sets, individual profiles of relevant concepts and
215 event types could be created per use case to define actionability in future research. These profiles could even be inferred by automatic models.

4 Approaches

As described above, users generate huge amounts of data on Twitter every second, and finding tweets related to an ongoing event is not trivial (Landwehr and Carley, 2014). Several detection approaches have been presented in literature so far. We
220 will group them into three categories: Filtering by characteristics, crowdsourcing, and machine learning approaches. As researchers have only started to focus on detecting actionable information in recent years, many of the presented methods do

not offer the necessary flexibility yet, instead only offering solutions for specific use cases or the generalized task of finding related/relevant/informative tweets in a crisis event. Nevertheless, we will present them here as a very useful basis for future work, and will point out whether the described approaches are already useful for detecting actionable information or how they can be adapted accordingly. These questions are somewhat easier to answer for filtering by characteristics and crowdsourcing (sections 4.1 and 4.2) because such systems need to be invoked for specific tasks in a new event anyway. For machine learning methods however, models are usually trained on data from past events or tasks and then statically used in newly occurring events, as described in section 4.3. In section 4.4, we point out novel directions of research for also adapting machine learning algorithms to desired new tasks, implementing the actionability concept.

230 4.1 Filtering by characteristics

The most obvious strategy is the filtering of tweets by various surface characteristics. An example is *TweetTracker*, which was first presented in 2011 (Kumar et al., 2011) and is still available¹⁰. This platform is able to collect tweets by hashtag, keyword, or location, perform trend analysis, and provide visualizations.

Keywords and hashtags are used most frequently for this and often serve as a useful pre-filter for data collection (e.g. in (Lorini et al., 2019) where tweets are pre-filtered by geographic keywords). The Twitter API allows searching directly for keywords and hashtags or recording the live stream of tweets containing those, meaning that this approach is often a good starting point for researchers. This is especially relevant because only 1 % of the live stream can be collected for free (also see section 5) - when a keyword filter is employed, this 1 % is more likely to contain relevant tweets.

Olteanu et al. (2014) developed a lexicon called *CrisisLex* for this purpose. However, the keyword-filtering approach easily misses tweets that do not mention the keywords specified in advance, particularly when changes occur or the attention focus shifts during the event. To tackle this recall-related problem, Olteanu et al. (2014) propose a method to update the keyword list based on query expansion using new messages. A further, semi-supervised dynamic keyword generation approach, utilizing incremental clustering, SVMs, expectation maximization and word graph generation, is proposed in (Zheng et al., 2017).

Another problem with keyword lists is that unrelated data that contains the same keywords may be retrieved (Imran et al., 2015). In general, filtering by keywords is not a very flexible approach to tackle different use cases and therefore implement actionability. Nevertheless, such approaches have been used in insightful studies, e.g. in (de Albuquerque et al., 2015), where keyword-filtered tweets during a flood event were correlated with flooding levels.

Geolocation is another frequently employed feature that can be useful for retrieving tweets from an area affected by a disaster. However, this approach misses important information that could be coming from a source outside the area, such as help providers or news sources. Additionally, only a small fraction of tweets is geo-tagged at all, leading to a large amount of missed tweets from the area (Sloan et al., 2013).

¹⁰<http://tweettracker.fulton.asu.edu/>

4.2 Crowdsourcing approaches

To resolve the problems mentioned above, other strategies were developed. One solution lies in crowdsourcing the analysis of tweets, i.e. asking human volunteers to manually label the data (Poblet et al., 2014). From an actionability standpoint, this may seem ideal because human subjects are fairly good judges of whether a tweet is relevant in a specific use case. However, this seemingly easy task can easily turn into a complex problem that is subject to the individual volunteers' interpretation depending on the situation. Partitioning the problem into sub-tasks that can be judged more easily can be a remedy to this (Xu et al., 2020).

The main disadvantage of crowdsourcing lies in the necessity for many helpers due to the large amount of incoming tweets, and the resulting effort necessary to organize tasks and consolidate results. Nevertheless, volunteers can be extremely helpful in crisis situations. Established communities of such volunteers exist and can be activated quickly in a disaster event, for example the *Standby Task Force*¹¹.

To facilitate their work, platforms have been developed over the years. One of the most well-known systems is *Ushahidi*¹². This platform allows people to share situational information in various media, e.g. by text message, by e-mail, and of course by Twitter. Messages can then be tagged with categories relevant to the event. *Ushahidi* was started by a team of Kenyan citizens during the 2007 Kenyan election crisis, and has since been used successfully in a number of natural disasters, humanitarian crises, and elections (for monitoring). Both the server and the platform software are available open-source¹³. Efforts were made to integrate automatic analysis tools into the platform (named "SwiftRiver"), but discontinued in 2015.

Such automatic analysis tools are the motivation for *AIDR* (Imran et al., 2015). *AIDR* was first developed as a quick response to the 2013 Pakistan earthquake. Its main purpose lies in facilitating machine learning methods to streamline the annotation process. In a novel situation, users first choose their own keywords and regions to start collecting a stream of tweets. Then, volunteers annotate relevant categories. A supervised classifier is then trained on these given examples, and is automatically applied to new incoming messages. A front-end platform named *MicroMappers*¹⁴ also exists. *AIDR* is available in an open-source format as well¹⁵. It has been used in the creation of various data sets and experiments.

Another contribution to crowdsourcing crisis tweets is *CrisisTracker* (Rogstadius et al., 2013). In *CrisisTracker*, tweets are also collected in real-time. Local Sensitive Hashing (LSH) is then applied to detect clusters of topics (so-called stories), so that volunteers can consider these stories jointly instead of single tweets. The *AIDR* engine has also been integrated to provide topic filtering. As a field trial, the platform was used in the 2012 Syrian civil war. *CrisisTracker* is also available free and open-source¹⁶, but maintenance stopped in 2016.

¹¹<https://www.standbytaskforce.org/>

¹²<https://www.ushahidi.com/>

¹³https://github.com/ushahidi/Ushahidi_Web

¹⁴<https://micromappers.wordpress.com/>

¹⁵<https://github.com/qcri-social/AIDR>

¹⁶<https://github.com/JakobRogstadius/CrisisTracker/>

4.3 Machine learning approaches

To forgo the need for many human volunteers while still intelligently detecting crisis-related tweets, various machine learning approaches have been developed over the years. We distinguish between two categories here: “Traditional” machine learning approaches that put an emphasis on NLP feature engineering, and deep learning approaches with Neural Networks that often
285 utilize automatically learned word or sentence embeddings. An overview of proposed methods of both types is given in table 3.

Generally, machine learning approaches all follow the same rough processing pipeline which is outlined in figure 1. Pre-processed text data is fed into a feature extraction method, and the generated features are forwarded to a model that then outputs a result. In deep learning approaches, this model is a neural network. Feature extraction and model training/inference
290 used to be separate processes in classical NLP, but have become increasingly combined over the past years with the arrival of word and sentence embeddings that can be integrated into the training process.

In both flavors of machine learning, research has mainly focused on static general-purpose models trained a single time on known data to reduce social media information overload. These models are usually intended to detect messages that are potentially relevant to crisis situations. An immediate applicability comes at the cost of a limited generalization capability, i.e.
295 in case of new events and especially new event types, the models may fail dramatically (see for example experimental results in (Wiegmann et al., 2020b)). Furthermore, a decision is usually made on tweet-level without taking into account thematically, spatially or temporally adjacent information. As pointed out in section 2, it is now becoming apparent that more user-centric perspectives need to be taken into account (i.e. defining actionability for a certain task). Hence, more adjustable and flexible methods that allow for more interactive data filtering by actionability are also reviewed here (see section 4.4). These methods
300 do not necessarily focus on the filtering task itself, but can be used in this context and may provide additional valuable capabilities, like an aggregation of semantically similar messages, to support the understanding of contained information and their changes.




Figure 1. General processing pipeline for machine learning approaches.

4.3.1 Machine learning based on feature engineering

Linguistic features

305 A crucial component of a social media classification model is the representation of the text data at the input (i.e. how words or sentences are mapped to numeric values that the model can process). Classical NLP features are based in linguistics and may employ additional models, e.g. for sentiment analysis or topic modeling.

Table 3. Overview of the related work proposing filtering algorithms, ordered by the employed method, and listing the data sets used.

| Reference | Features | Method | Data |
|--|--|---|---|
| <i>Machine learning based on feature engineering</i> | | | |
| Parilla-Ferrer et al. (2014) | BoW | NB, SVM | Own data |
| Stowe et al. (2016) | Time, retweet, URLs, unigrams, NER, POS, <i>Word2vec</i> | NB, Maximum Entropy, SVM | Own data |
| To et al. (2017) | BoW, TF-IDF (with PCA) | LR | <i>CrisisLexT26, DSM</i> |
| Win and Aung (2017) | POS, n-grams, emotions, word cluster, lexicon-based features, URLs, hashtags | Linear classification, SVM | <i>CrisisLexT6</i> |
| Habdank et al. (2017) | Term counts, TF-IDF, n-grams | NB, Decision Tree, SVM, RF, ANN | Own data |
| Resch et al. (2018) | BoW | LDA | Own data |
| Li et al. (2018) | Term occurrence | NB, semi-supervised domain adaptation | <i>CrisisLexT6</i> |
| Mazloom et al. (2019) | BoW | NB, RF, domain adaptation | <i>CrisisLexT6, IRTD</i> |
| Kejriwal and Zhou (2019) | <i>fastText</i> | Linear Regression ensemble, semi-supervised | Own data |
| Kaufhold et al. (2020) | BoW, TF-IDF, NER, author-event distance, RT, URLs, media, tweet length, language | RF: active, incremental and online learning | Own data |
| <i>Neural networks</i> | | | |
| Caragea et al. (2016) | BoW, n-grams | CNN | <i>CrisisLexT26</i> |
| Nguyen et al. (2016b) | Domain-specific <i>Word2vec</i> | CNN, online training | <i>CrisisNLP</i> |
| Nguyen et al. (2017a) | <i>Word2vec</i> (Mikolov et al., 2013), own crisis word embedding, | CNN | <i>CrisisLexT6, CrisisNLP</i> |
| Alam et al. (2018a) | <i>Word2vec</i> , graph embedding | CNN, adversarial and semi-supervised learning | <i>CrisisNLP</i> |
| Burel and Alani (2018) | <i>Word2vec</i> | CNN | <i>CrisisLexT26</i> |
| Kersten et al. (2019) | <i>Word2vec</i> (Imran et al., 2016a) | CNN | <i>CrisisLexT26, CrisisNLP, Epic, Events2012, CrisisMMD</i> |
| Kruspe et al. (2019) | <i>Word2vec</i> (Nguyen et al., 2016a) | CNN few-shot model | <i>CrisisLexT26, CrisisNLP</i> |
| Ning et al. (2019) | Autoencoder: Linguistic, emotional, symbolic, NER, LDA | CNN | <i>CrisisLexT26</i> |
| Lorini et al. (2019) | Multilanguage-adapted <i>GloVe</i> | CNN | Own data (floods) |
| Wiegmann et al. (2020b) | USE, BERT, (Imran et al., 2016a) | CNN, NN | <i>DTC</i> |
| Snyder et al. (2020) | <i>Word2vec</i> | CNN, RNN, LSTM | <i>CrisisLexT26, Own data</i> |
| de Bruijn et al. (2020) | Multilanguage-adapted <i>fastText</i> | CNN + multimodal NN | Own data (floods) |

A corpus (i.e. set) of documents (i.e. tweets) is built up by a vocabulary of N words. A straightforward approach to represent each word is a “one-hot” vector of length N . Given the i^{th} word of the vocabulary, the corresponding one-hot vector is 1 at position i and zero otherwise. Depending on the vocabulary size, these vectors might be quite large and the one-hot representation does not allow for direct comparison of different words, e.g. with Euclidean or Cosine similarity.

Within this framework, a Bag-of-Words (BoW) model simply counts the occurrence of each term (term frequency–TF) in a document or corpus independently of its position. In order to reduce the impact of frequently occurring but not descriptive terms, like “a” or “and”, these so-called stop words can be removed in advance or the term frequencies are normalized, for example by the commonly used inverse document frequency (IDF). TF-IDF results in high weights in case of a high term frequency (in a document) along with a low term frequency over the whole corpus. Even though this approach proved to be suitable in many studies (Parilla-Ferrer et al., 2014; To et al., 2017; Resch et al., 2018; Mazloom et al., 2019; Kaufhold et al., 2020), contextual information is neglected. The concept of n-grams accounts for context in terms of n adjacent terms. However, this approach may drastically increase the vocabulary dimensionality.

Further commonly used features (see for example (Stowe et al., 2016; Kaufhold et al., 2020)) result from part-of-speech (POS) tagging and named entity recognition (NER). POS tagging finds the syntactic category of each words (e.g., noun, verb or adjective) in written text, whereas NER allows for tagging all words representing given names, for example of countries, places, companies, and persons. The extracted features are sometimes subjected to dimensionality reduction procedures such as Principal Component Analysis (PCA) before the model input.

Finally, Twitter-specific features, like tweet length, timestamp, whether a tweet is a retweet, whether a tweet contains media, links, emojis, usernames, or hashtags, have been found to be useful features (see for example (Stowe et al., 2016; Win and Aung, 2017; Kaufhold et al., 2020)).

A few approaches also use neural network-based word embeddings, e.g. *Word2vec* and *fastText*, which are described below.

Models

Based on the feature vectors that represent a tweet, several methods are available to train models that seek to assign each tweet to pre-defined classes. The task of distinguishing crisis- or incident-related content from all other types of tweets is a binary problem, for which generative and discriminative approaches exists. Generative approaches attempt to model the joint probability of the features and the corresponding labels. Even the relatively simple Naïve Bayes approach produces promising results, for example in (Parilla-Ferrer et al., 2014; Stowe et al., 2016; Habdank et al., 2017; Mazloom et al., 2019).

In contrast, discriminative methods, like Support Vector Machines (SVMs), decision trees, Random Forests (RFs) and Logistic Regression (LR), are commonly used to directly distinguish between classes (see for example (Win and Aung, 2017; Kejriwal and Zhou, 2019)). For instance, a linear SVM estimates the hyperplane that separates the two classes in the feature space without modeling the distribution of these classes.

Some proposed methods also take an indirect approach to the binary classification task, such as (Resch et al., 2018) where Latent Dirichlet Allocation (LDA) (Blei et al., 2003) is used for topic modeling, and the resulting topic clusters are then analyzed further.

4.3.2 Neural networks

In recent years, neural networks have come to the forefront of research. In contrast to the models in the previous section, deep neural networks allow for more powerful and complex modeling, but also require more data and computational resources to train them, and their decisions are often less transparent. The last point can be particularly grave if critical decisions are made based on these models. Another difference is that they commonly do not use linguistically motivated features as their inputs, but instead use word or sentence embedding layers at the inputs, which are neural networks themselves. These embeddings are often pre-trained on even larger data sets, but can also be integrated into the training process for finetuning or training from scratch.

350 Neural network features & embeddings

As mentioned, hand-crafted features have become more and more replaced with automatically trained word embeddings since their inception in 2011 (Collobert et al., 2011). These embeddings are neural networks themselves, and are part of the complete classification network. Multiple refinements have been proposed over the years. Many approaches for crisis tweet detection employ *Word2vec*, a pre-trained word embedding that was first presented in 2013 (Mikolov et al., 2013) and has since been expanded in various ways. A version specifically trained on crisis tweets is presented in (Imran et al., 2016b). Burel et al. (2017a) integrate semantic concepts and entities from *DBPedia*¹⁷. *GloVe* (Pennington et al., 2014) and *fastText* (Joulin et al., 2016) embeddings follow a similar idea, and are expanded for multilingual tweet classification in (Lorini et al., 2019) and (de Bruijn et al., 2020) respectively, based on the adaptation method proposed by Lample et al. (2018).

In the past two years, BERT (Devlin et al., 2019) embeddings and their various offshoots have become very popular (McCreadie et al., 2020). These embeddings still function on the word level, but take complex contexts into account. A crisis-specific version is proposed in (Liu et al., 2020). In another direction, embeddings that do not represent words but whole sentences are also becoming used more widely, e.g. in (Kruspe, 2020; Kruspe et al., 2020; Wiegmann et al., 2020b). The most prominent example is the Universal Sentence Encoder (USE) (Cer et al., 2018) and its multilingual version (MUSE) (Yang et al., 2019). In most cases, versions of embeddings that are pre-trained on large text corpora are used. These corpora are not necessarily social media texts or crisis-related, but the models have been shown to produce good results anyway. The advantage of using pre-trained models is that they are easy to apply, and do not require as much training data (Wiegmann et al., 2020b). In the case of sentence-level embeddings, their usage also leads to a simplification of the subsequent network layers as the embeddings themselves already capture the context of the whole sentence. As mentioned above, versions finetuned to the task are also available for many common embeddings. A comparison of various word and sentence embeddings for crisis tweet classification can be found in (ALRashdi and O’Keefe, 2019).

It should also be mentioned that occasionally, deep models also utilize the linguistic features described above, e.g. (Ning et al., 2019). In the first iteration of the TREC-IS challenge, several approaches produced good results with such hand-crafted features as well (McCreadie et al., 2019). Their advantage lies in the fact that they do not need to be trained, and can therefore

¹⁷<https://wiki.dbpedia.org/>

work with a small amount of data, which may sometimes be the case in new crises.

375

Classification networks

Extracted features, which may be embeddings are then fed into a subsequent neural network. In most crisis-related use cases, these will be classification models, although regression models are occasionally used for binary concepts like relevance, priority, or similarity, as well as sentiment. Commonly, text processing tasks employ Recurrent Neural Networks to leverage longer context, but in short text tasks, Convolutional Neural Networks (CNNs) are more popular.

Caragea et al. (2016) first employed CNNs for the classification of tweets into those informative with regards to flood events and those not informative. Lin et al. (2016) also applied CNNs to social media messages, but for the *Weibo* platform instead of Twitter. In many of the following approaches, a type of CNN developed by Kim for text classification is used (Kim, 2014), such as in (Burel and Alani, 2018; de Bruijn et al., 2020; Kersten et al., 2019). A schematic is shown in figure 2. These methods achieve accuracies of around 80% for the classification into related and unrelated tweets. In (Burel and Alani, 2018) as well as in (Burel et al., 2017a) and (Nguyen et al., 2016b), this kind of model is also used for information type classification. Recently, these CNN architectures have been expanded in different directions. Ning et al. (2019) show a multi-task variant. In (Burel et al., 2017a), a CNN with word embedding inputs is combined with one for semantic document representations. The resulting system is packaged as *CREES* (Burel and Alani, 2018), a service that can be integrated into other platforms similar to *AIDR*. Snyder et al. (2020) and Nguyen et al. (2016b) show active learning approaches that allow adapting the CNN over the progress of a crisis as new tweets arrive, dovetailing with the crowdsourcing systems described above. More novel approaches for adaptation to actionability are described in the next section.




Figure 2. CNN for text classification as proposed by Kim (2014).

4.4 Adaptation to actionability

All of the approaches mentioned above aim to generalize to any kind of event on tweet level without any *a priori* information, and can therefore not easily adapt to specific use cases. The transferability of pre-trained models to new events and event types is thoroughly investigated in (Wiegmann et al., 2020b). A real-world system may not need to be restricted in this way; in many cases, its users will already have some information about the event, and may already have spotted tweets of the required type. This removes the need to anticipate any type of event. It also directs the system towards a specific event rather than any event happening at that time.

As a consequence, a shift from static pre-trained models to more adaptable and flexible machine learning methods is required. Approaches such as semi-supervised learning of regression model ensembles (Kejriwal and Zhou, 2019), domain adaptation (Mazloom et al., 2019; Poblete et al., 2018), as well as active, incremental and online learning using Random Forests (Kaufhold et al., 2020) demonstrate that traditional pre-trained models can also be utilized in a more interactive fashion and therefore have the potential to better fit to needs of emergency responders. With respect to deep learning, Li et al. (2018) and Mazloom et al. (2019) show that models adapted to the domain of the event can perform better than generalized models. Alam et al. (2018a) propose an interesting variant for neural networks: Their system includes an adversarial component which can be used to adapt a model trained on a specific event to a new one (i.e. a new domain). Pre-trained embeddings play a key role in transfer learning or finetuning to new events, as they provide a large amount of pre-existing linguistic knowledge to the model, and therefore reduce the necessity for large amounts of training data (Snyder et al., 2020; Wiegmann et al., 2020b). In addition to their usage as classification inputs, embeddings can also be used in other ways, such as key- or descriptive word expansion (Viegas et al., 2019; Qiang et al., 2019), clustering (Hadifar et al., 2019; Comito et al., 2019), queries, or summarization (Singh and Shashi, 2019).

Kruspe et al. (2019) propose a system that does not assume an explicit notion of relatedness vs. unrelatedness (or relevance vs. irrelevance) to a crisis event. As described above, these qualities are not easy to define, and might vary for different users or different types of events. The presented method is able to determine whether a tweet belongs to a class (i.e. a crisis event or a desirable topic in a certain use case) implicitly defined by a small selection of example tweets by employing few-shot models. The approach is evaluated in more detail in (Kruspe, 2019).

In the broader picture of detecting actionable information, a trade-off between the flexibility of automatic data stream analysis methods and the available expertise and resources is required. Even though analysis on the tweet level may be fast and can be automated, this approach is quite restrictive because contextual information in terms of semantically similar message contents as well as developments over time and location are not taken into account. As a consequence, parallel events and discussions are difficult to distinguish at this stage. We therefore propose to split the task of identifying actionable information with a specific thematic focus into two steps: (1) Data stream overload reduction with a general, potentially automated and pre-trained model for classifying disaster- or incident-related tweets, and (2) applying (one or even more subsequent) methods that allow for tailored contextual, semantic and/or interactive analyses of the filtered results. This type of approach has, for example, been investigated in (Alam et al., 2020; Kersten and Klan, 2020), and is intended to offer a modular and flexible set of well-

understood methods addressing user-specific sub-tasks, and to provide insights on different granularity levels. Compared to an end-to-end ("black box") approach comprising multiple tasks, modularity helps to keep the complexity low for each sub-task. Furthermore, this workflow supports process interpretability and offers the ability to transparently fuse, combine, or jointly interpret the results from each actionability sub-task.

Methods suitable for in-depth analyses of pre-filtered (i.e. crisis-related) tweets can be grouped into supervised, unsupervised, and hybrid ones. One straightforward approach is the tweet-wise classification into information classes described earlier. The aforementioned data sets *CrisisLexT26*, *CrisisMMD*, *TREC-IS 2019B*, and *SSM* provide example tweets for such classes or ontologies, which were defined in cooperation with emergency managers or agencies. As suggested in (McCreadie et al., 2020), a specific (but not necessarily fixed) subset of information classes can then be analyzed more closely as they represent actionable topics, like "Request-SearchAndRescue" or "Report-EmergingThreats". Additionally, a tweet-wise ranking according to a priority level (e.g. "low", "medium" and "high"), either through classification or through regression, is useful for information prioritization. Ranking tweets via deep learning-based and handcrafted features describing the quality of content (Ibtihel et al., 2019) in order to find fact-checkable messages (Barnwal et al., 2019) or informative content based on multi-modal analyses (Nalluru et al., 2019) are further promising options.

However, tweet-wise analyses alone do not exploit the full potential offered by the Twitter data stream. Important aspects, like aggregating messages, assessing the credibility or geolocation accuracy of a single message/information, and understanding the "big picture" of a situation can significantly be supported by integrating context. In particular, the utilization of unsupervised methods enables a flexible capturing of unforeseen events, discussions, developments, and situations that indicate the need for action.

Identifying significant increases of "bursty keywords" might be a first option for detecting events, like earthquakes (Poblete et al., 2018), but this approach alone tends to produce quite noisy results (Ramachandran and Ramasubramanian, 2018). Topic modeling techniques, like Non-negative Matrix Factorization (NMF) and Latent Dirichlet Allocation (LDA), are commonly used to identify discussed topics (e.g. (Casalino et al., 2018) and (Resch et al., 2018)). Furthermore, clustering techniques that utilize spatial (Ester et al., 1996), spatio-temporal (Birant and Kut, 2007; Lee et al., 2017), and content-based features (Mendonça et al., 2019; Comito et al., 2019; Singh and Shashi, 2019; Fedoryszak et al., 2019) as well as combinations of these (Nguyen and Shin, 2017; Zhang and Eick, 2019) are available. A quite interesting and effective approach lies in directly using word or sentence-embeddings to semantically cluster tweets for various tasks, like the detection of topics (de Miranda et al., 2020), events (Ertugrul et al., 2017), or novelty during crises (Kruspe, 2020). A further promising direction is the combination of pre-trained models and unsupervised methods like the aforementioned clustering. In (Bongard, 2020; Kersten et al., 2021), for example, an unsupervised grouping of incoming tweets helps to keep track of all discussed topics. A simple list of keywords or hashtags together with pre-trained models then support the automated identification of topic-specific, crisis-related, or actionable clusters. An exemplary result based on the *Events2012* data set is depicted in figure 3.

The methodological improvements mentioned above may still not be sufficient for real-world scenarios. Limited personal or computational resources and expert domain knowledge paired with time pressure and data uncertainty motivate the integration of machine learning methods into "systems" that allow to better interact, adjust, summarize, and visualize data analysis results.




Figure 3. Top: 2D visualization of clusters containing the keyword “Long Island” identified on October 14, 2012 (arbitrary dimensions). Bottom: Tweet counts over time (GMT) per cluster. Source: Bongard (2020)

In this regard, McCreadie et al. (2016) propose an Emergency Analysis Identification and Management System (EAIMS) to enable civil protection agencies to easily make use of social media. The system comprises a crawler, service, and user interface layer and enables real-time detection of emergency events, related information finding, and credibility analysis. Furthermore, machine learning is utilized over data gathered from past disasters to build effective models for identifying new events, tracking developments within those events, and analyzing those developments to enhance the decision-making processes of emergency response agencies. The recently proposed decision support system Event Tracker (Thomas et al., 2019) aims at providing a unified view of an event, integrating information from news sources, emergency response officers, social media, and volunteers.

5 Challenges

None of the approaches presented are able to solve the problem of detecting tweets in disaster events perfectly. In some respects, this is due to technical limitations; however, there are several difficulties immanent to the task itself, which we will discuss in this section.

Ambiguous problem definition As stated throughout the paper, the task of tweet detection in disasters is ill-defined and heavily dependent on the use case. Annotation experiments also show that even if the goal is clearly stated, inter-rater agreement is commonly low, with raters often interpreting both the problem statement as well as tweet content very

differently (Stowe et al., 2018). This problem becomes even more emphasized when annotating more fine-grained labels, e.g. for content type classes or for priority. Current research suggests a shift from the target of situational awareness to user-specific actionability.

480 **Linguistic difficulties and language variety** As mentioned above, most data sets and, accordingly, methods for automatic tweet detection focus on English-language data. This would often not be the best choice in a real-world scenario; multi-lingual methods are necessary.

Apart from the question of the language itself, Twitter users frequently utilize an highly idiosyncratic style of writing. Due to the character limitation, words are often abbreviated and grammar is compressed. In contrast to e.g. newspaper articles, user-generated content is relatively noisy, containing lots of erroneous or specialized spelling variations.
485 Additionally, interpretation of tweet content frequently requires (cultural) context knowledge.

Data limitations, legal and privacy issues As mentioned above, Twitter is one of the few popular social media platforms providing an access API to its data to outside users. Despite this, however, limitations exist. For non-paying users, only 1% of the live data of each second can be collected automatically via Twitter’s streaming API. For past events, the search API can be utilized, but this only returns tweets still in the search index, which is usually valid for around one week.
490 Older tweets can be retrieved by their ID, but this does not allow for a flexible search. As a free user, the download rate is limited to 18,000 tweets per 15 minutes. Twitter also offers several paid options (called “firehoses”) to access more live data, but these are somewhat intransparent. An in-depth analysis of the effect that these limitations can have on research is given in (Valkanas et al., 2014).

Twitter also forbids direct redistribution of tweet content, meaning that the described data sets are only available as lists of tweet IDs. This introduces two difficulties: One, retrieving the actual tweet content (“hydrating”) can take a very long time for large data sets due to the rate limit. Two, tweets may become unavailable over time because their creator deleted them or their whole account, or because they were banned. In some cases of older data sets, this means that a significant portion of the corpus cannot be used anymore, impeding reproducibility and comparability of published research.

500 Apart from access limitations, Twitter and legal restrictions also regulate what researchers are allowed to do with this data. As an example, the Twitter user agreement states (Twitter, Inc., 2020):

”Unless explicitly approved otherwise by Twitter in writing, you may not use, or knowingly display, distribute, or otherwise make Twitter Content, or information derived from Twitter Content, available to any entity for the purpose of: (a) conducting or providing surveillance or gathering intelligence, including but not limited to investigating or tracking Twitter users or Twitter Content; (b) conducting or providing analysis or research
505 for any unlawful or discriminatory purpose, or in a manner that would be inconsistent with Twitter users’ reasonable expectations of privacy; (c) monitoring sensitive events (including but not limited to protests, rallies, or community organizing meetings); or (d) targeting, segmenting, or profiling individuals based on sensitive personal information, including their health (e.g., pregnancy), negative financial status or condition, political affiliation or beliefs, racial or ethnic origin, religious or philosophical affiliation or beliefs, sex life

510 or sexual orientation, trade union membership, Twitter Content relating to any alleged or actual commission
of a crime, or any other sensitive categories of personal information prohibited by law.”

Many interesting research questions are not identical, but related to problematic usages described in this statement, e.g. inference on a user basis or monitoring of protests. Researchers must therefore be careful not to step into prohibited territory.

515

Lack of geolocation In a disaster context, knowing exactly where a tweet was sent is often crucial to the usability of this information. Twitter provides several ways of detecting geolocation. The most precise of them is the option for users to send their coordinates along with the tweet. However, only about 1% of tweets contain this information (Sloan et al., 2013). A tweet’s location can also be estimated from the location stated in the user profile, or by analyzing the tweet’s
520 content with regards to mention of geolocation. For operationalization, a geocoding to coordinates is then required, which can be provided by services such as Google Maps or OpenStreetMap’s Nominatim. Unfortunately, these geolocations are prone to errors, e.g. because a user mentions a position other than their own, because they might be traveling, or because the center coordinates of a city are too imprecise to be usable. Geocoding, i.e. the prediction of tweet locations from other sources such as the text content, is also an active area of research (e.g. (Qazi et al., 2020; Brouwer et al., 2017)).

525 **6 Related tasks**

Once tweets related to a disaster event have been discovered, many further analysis steps are possible. We will only touch upon those briefly here. As described in section 3, some of the available data sets have already been annotated with these additional concepts.

A popular next step that many automatic approaches already include is the classification into semantic or information type
530 classes. Such classes may include sentiments, affected people seeking various types of assistance, media reports, warnings and advice etc. No common set of such classes exists; in the *CrisisNLP* and *CrisisLexT26* corpora, 9 and 7 classes are used respectively with some overlap. For the TREC Incident Streams challenge, potential end users were questioned about their classes of interests, resulting in a two-tier ontology with 25 classes on the lower tier. As an added difficulty, classes often overlap in tweets; for these reasons, TREC allows multiple labels per tweet. Furthermore, annotators often disagree whether
535 an information type is present in a tweet.

Another way of further discerning between tweets is a distinction between levels of informativeness or priority. This can be implemented either with discrete classes (low/medium/high importance), on a continuous numerical scale, or as a ranking of tweets. The *CrisisLexT26* and *TREC-IS 2019A* data sets contain such annotations.

Apart from approaches processing single tweets, the analysis of the spatio-temporal distribution and development of discussed
540 topics within affected areas at different scales may provide valuable insights (Kersten and Klan, 2020). Other research focuses on the detection of specific events, or types of events (e.g. floods, wildfires, or man-made disasters) (e.g. Burel et al. (2017b)).

This can often be helpful when social media is used as an alerting system. Additionally, models specialized to event types can be more precise and allow for different distinctions than general-purpose models (Kersten et al., 2019; Wiegmann et al., 2020b); detection of the event type enables the automatic selection of such a more specialized method.

545 Apart from these text-based tasks, image analysis can also be a helpful source of information. As an example, images posted on social media can be used to determine the degree of destruction in the aftermath of a disaster (Alam et al., 2017; Nguyen et al., 2017b).

As suggested in section 4, taking a larger variety of semantic concepts into account could lead to a possible solution of the problem of automatic actionability detection. These concepts can be combined in intelligent and adaptable ways to zone in on
550 what exactly are relevant tweets to a user.

7 Future work

Many very interesting new analysis tasks are thinkable based on the detection methods described so far, particularly when employing automatic methods. A good starting point to identify relevant practical issues related to acquisition tasks that could potentially be solved by analyzing social media data is provided in (Wiegmann et al., 2020c). Here, opportunities and risks of
555 disaster data from social media are investigated by means of a systematic review of currently available incident information.

One aspect that has not been considered in research so far is how an event changes over time. New approaches could be used to analyze the spatiotemporal development of disasters, and how this could be utilized in disaster prevention. During the course of an event, clustering methods could be employed to rapidly detect novel developments such as sub-events or new topics. This is particularly relevant for relief providers, who require extremely fast situation monitoring.

560 As described in section 5, localizing information coming from Twitter is often a challenge. Approaches that are able to deal with this lack of information are necessary. This could be implemented either by deriving location by some other means, or by spatiotemporal and semantic analysis of large sets of tweets to cross-reference and check information.

As mentioned above, languages other than English have also usually not been included in research on this topic. Multilingual approaches would be a very helpful next step to facilitate usage of such methods in regions of the world where English is not
565 the main language. Another aspect of the data that has not been used often so far are images posted by users. In particular, a multimodal joint analysis of text and images is very interesting from both the research as well as the usage perspective. The *CrisisMMD* data set is an interesting first step in this direction.

As described in section 4, some crowdsourcing approaches already integrate machine learning-based methods. In future work, expanding human-in-the-loop approaches would be very useful.

570 In general, social media is usually not the only source of information and cannot provide a full picture of the situation. Therefore, an integration with other information sources, such as earth observation data, media information, or governmental data, is highly relevant. de Bruijn et al. (2020) present a first foray by combining social media information with hydrological data.

As described, a large step towards making automatic tweet detection approaches more useful in real-life systems lies in their adaptability to the desired use cases. We have identified three promising research directions in this paper:

- 575 1. Exploiting various concepts and other analysis methods suggested in this section to allow users to flexibly define actionability and detect tweets based on this definition.
2. Machine learning models that can adapt to new use cases, e.g. through active learning or few-shot modeling with the involvement of users, through domain adaptation, or through novelty detection.
- 580 3. Complex systems that integrate automatic tweet analysis with available expertise and other resources, e.g. by combining an automatic pre-filtering step with dedicated methods for a specific actionability scenario (either manual or also automatic).

8 Conclusions

In this review paper, we gave an overview over current methods to detect tweets pertaining to disaster events. As a major hindrance, we identified the necessity for an exact definition of the desired tweets. Conventionally, automatic recognition of tweets aims to achieve a generalized situational awareness, utilizing the ill-defined concepts of “relatedness”, “informativeness”, or “relevance”. In real-world scenarios, however, the question which tweets should be detected depends on the use case, and has been framed as the concept of actionability in recent research. Most data sets and applications do not yet offer this flexibility. We compare various crisis tweet data sets available online. Unfortunately, these usually only provide ID’s of the tweets, which leads to changes in the data sets over time. In addition, labels are usually only provided for the described static binary concepts (related/informative/relevant), and definitions do not match across data sets. Nevertheless, these collections are a very useful basis for analyzing user behavior and for developing new models. They also frequently offer annotations for other concepts, such as information types or sources. We believe integrating these concepts in future approaches could lead to more flexibility in the domain of actionability.

585

590

On the methodical side, there are three main ways to approach the problem: Filtering tweets by characteristics such as location and contained keywords or hashtags, crowdsourcing, and machine-learning based methods. Each of these has its advantages and disadvantages, but machine learning appears to be the current main avenue of research with big improvements in the past few years. Once again, most methods from the past few years follow a static ontology, but there is now a development towards novel approaches that allow for flexible adaptation to user-based actionability definitions, e.g. via few-shot learning based on a small number of example tweets or by detecting specific topical clusters of tweets.

595

Besides the definition problem, other difficulties include the subjectivity of classes and tweet interpretations, data limitations, linguistic difficulties, and legal issues. Nevertheless, large strides have been made in the past years to tackle this problem, and research in this area remains highly active. Many related and novel analysis tasks are possible in the future, for instance, to support impact assessment during long-term crises such as droughts but also during humanitarian conflicts. For further reading on the topic of crisis informatics as a whole, we recommend the bibliography provided in (Palen et al., 2020).

600

605 *Author contributions.* Anna Kruspe wrote this paper with assistance and input from Jens Kersten and Friederike Klan. Jens Kersten contributed ideas and text to sections 2, 3, 4.3, and 4.4.

Competing interests. The authors declare that they have no conflict of interest.

Disclaimer. TEXT

Acknowledgements. TEXT

610 References

- Alam, F., Imran, M., and Ofli, F.: Image4Act: Online Social Media Image Processing for Disaster Response, in: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, ASONAM '17, p. 601–604, 2017.
- Alam, F., Joty, S., and Imran, M.: Domain Adaptation with Adversarial Training and Graph Embeddings, in: 56th Annual Meeting of the Association for Computational Linguistics (ACL), Melbourne, Australia, 2018a.
- 615 Alam, F., Ofli, F., and Imran, M.: CrisisMMD: Multimodal Twitter Datasets from Natural Disasters, in: Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM), 2018b.
- Alam, F., Ofli, F., and Imran, M.: Descriptive and visual summaries of disaster events using artificial intelligence techniques: case studies of Hurricanes Harvey, Irma, and Maria, *Behaviour & Information Technology*, 39, 288–318, <https://doi.org/10.1080/0144929X.2019.1610908>, <https://doi.org/10.1080/0144929X.2019.1610908>, 2020.
- 620 ALRashdi, R. and O’Keefe, S.: Deep Learning and Word Embeddings for Tweet Classification for Crisis Response, 2019.
- American Red Cross: Social Media in Disasters and Emergencies, Tech. rep., 2010.
- Appen Ltd.: Multilingual Disaster Response Messages, <https://appen.com/datasets/combined-disaster-response-data/>, 2020.
- Barnwal, D., Ghelani, S., Krishna, R., Basu, M., and Ghosh, S.: Identifying Fact-Checkable Microblogs during Disasters: A Classification-Ranking Approach, *ICDCN '19*, p. 389–392, Association for Computing Machinery, New York, NY, USA, 625 <https://doi.org/10.1145/3288599.3295587>, <https://doi.org/10.1145/3288599.3295587>, 2019.
- Birant, D. and Kut, A.: ST-DBSCAN: An algorithm for clustering spatial-temporal data, *Data and Knowledge Engineering*, 60, 208–221, *intelligent Data Mining*, 2007.
- Blei, D. M., Ng, A. Y., and Jordan, M. I.: Latent dirichlet allocation, *J. Mach. Learn. Res.*, 3, 993–1022, 2003.
- Bongard, J. H.: Twitter Stream Clustering for the identification and contextualization of event related Tweets, Master’s thesis, Friedrich 630 Schiller Universität Jena, Institute of Geography, Grietgasse 6, 07743 Jena, 2020.
- Brouwer, T., Eilander, D., van Loenen, A., Booij, M. J., Wijnberg, K. M., Verkade, J. S., and Wagemaker, J.: Probabilistic flood extent estimates from social media flood observations, *Natural Hazards and Earth System Sciences*, 17, 735–747, <https://doi.org/10.5194/nhess-17-735-2017>, 2017.
- Burel, G. and Alani, H.: Crisis Event Extraction Service (CREES) - Automatic Detection and Classification of Crisis-related Content on 635 Social Media, in: Proceedings of the 15th International Conference on Information Systems for Crisis Response and Management (IS-CRAM), p. 12, 2018.
- Burel, G., Saif, H., and Alani, H.: Semantic Wide and Deep Learning for Detecting Crisis-Information Categories on Social Media, in: *International Semantic Web Conference (ISWC)*, Vienna, Austria, 2017a.
- Burel, G., Saif, H., Fernandez, M., and Alani, H.: On semantics and deep learning for event detection in crisis situations, in: *Workshop on 640 Semantic Deep Learning (SemDeep)*, at *ESWC 2017*, 2017b.
- Caragea, C., Silvescu, A., and Tapia, A.: Identifying informative messages in disaster events using Convolutional Neural Networks, in: 13th *International Conference on Information Systems for Crisis Response and Management*, 2016.
- Casalino, G., Castiello, C., Buono, N. D., and Mencar, C.: A framework for intelligent Twitter data analysis with non-negative matrix factorization, *IJWIS*, 14, 334–356, 2018.
- 645 Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N. L. U., John, R. S., Constant, N., Guajardo-Céspedes, M., Yuan, S., Tar, C., Sung, Y., Strope, B., and Kurzweil, R.: Universal Sentence Encoder, *arXiv:1803.11175*, 2018.

- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P.: Natural Language Processing (Almost) from Scratch, *J. Mach. Learn. Res.*, 999888, 2493–2537, 2011.
- Comito, C., Forestiero, A., and Pizzuti, C.: Word Embedding Based Clustering to Detect Topics in Social Media, in: *IEEE/WIC/ACM International Conference on Web Intelligence, WI'19*, pp. 192–199, Association for Computing Machinery, New York, NY, USA, 2019.
- 650 Crowdflower: <https://data.world/crowdflower/disasters-on-social-media>, 2015.
- de Albuquerque, J. P., Herfort, B., Brenning, A., and Zipf, A.: A geographic approach for combining social media and authoritative data towards identifying useful information for disaster management, *International Journal of Geographical Information Science*, 29, 667–689, 2015.
- 655 de Bruijn, J. A., de Moel, H., Weerts, A. H., de Ruiter, M. C., Basar, E., Eilander, D., and Aerts, J. C. J. H.: Improving the classification of flood tweets with contextual hydrological information in a multimodal neural network, *Comput. Geosci.*, 140, 104 485, 2020.
- de Miranda, G. R., Pasti, R., and de Castro, L. N.: Detecting Topics in Documents by Clustering Word Vectors, in: *Distributed Computing and Artificial Intelligence, 16th International Conference*, edited by Herrera, F., Matsui, K., and Rodríguez-González, S., pp. 235–243, Springer International Publishing, Cham, 2020.
- 660 Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019.
- Ertugrul, A. M., Velioglu, B., and Karagoz, P.: Word Embedding Based Event Detection on Social Media, in: *Hybrid Artificial Intelligent Systems*, edited by Martínez de Pisón, F. J., Urraca, R., Quintián, H., and Corchado, E., pp. 3–14, Springer International Publishing, Cham,
- 665 2017.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X.: A Density-based Algorithm for Discovering Clusters a Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise, in: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96*, pp. 226–231, AAAI Press, 1996.
- Fathi, R., Thom, D., Koch, S., Ertl, T., and Fiedrich, F.: VOST: A case study in voluntary digital participation for collaborative emergency
- 670 management, *Information Processing & Management*, 57, 102 174, 2020.
- Fedoryszak, M., Frederick, B., Rajaram, V., and Zhong, C.: Real-Time Event Detection on Social Data Streams, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, pp. 2774–2782, Association for Computing Machinery, New York, NY, USA, <https://doi.org/10.1145/3292500.3330689>, 2019.
- Grace, R.: Crisis social media data labeled for storm-related information and toponym usage, *Data in Brief*, 30, 2020.
- 675 Habdank, M., Rodehutsors, N., and Koch, R.: Relevancy assessment of tweets using supervised learning techniques: Mining emergency related tweets for automated relevancy classification, 2017 4th ICT-DM, 2017.
- Hadifar, A., Sterckx, L., Demeester, T., and Develder, C.: A Self-Training Approach for Short Text Clustering, in: *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pp. 194–199, Association for Computational Linguistics, Florence, Italy, <https://doi.org/10.18653/v1/W19-4322>, <https://www.aclweb.org/anthology/W19-4322>, 2019.
- 680 Havas, C., Resch, B., Francalanci, C., Pernici, B., Scalia, G., Fernandez-Marquez, J. L., Van Achte, T., Zeug, G., Mondardini, M. R. R., Grandoni, D., Kirsch, B., Kalas, M., Lorini, V., and Rüping, S.: E2mC: Improving Emergency Management Service Practice through Social Media and Crowdsourcing Analysis in Near Real Time, *Sensors*, 17, 2017.

- He, X., Lu, D., Margolin, D., Wang, M., Idrissi, S. E., and Lin, Y.-R.: The Signals and Noise: Actionable Information in Improvised Social Media Channels During a Disaster, in: Proceedings of the 2017 ACM on Web Science Conference, WebSci '17, p. 33–42, Association for Computing Machinery, New York, NY, USA, 2017.
- 685 Ibtihel, B. L., Lobna, H., and Lotfi, B. R.: A Deep Learning-based Ranking Approach for Microblog Retrieval, *Procedia Computer Science*, 159, 352–362, <https://doi.org/https://doi.org/10.1016/j.procs.2019.09.190>, <https://www.sciencedirect.com/science/article/pii/S1877050919313729>, *knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 23rd International Conference KES2019*, 2019.
- 690 Imran, M., Elbassuoni, S., Castillo, C., Diaz, F., and Meier, P.: Practical extraction of disaster-relevant information from social media, in: Proceedings of the 22nd International Conference on World Wide Web (WWW) Companion, pp. 1021–1024, ACM Press, Rio de Janeiro, Brazil, 2013.
- Imran, M., Castillo, C., Lucas, J., Meier, P., and Vieweg, S.: AIDR: Artificial Intelligence for Disaster Response, in: Proceedings of the 23rd International Conference on World Wide Web, WWW '14 Companion, p. 159–162, Association for Computing Machinery, New York, NY, USA, 2014.
- 695 Imran, M., Castillo, C., Diaz, F., and Vieweg, S.: Processing Social Media Messages in Mass Emergency: A Survey, *ACM Computing Surveys*, 47, 1–38, <https://doi.org/10.1145/2771588>, <http://dl.acm.org/citation.cfm?doid=2775083.2771588>, 2015.
- Imran, M., Mitra, P., and Castillo, C.: Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages, in: Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC), European Language Resources Association (ELRA), Paris, France, 2016a.
- 700 Imran, M., Mitra, P., and Castillo, C.: Twitter as a lifeline: Human-annotated twitter corpora for NLP of crisis-related messages, in: Tenth International Conference on Language Resources and Evaluation (LREC), Portoroz, Slovenia, 2016b.
- Imran, M., Mitra, P., and Srivastava, J.: Enabling Rapid Classification of Social Media Communications During Crises, *Int. J. Inf. Syst. Crisis Response Manag.*, 8, 1–17, 2016c.
- 705 Imran, M., Castillo, C., Diaz, F., and Vieweg, S.: Processing Social Media Messages in Mass Emergency: Survey Summary, in: Companion Proceedings of the The Web Conference 2018, WWW '18, pp. 507–511, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 2018.
- Joulin, A., Grave, E., Bojanowski, P., and Mikolov, T.: Bag of Tricks for Efficient Text Classification, 2016.
- Kaufhold, M.-A., Bayer, M., and Reuter, C.: Rapid relevance classification of social media posts in disasters and emergencies: A system and evaluation featuring active, incremental and online learning, *Information Processing & Management*, 57, 2020.
- 710 Kejriwal, M. and Zhou, P.: Low-supervision Urgency Detection and Transfer in Short Crisis Messages, arXiv:1907.06745, 2019.
- Kersten, J. and Klan, F.: What Happens Where During Disasters? A Workflow for the Multi-Faceted Characterisation of Crisis Events Based on Twitter Data, *Journal of Contingencies and Crisis Management*, special Issue: Knowledge, Semantics and AI for Risk and Crisis Management” of the *Journal of Contingencies and Crisis Management*, in press, 2020.
- 715 Kersten, J., Kruspe, A., Wiegmann, M., and Klan, F.: Robust Filtering of Crisis-related Tweets, in: International Conference on Information Systems for Crisis Response and Management (ISCRAM), Valencia, Spain, 2019.
- Kersten, J., Bongard, J. H., and Klan, F.: Combining Supervised and Unsupervised Learning to Detect and Semantically Aggregate Crisis-Related Twitter Content, in: International Conference on Information Systems for Crisis Response and Management (ISCRAM), under review, 2021.

- 720 Kim, Y.: Convolutional neural networks for sentence classification, in: 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 2014.
- Kropczynski, J., Grace, R., Coche, J., Halse, S., Obeysekare, E., Montarnal, A., Benaben, F., and Tapia, A.: Identifying Actionable Information on Social Media for Emergency Dispatch, in: ISCRAM Asia Pacific 2018: Innovating for Resilience – 1st International Conference on Information Systems for Crisis Response and Management Asia Pacific., pp. p.428–438, Wellington, New Zealand, 2018.
- 725 Kruspe, A.: Few-shot tweet detection in emerging disaster events, in: AI+HADR Workshop @ NeurIPS, 2019.
- Kruspe, A.: Detecting novelty in social media messages during emerging crisis events, in: International Conference on Information Systems for Crisis Response and Management (ISCRAM), 2020.
- Kruspe, A., Kersten, J., and Klan, F.: Detecting Event-Related Tweets by Example using Few-Shot Models, in: International Conference on Information Systems for Crisis Response and Management (ISCRAM), Valencia, Spain, 2019.
- 730 Kruspe, A., Häberle, M., Kuhn, I., and Zhu, X. X.: Cross-language sentiment analysis of European Twitter messages during the COVID-19 pandemic, in: Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020, Association for Computational Linguistics, 2020.
- Kumar, S., Barbier, G., Abbasi, M. A., and Liu, H.: TweetTracker: An Analysis Tool for Humanitarian and Disaster Relief, in: International AAI Conference on Weblogs and Social Media (ICWSM), Barcelona, Spain, 2011.
- Lample, G., Conneau, A., Ranzato, M., Denoyer, L., and Jégou, H.: Word translation without parallel data, in: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings, 2018.
- 735 Landwehr, P. M. and Carley, K. M.: Social Media in Disaster Relief - Usage Patterns, Data Mining Tools, and Current Research Directions, in: Data Mining and Knowledge Discovery for Big Data, edited by Chu, W. W., pp. 225–257, Springer Berlin Heidelberg, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-40837-3_7, 2014.
- Lang, S., Füreder, P., Riedler, B., Wendt, L., Braun, A., Tiede, D., Schoepfer, E., Zeil, P., Spröhnle, K., Kulesa, K., Rogenhofer, E., Bäuerl, M., Öze, A., Schwendemann, G., and Hochschild, V.: Earth observation tools and services to increase the effectiveness of humanitarian assistance, *European Journal of Remote Sensing*, 53, 67–85, 2020.
- 740 Lee, J., Gong, J., and Li, S.: Exploring spatiotemporal clusters based on extended kernel estimation methods, *International Journal of Geographical Information Science*, 31, 1154–1177, 2017.
- Li, H., Caragea, D., Caragea, C., and Herndon, N.: Disaster response aided by tweet classification with a domain adaptation approach, *Journal of Contingencies and Crisis Management*, 26, 16–27, 2018.
- 745 Lin, Z., Jin, H., Robinson, B., and Lin, X.: Towards an accurate social media disaster event detection system based on deep learning and semantic representation, in: Proceedings of the 14th Australasian Data Mining Conference, Canberra, Australia, pp. 6–8, 2016.
- Liu, J., Singhal, T., Blessing, L. T. M., Wood, K. L., and Lim, K. H.: CrisisBERT: a Robust Transformer for Crisis Classification and Contextual Crisis Embedding, 2020.
- 750 Lorini, V., Castillo, C., Dottori, F., Kalas, M., Nappo, D., and Salamon, P.: Integrating Social Media into a Pan-European Flood Awareness System: A Multilingual Approach, in: 16th International Conference on Information Systems for Crisis Response and Management (ISCRAM), València, Spain, 2019.
- Mazloom, R., Li, H., Caragea, D., Caragea, C., and Imran, M.: A Hybrid Domain Adaptation Approach for Identifying Crisis-Relevant Tweets, *International Journal of Information Systems for Crisis Response and Management*, 11, 2019.
- 755 McCreddie, R., Macdonald, C., and Ounis, I.: EAIMS: Emergency Analysis Identification and Management System, in: SIGIR, pp. 1101–1104, ACM, 2016.

- McCreadie, R., Buntain, C., and Soboroff, I.: TREC Incident Streams: Finding Actionable Information on Social Media, in: International Conference on Information Systems for Crisis Response and Management (ISCRAM), Valencia, Spain, 2019.
- 760 McCreadie, R., Buntain, C., and Soboroff, I.: Incident Streams 2019: Actionable Insights and How to Find Them, in: 17th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2020), pp. 744–760, 2020.
- McMinn, A. J., Moshfeghi, Y., and Jose, J. M.: Building a large-scale corpus for evaluating event detection on twitter, in: ACM International Conference on Information and Knowledge Management (CIKM), pp. 409–418, San Francisco, CA, USA, 2013.
- Mendonça, I., Trouvé, A., Fukuda, A., Murakami, K., Tsai, C.-F., Hu, Y.-H., Wang, M.-C., Liu, K. E., Yu, X., Yuan, Y., et al.: On Clustering Algorithms: Applications in Word-Embedding Documents., JCP, 14, 88–92, 2019.
- 765 Mikolov, T., Chen, K., Corrado, G., and Dean, J.: Efficient Estimation of Word Representations in Vector Space, CoRR, abs/1301.3781, <http://arxiv.org/abs/1301.3781>, 2013.
- Nalluru, G., Pandey, R., and Purohit, H.: Relevancy Classification of Multimodal Social Media Streams for Emergency Services, IEEE International Conference on Smart Computing, <https://par.nsf.gov/biblio/10098177>, 2019.
- 770 Nguyen, D. T., Joty, S. R., Imran, M., Sajjad, H., and Mitra, P.: Applications of Online Deep Learning for Crisis Response Using Social Media Information, in: International Workshop on Social Web for Disaster Management (SWDM), Los Angeles, CA, USA, 2016a.
- Nguyen, D. T., Joty, S. R., Imran, M., Sajjad, H., and Mitra, P.: Applications of Online Deep Learning for Crisis Response Using Social Media Information, CoRR, abs/1610.01030, 2016b.
- 775 Nguyen, D. T., Al-Mannai, K. A., Joty, S., Sajjad, H., Imran, M., and Mitra, P.: Robust classification of crisis-related data on social networks using convolutional neural networks, in: Proceedings of the 11th International Conference on Web and Social Media (ICWSM), pp. 632–635, AAAI press, 2017a.
- Nguyen, D. T., Ofli, F., Imran, M., and Mitra, P.: Damage Assessment from Social Media Imagery Data During Disasters, in: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, ASONAM '17, p. 569–576, Association for Computing Machinery, New York, NY, USA, 2017b.
- 780 Nguyen, M. D. and Shin, W.-Y.: DBSTexC: Density-Based Spatio-Textual Clustering on Twitter, in: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, ASONAM '17, pp. 23–26, Association for Computing Machinery, New York, NY, USA, <https://doi.org/10.1145/3110025.3110096>, <https://doi.org/10.1145/3110025.3110096>, 2017.
- Niles, M. T., Emery, B. F., Reagan, A. J., Dodds, P. S., and Danforth, C. M.: Social media usage patterns during natural hazards, PLOS ONE, 14, 1–16, 2019.
- 785 Ning, X., Yao, L., Benatallah, B., Zhang, Y., Sheng, Q. Z., and Kanhere, S. S.: Source-Aware Crisis-Relevant Tweet Identification and Key Information Summarization, ACM Trans. Internet Technol., 19, 2019.
- Olteanu, A., Castillo, C., Diaz, F., and Vieweg, S.: CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises, in: AAAI Conference on Weblogs and Social Media (ICWSM), Ann Arbor, MI, USA, 2014.
- Olteanu, A., Vieweg, S., and Castillo, C.: What to Expect When the Unexpected Happens: Social Media Communications Across Crises, in: Conference on Computer Supported Cooperative Work and Social Computing (ACM CSCW), Vancouver, BC, Canada, 2015.
- 790 Palen, L., Anderson, J., Bica, M., Castillos, C., Crowley, J., Díaz, P., Finn, M., Grace, R., Hughes, A., Imran, M., Kogan, M., LaLone, N., Mitra, P., Norris, W., Pine, K., Purohit, H., Reuter, C., Rizza, C., St Denis, L., Semaan, B., Shalin, V., Shanley, L., Shih, P., Soden, R., Starbird, K., Stephen, K., Toups, Z. O., and Wilson, T.: Crisis Informatics: Human-Centered Research on Tech & Crises, <https://hal.archives-ouvertes.fr/hal-02781763>, working paper or preprint, 2020.

- 795 Parilla-Ferrer, B. E., Fernandez, P., and T. Ballena IV, J.: Automatic Classification of Disaster-Related Tweets, in: International conference on Innovative Engineering Technologies (ICIET), 2014.
- Pennington, J., Socher, R., and Manning, C.: GloVe: Global Vectors for Word Representation, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1532–1543, Association for Computational Linguistics, Doha, Qatar, 2014.
- 800 Poblet, M., García-Cuesta, E., and Casanovas, P.: Crowdsourcing Tools for Disaster Management: A Review of Platforms and Methods, in: AI Approaches to the Complexity of Legal Systems, edited by Casanovas, P., Pagallo, U., Palmirani, M., and Sartor, G., pp. 261–274, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- Poblete, B., Guzman, J., Flores, J. A. M., and Tobar, F. A.: Robust Detection of Extreme Events Using Twitter: Worldwide Earthquake Monitoring, *IEEE Trans. Multim.*, 20, 2551–2561, 2018.
- 805 Purohit, H., Castillo, C., Imran, M., and Pandey, R.: Social-EOC: Serviceability Model to Rank Social Media Requests for Emergency Operation Centers, in: IEEE/ACM 2018 International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018, Barcelona, Spain, August 28-31, 2018, edited by Brandes, U., Reddy, C., and Tagarelli, A., pp. 119–126, IEEE Computer Society, 2018.
- Qazi, U., Imran, M., and Ofli, F.: GeoCoV19: A Dataset of Hundreds of Millions of Multilingual COVID-19 Tweets with Location Information, *ACM SIGSPATIAL Special*, 12, 6–15, 2020.
- 810 Qiang, J., Li, Y., Yuan, Y., Liu, W., and Wu, X.: A practical algorithm for solving the sparseness problem of short text clustering, *Intell. Data Anal.*, 23, 701–716, 2019.
- Ramachandran, D. and Ramasubramanian, P.: Event detection from Twitter - a survey, *International Journal of Web Information Systems*, 14, 262–280, <https://doi.org/10.1108/IJWIS-11-2017-0075>, <https://doi.org/10.1108/IJWIS-11-2017-0075>, 2018.
- 815 Resch, B., Usländer, F., and Havas, C.: Combining machine-learning topic models and spatiotemporal analysis of social media data for disaster footprint and damage assessment, *Cartography and Geographic Information Science*, 45, 362–376, 2018.
- Reuter, C. and Kaufhold, M.-A.: Fifteen Years of Social Media in Emergencies: A Retrospective Review and Future Directions for Crisis Informatics, *Journal of Contingencies and Crisis Management (JCCM)*, 26, 41–57, special Issue: Human-Computer-Interaction and Social Media in Safety-Critical Systems, 2018.
- 820 Rogstadius, J., Vukovic, M., Teixeira, C. A., Kostakos, V., Karapanos, E., and Laredo, J. A.: CrisisTracker: Crowdsourced social media curation for disaster awareness, *IBM Journal of Research and Development*, 57, 4:1–4:13, 2013.
- Rossi, C.: I-REACT: Improving Resilience to Emergencies through Advanced Cyber Technologies, Tech. rep., Istituto Superiore Mario Boella (ISMB), 2019.
- Sakaki, T., Okazaki, M., and Matsuo, Y.: Earthquake shakes Twitter users: Real-time event detection by social sensors, in: Proceedings of the 19th International Conference on World Wide Web, pp. 851–860, ACM, 2010.
- 825 Schulz, A. and Guckelsberger, C.: <http://www.doc.gold.ac.uk/~cguck001/IncidentTweets/>, 2016.
- Singh, A. K. and Shashi, M.: Vectorization of Text Documents for Identifying Unifiable News Articles, *International Journal of Advanced Computer Science and Applications*, 10, 2019.
- Sloan, L., Morgan, J., Housley, W., Williams, M., Edwards, A., Burnap, P., and Rana, O.: Knowing the Tweepers: Deriving Sociologically Relevant Demographics from Twitter, *Sociological Research Online*, 18, 7, 2013.
- 830 Snyder, L. S., Lin, Y., Karimzadeh, M., Goldwasser, D., and Ebert, D. S.: Interactive Learning for Identifying Relevant Tweets to Support Real-time Situational Awareness, *IEEE Trans. Vis. Comput. Graph.*, 26, 558–568, 2020.

- Stieglitz, S., Mirbabaie, M., Fromm, J., and Melzer, S.: The Adoption of Social Media Analytics for Crisis Management – Challenges and Opportunities, in: *Twenty-Sixth Eur. Conf. Inf. Syst. (ECIS2018)*, 2018.
- 835 Stowe, K., Paul, M. J., Palmer, M., Palen, L., and Anderson, K.: Identifying and Categorizing Disaster-Related Tweets, in: *Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media*, ACL, Austin, TX, USA, 2016.
- Stowe, K., Palmer, M., Anderson, J., Kogan, M., Palen, L., Anderson, K. M., Morss, R., Demuth, J., and Lazrus, H.: 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, Association for Computational Linguistics, 2018.
- 840 Thomas, C., McCreadie, R., and Ounis, I.: Event Tracker: A Text Analytics Platform for Use During Disasters, in: *SIGIR*, pp. 1341–1344, ACM, 2019.
- To, H., Agrawal, S., Kim, S. H., and Shahabi, C.: On Identifying Disaster-Related Tweets: Matching-Based or Learning-Based?, in: *2017 IEEE Third BigMM*, 2017.
- Twitter, Inc.: Developer Agreement and Policy, <https://developer.twitter.com/en/developer-terms/agreement-and-policy>, 2020.
- 845 Valkanas, G., Katakis, I., Gunopulos, D., and Stefanidis, A.: Mining Twitter Data with Resource Constraints, in: *2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, Warsaw, Poland, August 11-14, 2014 - Volume II, pp. 157–164, IEEE Computer Society, 2014.
- Viegas, F., Canuto, S., Gomes, C., Luiz, W., Rosa, T., Ribas, S., Rocha, L., and Gonçalves, M. A.: CluWords: Exploiting Semantic Word Clustering Representation for Enhanced Topic Modeling, in: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM '19*, pp. 753–761, Association for Computing Machinery, New York, NY, USA, <https://doi.org/10.1145/3289600.3291032>, 2019.
- 850 Voigt, S., Giulio-Tonolo, F., Lyons, J., Kučera, J., Jones, B., Schneiderhan, T., Platzeck, G., Kaku, K., Hazarika, M. K., Czarán, L., Li, S., Pedersen, W., James, G. K., Proy, C., Muthike, D. M., Bequignon, J., and Guha-Sapir, D.: Global trends in satellite-based emergency mapping, *Science*, 353, 247–252, 2016.
- Wiegmann, M., Kersten, J., Klan, F., Potthast, M., and Stein, B.: Disaster Tweet Corpus 2020, <https://doi.org/10.5281/zenodo.3713920>, 2020a.
- 855 Wiegmann, M., Kersten, J., Klan, F., Potthast, M., and Stein, B.: Analysis of Detection Models for Disaster-Related Tweets, in: *17th ISCRAM Conference*, edited by Hughes, A., McNeill, F., and Zobel, C., ISCRAM, 2020b.
- Wiegmann, M., Kersten, J., Potthast, M., Klan, F., and Stein, B.: Opportunities and Risks of Disaster Data from Social Media: A Systematic Review of Incident Information, *NHESS Special Issue: Groundbreaking technologies, big data, and innovation for disaster risk modelling and reduction*, under review, 2020c.
- 860 Win, S. S. M. and Aung, T. N.: Target oriented tweets monitoring system during natural disasters, in: *Proceedings of the IEEE/ACIS 16th International Conference on Computer and Information Science (ICIS)*, pp. 143–148, 2017.
- Xu, Z., Liu, Y., Yen, N. Y., Mei, L., Luo, X., Wei, X., and Hu, C.: Crowdsourcing Based Description of Urban Emergency Events Using Social Media Big Data, *IEEE Transactions on Cloud Computing*, 8, 387–397, 2020.
- 865 Yang, Y., Cer, D., Ahmad, A., Guo, M., Law, J., Constant, N., Abrego, G. H., Yuan, S., Tar, C., Sung, Y.-H., Strope, B., and Kurzweil, R.: Multilingual Universal Sentence Encoder for Semantic Retrieval, *arXiv:1907.04307*, 2019.
- Zade, H., Shah, K., Rangarajan, V., Kshirsagar, P., Imran, M., and Starbird, K.: From Situational Awareness to Actionability: Towards Improving the Utility of Social Media Data for Crisis Response, *Proc. ACM Hum.-Comput. Interact.*, 2, 2018.

- Zhang, Y. and Eick, C. F.: Tracking Events in Twitter by Combining an LDA-Based Approach and a Density-Contour Clustering Approach, 870 International Journal of Semantic Computing, 13, 87–110, <https://doi.org/10.1142/S1793351X19400051>, 2019.
- Zheng, X., Sun, A., Wang, S., and Han, J.: Semi-Supervised Event-related Tweet Identification with Dynamic Keyword Generation, in: Proceedings of the 2017 ACM CIKM, ACM, 2017.