

Comprehensive Wikipedia Monitoring for Global and Realtime Natural Disaster Detection

Thomas Steiner

Google Germany GmbH, Hamburg, Germany and
CNRS, Université de Lyon, LIRIS – UMR5205, Université Lyon 1, France
tsteiner@{liris.cnrs.fr, google.com}

Abstract. Natural disasters are harmful events resulting from natural processes of the Earth. Examples of natural disasters include tsunamis, volcanic eruptions, earthquakes, floods, droughts, and other geologic processes. If they affect populated areas, natural disasters can cause economic damage, injuries, or even losses of lives. It is thus desirable that natural disasters be detected as early as possible and potentially affected persons be notified via emergency alerts. By their pure nature, natural disasters are global phenomena that people refer to by different names, for example, the 2014 typhoon *Rammasun*¹ is known as typhoon *Glenda* in the Philippines. In this paper, we present our ongoing early-stage research on a realtime Wikipedia-based monitoring system for the detection of natural disasters around the globe. The long-term objective is to make data about natural disasters detected by this system available through public alerts following the Common Alerting Protocol (CAP).

Keywords: Natural disaster detection, crisis response, Wikipedia

1 Introduction

1.1 Natural Disaster Detection and Response: A Global Challenge

According to a study [4] published by the International Monetary Funds (IMF) in 2012, about 700 natural disasters were registered worldwide between 2010 and 2012, affecting more than 450 million people. According to the study, “[d]amages have risen from an estimated US\$20 billion on average per year in the 1990s to about US\$100 billion per year during 2000–10.” The authors expect this upward trend to continue “as a result of the rising concentration of people living in areas more exposed to natural disasters, and climate change.” In consequence, public emergency alerting systems become more and more crucial in the future.

National agencies like the *Federal Emergency Management Agency* (FEMA)² in the United States of America or the *Bundesamt für Bevölkerungsschutz und*

¹ Rammasun: [http://en.wikipedia.org/wiki/Typhoon_Rammasun_\(2014\)](http://en.wikipedia.org/wiki/Typhoon_Rammasun_(2014))

² FEMA: <http://www.fema.gov/>

Katastrophenhilfe (BBK,³ “Federal Office of Civil Protection and Disaster Assistance”) in Germany work to ensure the safety of the population on a national level, combining and providing relevant tasks and information in a single place. The *United Nations Office for the Coordination of Humanitarian Affairs* (OCHA)⁴ is a United Nations (UN) body formed to strengthen the UN’s response to complex emergencies and natural disasters. The *Global Disaster Alert and Coordination System* (GDACS)⁵ is “a cooperation framework between the United Nations, the European Commission, and disaster managers worldwide to improve alerts, information exchange, and coordination in the first phase after major sudden-onset disasters.” Global companies like Facebook,⁶ Airbnb,⁷ or Google⁸ have dedicated crisis response teams that work on making critical emergency information accessible in times of disaster. As can be seen from the (incomprehensive) list above, natural disaster detection and response is a problem tackled on national, international, and global levels; both from the public and private sectors. To facilitate collaboration, a common protocol is essential.

1.2 The Common Alerting Protocol

The *Common Alerting Protocol* (CAP) [10] is an XML-based general data format for exchanging public warnings and emergencies between alerting technologies. CAP allows a warning message to be consistently disseminated simultaneously over many warning systems to many applications. The protocol increases warning effectiveness and simplifies the task of activating a warning for officials. CAP also provides the capability to include multimedia data, such as photos, maps, or videos. Alerts can be geographically targeted to a defined warning area. An exemplary flood warning CAP feed stemming from GDACS is shown in Listing 1.

1.3 Contributions, Hypotheses, and Research Questions

In this paper, we present first results of our ongoing early-stage research on a realtime comprehensive Wikipedia-based monitoring system for the detection of natural disasters around the globe. We are steered by the following hypotheses.

- Ⓕ1 Content about natural disasters gets added to Wikipedia in a timely fashion.
- Ⓕ2 Natural disasters being geographically constrained, textual and multimedia content about them gets added to local, *i.e.*, non-English Wikipedias as well.
- Ⓕ3 Link structure dynamics of Wikipedia provide for a meaningful way to detect future natural disasters, *i.e.*, disasters unknown at system creation time.

³ BBK: <http://www.bbk.bund.de/>

⁴ OCHA: <http://www.unocha.org/>

⁵ GDACS: <http://www.gdacs.org/>

⁶ Facebook Disaster Relief: <https://www.facebook.com/DisasterRelief>

⁷ Airbnb Disaster Response: <https://www.airbnb.com/disaster-response>

⁸ Google Crisis Response: <https://www.google.org/crisisresponse/>

These hypotheses lead us to the following research questions.

- Q1 How timely and accurate for the purpose of natural disaster detection is content from Wikipedia compared to authoritative sources mentioned above?
- Q2 Does the disambiguated nature of Wikipedia surpass keyword-based natural disaster detection approaches, *e.g.*, via online social networks or search logs?

2 Related Work

Digitally crowdsourced data for disaster detection and response has gained momentum in recent years, as the Internet has proven resilient in times of crises, compared to other infrastructure. Ryan Falor, Crisis Response Product Manager

```

<?xml version="1.0" encoding="utf-8"?>
<alert xmlns="urn:oasis:names:tc:emergency:cap:1.2">
  <identifier>GDACS_FL_4159_1</identifier>
  <sender>info@gdacs.org</sender>
  <sent>2014-07-14T23:59:59-00:00</sent>
  <status>Actual</status>
  <msgType>Alert</msgType>
  <scope>Public</scope>
  <incidents>4159</incidents>
  <info>
    <category>Geo</category><event>Flood</event>
    <urgency>Past</urgency><severity>Moderate</severity>
    <certainty>Unknown</certainty>
    <senderName>Global Disaster Alert and Coordination System</senderName>
    <headline /><description />
    <web>http://www.gdacs.org/reports.aspx?eventtype=FL&amp;amp;eventid=4159</web>
    <parameter><valueName>eventid</valueName><value>4159</value></parameter>
    <parameter><valueName>currentepisodeid</valueName><value>1</value></parameter>
    <parameter><valueName>glide</valueName><value /></parameter>
    <parameter><valueName>version</valueName><value>1</value></parameter>
    <parameter><valueName>fromdate</valueName>
      <value>Wed, 21 May 2014 22:00:00 GMT</value></parameter>
    <parameter><valueName>todate</valueName>
      <value>Mon, 14 Jul 2014 21:59:59 GMT</value></parameter>
    <parameter><valueName>eventtype</valueName><value>FL</value></parameter>
    <parameter><valueName>alertlevel</valueName><value>Green</value></parameter>
    <parameter><valueName>alerttype</valueName><value>automatic</value></parameter>
    <parameter><valueName>link</valueName>
      <value>http://www.gdacs.org/report.aspx?eventtype=FL&amp;amp;eventid=4159</value>
    </parameter>
    <parameter><valueName>country</valueName><value>Brazil</value></parameter>
    <parameter><valueName>eventname</valueName><value /></parameter>
    <parameter><valueName>severity</valueName><value>Magnitude 7.44</value></parameter>
    <parameter><valueName>population</valueName><value>0 killed and 0 displaced</value>
    </parameter>
    <parameter><valueName>vulnerability</valueName><value /></parameter>
    <parameter><valueName>sourceid</valueName><value>DFO</value></parameter>
    <parameter><valueName>iso3</valueName><value /></parameter>
    <parameter><valueName>hazardcomponents</valueName>
      <value>FL,dead=0,displaced=0,main_cause=Heavy Rain,severity=2,sqkm=256564.57
    </value></parameter>
    <parameter><valueName>datemodified</valueName>
      <value>Mon, 01 Jan 0001 00:00:00 GMT</value></parameter>
    <area><areaDesc>Polygon</areaDesc><polygon>,,100</polygon></area>
  </info>
</alert>

```

Listing 1. Common Alerting Protocol feed via the Global Disaster Alert and Coordination System (http://www.gdacs.org/xml/gdacs_cap.xml, 2014-07-16)

at Google in 2011, remarks in [2] that “a substantial [...] proportion of searches are directly related to the crises; and people continue to search and access information online even while traffic and search levels drop temporarily during and immediately following the crises.” In the following, we provide a non-exhaustive list of related work on digitally crowdsourced natural disaster detection and response. Sakaki *et al.* consider in [7] each user of the online social networking (OSN) site Twitter⁹ a sensor for the purpose of earthquake detection in Japan. Goodchild *et al.* show in [3] how crowdsourced geodata from Wikipedia and Wikimapia,¹⁰ “a multilingual open-content collaborative map,” can help complete authoritative data about natural disasters. In [1], Abel *et al.* describe a crisis monitoring system that extracts relevant content about known disasters from Twitter. Liu *et al.* examine in [6] common patterns and norms of natural disaster coverage on the photo sharing site Flickr.¹¹ We have developed [9] a monitoring system that detects news events from concurrent Wikipedia edits and auto-generates related multimedia galleries based on content from various OSN sites and Wikimedia Commons.¹² Finally, Lin and Mishne examine realtime search query churn on Twitter [5] including in the context of natural disasters.

3 Proposed Methodology

3.1 Leveraging Wikipedia Link Structure

Wikipedia is an international online encyclopedia currently available in 287 languages.¹³ (i) Articles in one language are interlinked with versions of the same article in other languages, *e.g.*, the article “Natural disaster” on the English Wikipedia (http://en.wikipedia.org/wiki/Natural_disaster) links to 74 versions of this article in other languages.¹⁴ (ii) Each article can have redirects, *i.e.*, alternative URLs that point to the article. For the English “Natural disaster” article, there are eight redirects,¹⁵ *e.g.*, “Natural Hazard” (synonym), “Examples of natural disaster” (refinement), or “Natural disasters” (plural). (iii) For each article, the list of back links that link to the current article is available, *i.e.*, inbound links other than redirects. The article “Natural disaster” has more than 500 articles that link to it.¹⁶ Likewise, the list of out-

⁹ Twitter: <https://twitter.com/>

¹⁰ Wikimapia: <http://wikimapia.org/>

¹¹ Flickr: <https://www.flickr.com/>

¹² Wikimedia Commons: <https://commons.wikimedia.org/>

¹³ All Wikipedias: http://meta.wikimedia.org/wiki/List_of_Wikipedias

¹⁴ Article language links: http://en.wikipedia.org/w/api.php?action=query&prop=langlinks&limit=max&titles=Natural_disaster

¹⁵ Article redirects: http://en.wikipedia.org/w/api.php?action=query&list=backlinks&blfilter=redirects&bllimit=max&bltitle=Natural_disaster

¹⁶ Article inbound links: http://en.wikipedia.org/w/api.php?action=query&list=backlinks&bllimit=max&blnamespace=0&bltitle=Natural_disaster

bound links, *i.e.*, other articles that the current article links to, is available.¹⁷ By combining an article’s in- and outbound links, we determine the set of mutual links, *i.e.*, the set of articles that the current article links to (outbound links) and at the same time receives links from (inbound links).

3.2 Identification of Wikipedia Articles for Monitoring

Starting with the well-curated English seed article “Natural disaster”, we programmatically follow each of the therein contained links of type “Main article:”, which leads to an exhaustive list of English articles of concrete types of natural disasters, *e.g.*, “Tsunami” (<http://en.wikipedia.org/wiki/Tsunami>), “Flood” (<http://en.wikipedia.org/wiki/Flood>), “Earthquake” (<http://en.wikipedia.org/wiki/Earthquake>), *etc.* In total, we obtain links to 20 English articles about different types of natural disasters.¹⁸ For each of these English natural disasters articles, we obtain all versions of each article in different languages [step (i) above], and of the resulting list of international articles in turn all their redirect URLs [step (ii) above]. The intermediate result is a complete list of all (currently 1,270) articles in all Wikipedia languages and all their redirects that have any type of natural disaster as their subject. We call this list the “natural disasters list” and make it publicly available in different formats (.txt, .tsv, and .json), where the JSON version is the most flexible and recommended one.¹⁹ Finally, we obtain for each of the 1,270 articles in the “natural disasters list” all their back links, *i.e.*, their inbound links [step (iii) above], which serves to detect instances of natural disasters unknown at system creation time. For example, the article “Typhoon Rammasun (2014)” ([http://en.wikipedia.org/wiki/Typhoon_Rammasun_\(2014\)](http://en.wikipedia.org/wiki/Typhoon_Rammasun_(2014)))—which, as a concrete *instance of* a natural disaster of type tropical cyclone, is *not* contained in our “natural disasters list”—links back to “Tropical cyclone” (http://en.wikipedia.org/wiki/Tropical_cyclone), so we can identify “Typhoon Rammasun (2014)” as *related to* tropical cyclones (but not necessarily *identify as* a tropical cyclone), even if at the system’s creation time the typhoon did not exist yet. Analog to the inbound links, we obtain all outbound links of all articles in the “natural disasters list”, *e.g.*, “Tropical cyclone” has an outbound link to “2014 Pacific typhoon season” (http://en.wikipedia.org/wiki/2014_Pacific_typhoon_season), which also

¹⁷ Article outbound links: http://en.wikipedia.org/w/api.php?action=query&prop=links&plnamespace=0&format=json&pllimit=max&titles=Natural_disaster

¹⁸ “Avalanche”, “Blizzard”, “Cyclone”, “Drought”, “Earthquake”, “Epidemic”, “Extratropical cyclone”, “Flood”, “Gamma-ray burst”, “Hail”, “Heat wave”, “Impact event”, “Limnic eruption”, “Meteorological disaster”, “Solar flare”, “Tornado”, “Tropical cyclone”, “Tsunami”, “Volcanic eruption”, “Wildfire”

¹⁹ “Natural disasters list”: <https://github.com/tomayac/postdoc/blob/master/papers/comprehensive-wikipedia-monitoring-for-global-and-realtime-natural-disaster-detection/data/natural-disasters-list.json>

happens to be an inbound link of “Tropical cyclone”, so we have detected a mutual, circular link structure. Figure 1 shows the example in its entirety, starting from the seed level, to the disaster type level, to the in-/outbound link level. The end result is a large list called the “monitoring list” of all articles in all Wikipedia languages that are somehow—via a redirect, inbound, or outbound link (or resulting mutual link)—related to any of the articles in the “natural disasters list”. We make a snapshot of this dynamic “monitoring list” available for reference,²⁰ but note that it will be out-of-date soon and should be regenerated on a regular basis. The current version holds 141,001 different articles.

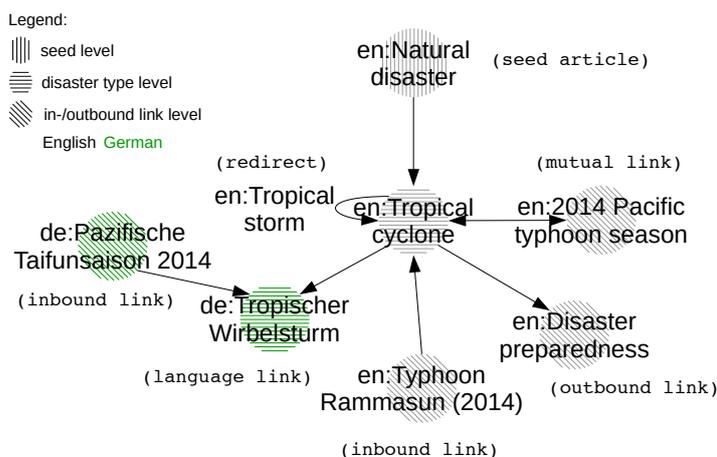


Fig. 1. Extracted Wikipedia link structure starting from seed article “Natural disaster”

3.3 Monitoring Process

In the past, we have worked on a Server-Sent Events (SSE) API [8] capable of monitoring realtime editing activity on all language versions of Wikipedia. This API allows us to easily analyze Wikipedia edits by reacting on events fired by the API. Whenever an edit event occurs, we check if it is for one of the articles on our “monitoring list”. We keep track of the historic one-day-window editing activity for each article on the “monitoring list” including their versions in other languages, and, upon a sudden spike of editing activity, trigger an alert about a potential new instance of a natural disaster type that the spiking article is an inbound or outbound link of (or both). To illustrate this, if, *e.g.*, the German article “Pazifische Taifunsaison 2014” including all of its language links is spiking,

²⁰ “Monitoring list”: <https://github.com/tomayac/postdoc/blob/master/papers/comprehensive-wikipedia-monitoring-for-global-and-realtime-natural-disaster-detection/data/monitoring-list.json>

we can infer that this is related to a natural disaster of type “Tropical cyclone” due to the detected mutual link structure mentioned earlier (Figure 1).

In order to detect spikes, we apply exponential smoothing to the last n edit intervals (we require $n \geq 5$) that occurred in the past 24 hours with a smoothing factor $\alpha = 0.5$. The therefore required edit events are retrieved programmatically via the Wikipedia API.²¹ As a spike occurs when an edit interval gets “short enough” compared to historic editing activity, we report a spike whenever the latest edit interval is shorter than half a standard deviation $0.5 \times \sigma$.

A subset of all Wikipedia articles are geo-referenced,²² so when we detect a spiking article, we try to obtain geo coordinates for the article itself (*e.g.*, “Pazifische Taifunsaison 2014”) or any of its language links that—as a consequence of the assumption in $\mathbb{H}2$ —may provide more local details (*e.g.*, “2014 Pacific typhoon season” in English or “2014年太平洋季” in Chinese). We then calculate the center point of all obtained latitude/longitude pairs.

In a final step, once a given confidence threshold has been reached and upon human inspection, we plan to send out a notification according to the *Common Alerting Protocol* following the format that (for GDACS) can be seen in Listing 1.

3.4 Implementation Details

We have created a publicly available prototypal demo application deployed²³ at <http://disaster-monitor.herokuapp.com/> that internally connects to the SSE API from [8]. It is implemented in Node.js on the server, and as a JavaScript Web application on the client. This application uses an hourly refreshed version of the “monitoring list” from Subsection 3.2 and whenever an edit event sent through the SSE API matches any of the articles in the list, it checks if, given this article’s and its language links’ edit history of the past 24 hours, the current edit event shows spiking behavior, as outlined in Subsection 3.3. The core source code snippet of the main monitoring loop can be seen in Listing 2, a screenshot of the application is shown in Figure 2.

4 Proposed Steps Toward an Evaluation

We recall our core research questions that were Q1 *How timely and accurate for the purpose of natural disaster detection is content from Wikipedia compared to authoritative sources mentioned above?* and Q2 *Does the disambiguated nature of Wikipedia surpass keyword-based natural disaster detection approaches,*

²¹ Wikipedia last revisions: [http://en.wikipedia.org/w/api.php?action=query&prop=revisions&rvlimit=6&rvprop=timestamp|user&titles=Typhoon_Rammasun_\(2014\)](http://en.wikipedia.org/w/api.php?action=query&prop=revisions&rvlimit=6&rvprop=timestamp|user&titles=Typhoon_Rammasun_(2014))

²² Article geo coordinates: http://en.wikipedia.org/w/api.php?action=query&prop=coordinates&format=json&colimit=max&coprop=dim|country|region|globe&coprimary=all&titles=September_11_attacks

²³ Source code: <https://github.com/tomayac/postdoc/tree/master/demos/disaster-monitor>

e.g., *via online social networks or search logs?* Regarding Q1, only a manual comparison covering several months worth of natural disaster data of the relevant authoritative data sources mentioned in Subsection 1.1 with the output of our system can help respond to the question. Regarding Q2, we propose an evaluation strategy for the OSN site Twitter, loosely inspired by the approach of Sakaki *et al.* in [7]. We choose Twitter as a data source due to the publicly

```

var init = function() {

  // fired whenever an edit event happens on any Wikipedia
  var parseWikipediaEdit = function(data) {
    var article = data.language + ':' + data.article;
    var disasterObj = monitoringList[article];
    // the article is on the monitoring list
    if (disasterObj) {
      showCandidateArticle(data.article, data.language, disasterObj);
    }
  };

  // fired whenever an article is on the monitoring list
  var showCandidateArticle = function(article, language, roles) {
    getGeoData(article, language, function(err, geoData) {
      getRevisionsData(article, language, function(err, revisionsData) {
        if (revisionsData.spiking) {
          // spiking article
        }
        if (geoData.averageCoordinates.lat) {
          // geo-referenced article, create map
        }
        // trigger alert if article is spiking
      });
    });
  };

  getMonitoringList(seedArticle, function(err, data) {
    // get the initial monitoring list
    if (err) {
      return console.log('Error_initializing_the_app.');
```

Listing 2. Main monitoring loop of the natural disaster monitor

Natural Disaster Monitor

Current status: Monitoring 141082 candidate Wikipedia articles.

Currently edited article: sv:Lysimachia_inconspicua

- ja:日本 (not spiking)



- **outbound link:** 1/8 Avalanche, 1/8 Flood, 1/8 Tsunami, 1/8 Earthquake, 1/8 Volcanic eruption, 1/8 Epidemic, 1/8 Wildfire, 1/8 Tornado
- **inbound link:** 1/2 Tsunami, 1/2 Earthquake
- **mutual link:** 1/2 Tsunami, 1/2 Earthquake

Fig. 2. Screenshot of the mobile-friendly “Natural Disaster Monitor” application prototype available at <http://disaster-monitor.herokuapp.com/> showing detected natural disaster types connected with the (currently non-spiking) article “Japan”

available user data through its streaming APIs,²⁴ which would be considerably harder, if not impossible, with other OSNs or search logs due to privacy concerns and API limitations. Based on the articles in the “monitoring list”, we put forward using article titles as search terms, but without disambiguation hints in parentheses, *e.g.*, instead of the complete article title “Typhoon Rammasun (2014)”, we suggest using “Typhoon Rammasun” alone. We advise monitoring the sample stream²⁵ for the appearance of any of the search terms, as the filtered stream²⁶ is too limited regarding the number of supported search terms. In order to avoid ambiguity issues with the international multi-language tweet stream, we recommend matching search terms only if the Twitter-detected tweet language equals the search term’s language, *e.g.*, English, as in “Typhoon Rammasun”.

5 Conclusions and Future Work

In this paper, we have presented first steps of our ongoing research on the creation of a Wikipedia-based natural disaster monitoring system, in particular, we have finished its underlying code scaffolding. While the system itself already works, a good chunk of work still lies ahead with the fine-tuning of its parameters. A first examples are the exponential smoothing parameters of the revision intervals, responsible for determining whether an article is spiking, and thus a potential new natural disaster, or not. A second example is the role that natural disasters play with articles: they can be inbound, outbound, or mutual links,

²⁴ Twitter streaming APIs: <https://dev.twitter.com/docs/streaming-apis/streams/public>

²⁵ Twitter sample stream: <https://dev.twitter.com/docs/api/1.1/get/statuses/sample>

²⁶ Twitter filtered stream: <https://dev.twitter.com/docs/api/1.1/post/statuses/filter>

and their importance for actual occurrences of disasters will vary. Future work will mainly focus on finding answers to our research questions Q1 and Q2 and the verification of the hypotheses H1–H3. We will focus on the evaluation of the system’s usefulness, accuracy, and timeliness in comparison to other keyword-based approaches. An interesting aspect of our work is that the monitoring system is not limited to natural disasters. Using an analog approach, we can monitor for human-made disasters (called “Anthropogenic hazard” on Wikipedia) like terrorism, war, power outages, air disasters, *etc.* We have created an exemplary “monitoring list” and made it available.²⁷ Concluding, we are excited about this research and look forward to putting the final system into operational practice.

References

1. F. Abel, C. Hauff, G.-J. Houben, R. Stronkman, and K. Tao. Twitcident: Fighting Fire with Information from Social Web Streams. In *Proceedings of the 21st International Conference Companion on World Wide Web, WWW ’12 Companion*, pages 305–308, New York, NY, USA, 2012. ACM.
2. R. Falor. Search data reveals people turn to the internet in crises, Aug. 2011. <http://blog.google.org/2011/08/search-data-reveals-people-turn-to.html>.
3. M. F. Goodchild and J. A. Glennon. Crowdsourcing Geographic Information for Disaster Response: A Research Frontier. *International Journal of Digital Earth*, 3(3):231–241, 2010.
4. N. Laframboise and B. Loko. Natural Disasters: Mitigating Impact, Managing Risks. IMF Working Paper, International Monetary Fund, 2012. <http://www.imf.org/external/pubs/ft/wp/2012/wp12245.pdf>.
5. J. Lin and G. Mishne. A Study of “Churn” in Tweets and Real-Time Search Queries (Extended Version). *CoRR*, abs/1205.6855, 2012.
6. S. B. Liu, L. Palen, J. Sutton, A. L. Hughes, and S. Vieweg. In Search of the Bigger Picture: The Emergent Role of On-line Photo Sharing in Times of Disaster. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*, 2008.
7. T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *Proceedings of the 19th International Conference on World Wide Web, WWW ’10*, pages 851–860, New York, NY, USA, 2010. ACM.
8. T. Steiner. Bots vs. Wikipedians, Anons vs. Logged-Ins (Redux): A Global Study of Edit Activity on Wikipedia and Wikidata. In *Proceedings of The International Symposium on Open Collaboration, OpenSym ’14*, pages 25:1–25:7, New York, NY, USA, 2014. ACM.
9. T. Steiner. *Enriching Unstructured Media Content About Events to Enable Semi-Automated Summaries, Compilations, and Improved Search by Leveraging Social Networks*. PhD thesis, Universitat Politècnica de Catalunya, 2014.
10. J. Westfall. Common Alerting Protocol Version 1.2. Standard, OASIS, 2010. <http://docs.oasis-open.org/emergency/cap/v1.2/CAP-v1.2-os.doc>.

²⁷ Anthropogenic hazard “monitoring list”: <https://github.com/tomayac/postdoc/blob/master/papers/comprehensive-wikipedia-monitoring-for-global-and-realttime-natural-disaster-detection/data/monitoring-list-anthropogenic-hazard.json>