# Natural disasters in Israel and USA as documented in the social media of Twitter

## Final report – December 2020

### The Israel National Knowledge and Research Center for Emergency Readiness

Prof. Avigdor Gal[1], Bar Genossar[1], Dr. Motti Zohar[1]

[1] Faculty of Industrial Engineering & Management, Technion

[2] The department of Geography and Environmental Studies, University of Haifa

**תקציר**

בעשור האחרון צמחה טוויטר כרשת חברתית בשיעור גבוה בקרב משתמשים רבים ומונה, נכון להיום מעל ל-500 מיליון משתמשים ברחבי העולם. בישראל עולה שיעור המשתמשים ברשת מאז 2014 בכ-100000 משתמשים בשנה ועומד על למעלה ממיליון. המסרים המועברים באמצעות טוויטר (״ציוצים״) מונים עד 280 תווי טקסט ויכולים להישלח על בסיס חיבור לרשת האינטרנט באמצעות טלפונים ניידים (׳סמארטפונים׳) ושלל אפליקציות צד-שלישי. במהלך התרחשות קטסטרופה סביבתית ואחריה עולה הנטייה ״לצייץ״ ועמה התעבורה ברשת. בעוד שהמידע הנגזר מטוויטר נמצא בשימוש ברחבי העולם למתן התרעה הקרובה לזמן אמת, חקר הפריסה הגיאוגרפית למיצוי תובנות יישומיות עדיין בחיתוליו. במחקר זה ביצענו רכשנו נתונים מטוויטר לפני ואחרי אירועי סיכונים טבעיים בולטים שהתרחשו בישראל וסביבתה ובארה״ב דוגמת שיטפונות, שריפות ורעידות אדמה. ניתוח נתונים אלו נועד לבחון האם הרשת החברתית מאפשרת מתן התרעה הקרובה לזמן אמת וניתן להשתמש בה למיפוי פריסת הנזקים. דו״ח זה מפרט את עיקרי המתודולוגיה ותוצאות ראשוניות של המחקר.



## ABSTRACT

During the last decade, the social network of Twitter has become a robust platform for distributing messages (tweets) among numerous subscribers worldwide. To date, Twitter is used by more than 500 million users worldwide. In Israel the growth of twitter subscribers is by ~100,000 per year since 2014 and to date consists of over 1,000,000 subscribers. The tweets, up to 280 characters only, can be sent via web pages, mobile devices, or third-party Twitter applications. During and around the occurrence of natural hazards, people tend to over-tweet and consequently, the number of tweets raise significantly. While Twitter is already in use for near real-time alerts, processes for extracting reported damage from tweets and examining the resulted spatial distribution are still under development. In this study we acquired tweets sent prior to and after natural hazards such as floods, fire and earthquakes that occurred in Israel and its close surroundings and the United states. The analysis of the tweets is aimed at verifying whether the social media can be used to near real time alerts and portraying the spatial distribution of the damage. This report elaborates on the methodology conducted and presents the preliminary results.

## 1. INTRODUCTION

During the last decade, the microblogging network of Twitter has become a robust platform for distributing messages (tweets) among numerous subscribers worldwide (Naaman et al., 2011). To date, Twitter is used by more than 500 million users worldwide. In Israel, the growth of twitter subscribers is by ~100,000 per year since 2014 and to date consists of over 1,000,000 subscribers.[1] The tweets, up to 280 characters only, can be sent via web pages, mobile devices, or third-party Twitter applications. A subscriber can tweet a message to his/her followers, transmit messages other subscribers have published (namely re-tweet) or reply to tweets he/she had received. The content of the tweets varies and may include any subject the subscriber has in mind (Java et al., 2007). Twitter is being used differently by communities residing in various regions that differ in their social, cultural, and economic characteristics. For instance, it was found that the intensity of the usage within a specific region correlates with the average income and education of the local residence (Li et al., 2013; Marwick & Boyd, 2011; Singh et al., 2009). Naturally, during crises or periods of stress some of the users tend to tweet more than usual, sharing their personal experience or distress, occasionally alongside descriptive information (Vieweg et al., 2010). This may also be the case in response to natural disasters and indeed, substantial rise in the number of tweets shortly after a disaster was apparent in many occurrences (De Longueville et al., 2009; Kent & Capello Jr, 2013; Sakaki et al., 2010; Saravanou et al., 2015). Obviously, due to the distressed situation the user may be in, the content of a tweet may not reflect the actual occurrence. it may contain descriptive information alongside personal expressions, but the content may certainly be suspected of being inaccurate, incomplete, or exaggerated, not to mention rumors that can spread worldwide by large number of users, whether deliberately or innocently (Takayasu et al., 2015).

One of the major challenges in the analysis of microblogging content like tweets is the extraction of geographic location to portray the spatial distribution of a phenomenon (Cheng et al., 2010; Ribeiro et al., 2012; White, 2010). As far as natural disasters are concerned, it is of great importance since many of the affected regions lack instrumental sensors or they cannot be installed. For instance, seismographs are sensitive to noise, thus cannot be located within a densely populated area and hydrographs are hard to be placed in steep Canyons. When extracting the

[1] The worldwide growth is based on https://www.statista.com/statistics/282087/number-of-monthly-activetwitter-users/ while for Israel see https://www.statista.com/statistics/558324/number-of-twitter-users-in-israel/ [accessed October 2019]

geographic location of a tweet, there are two major implications. The first, is the location the tweet was sent from while the second is the location the tweet is reporting of (Ribeiro & Pappa, 2018). Roughly, there are three approaches for extracting geographic locations. The first is the geographic coordinates and meta-information tags associated to each tweet that was sent from a mobile device, provide the location services (GPS) were active. In general, this is relevant for only 1-3% of the tweets depending on geographic region and social characteristics (Leetaru et al., 2013). The second approach is based on methods for extracting geographic locations out of digital text documents (Smith & Crane, 2001), web pages and blogs and news pages (Zhang et al., 2011) by referring the content tokens to well-known location lists (namely gazetteers). This approach should take into consideration that tweets are very much different from traditional text whereas they are short and occasionally contain misspellings and slangs, thus posing additional challenges for georeferencing process (Bouillot et al., 2012; Gelernter & Mushegian, 2011; Hahmann et al., 2014; Han et al., 2014). The last approach includes network-based methods that conclude locations by using the relations between users such as their mutual followers or other users mentioned inside the tweet's content (Ribeiro & Pappa, 2018).

This aim of this report is to describe the selected natural hazards as our inspected test cases, the data acquisition process, the methodology conducted to filter the associated tweets and preliminary results.

## 2. INSPECTED TEST CASES

Seven test cases were selected for the analysis. They represent earthquakes, fires and flashfloods events occurred in Israel and USA. These countries differ substantially in their geographic size and extent, population, culture, and social characteristics. Therefore, the comparison between the how Tweeter is used, and the way tweets are being sent in both countries is important for the understanding how one can use the data efficiently. The selected three and four test cases in Israel (Figure 1) and USA (Figure 2), respectively:

i. A series of 33 earthquakes ($M_W \geq 2.5$) occurred in northern Israel between 04-29, July 2018 (http://seis.gii.co.il/en/earthquake/searchEQS.php, accessed September 2019)
ii. The fire in Haifa and its close vicinity lasting between 23-26, November 2016 (Tessler et al., 2019)
iii. A series of flash floods in southern Israel between 25-26, April 2018 (Giva'ti et al., 2018)

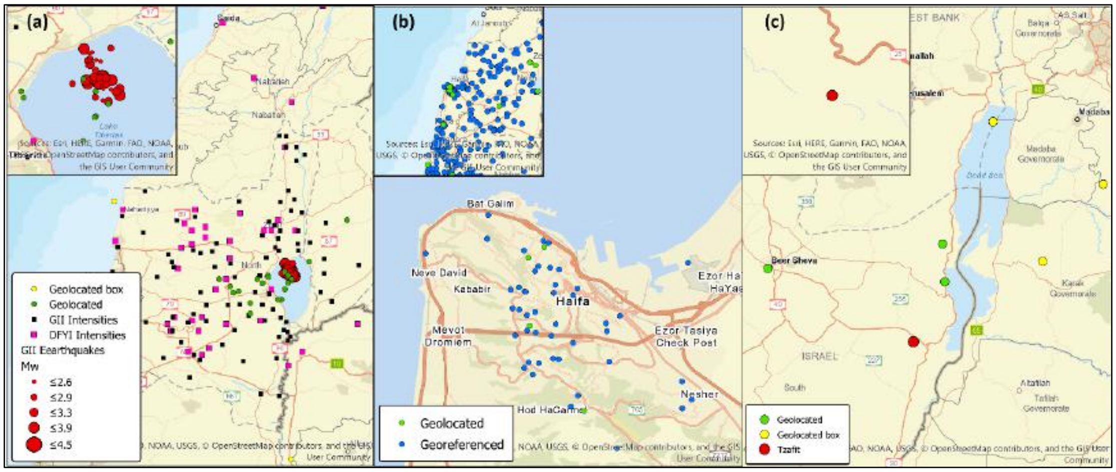


**Figure 1**: The research area of the inspected test cases in Israel. The distribution of tweets that are Geolocated (GPS meta coordinates), Geolocated box (a polygon bounding box meta coordinates, we use the upper left coordinates) and Georeferenced (manipulated location extraction using a gazetteer) are marked with yellow, green, and blue circles, respectively. Note that the tweets are already filtered by keyword/tokens during the acquisition process: (a) The sequence of earthquake in the Galilee (Israel). The sequence occurred at the north of the Lake of Galilee between 04.07.2018 and 29.07.2018, as is noted by red circles, scaled in accordance with the magnitude (see also the inset map). Seismic intensities of the two topmost earthquakes (the $M_W$=4.5 at 04.07.2018 19:15:39 and the $M_W$=4.2 at 04.07.2018 01:50:06) are noted by black (source: GII) and purple (source: USGS) squares; (b) The fire in Haifa and its vicinity on 24.11.2016. Location, reports, affected areas, emergency calls; and (c) the flash flood at Tzafit river on 26.04.2018. Location, reports, affected areas, emergency calls.

iv. The earthquake in California at 28.01.2015 21:08:53 (UTC) (https://earthquake.usgs.gov/earthquakes/eventpage/nc72387946/executive accessed September 2019)

v. The earthquake in Alaska at 29.07.2015 02:35:59 (UTC) (https://earthquake.usgs.gov/earthquakes/eventpage/ak11661615/executive, accessed September 2019).

vi. The floods in Missouri 23-29, December 2015 (Criss & Luo, 2016).

vii. Wildfires in California between 12-18, September 2015 (Gao et al., 2017)

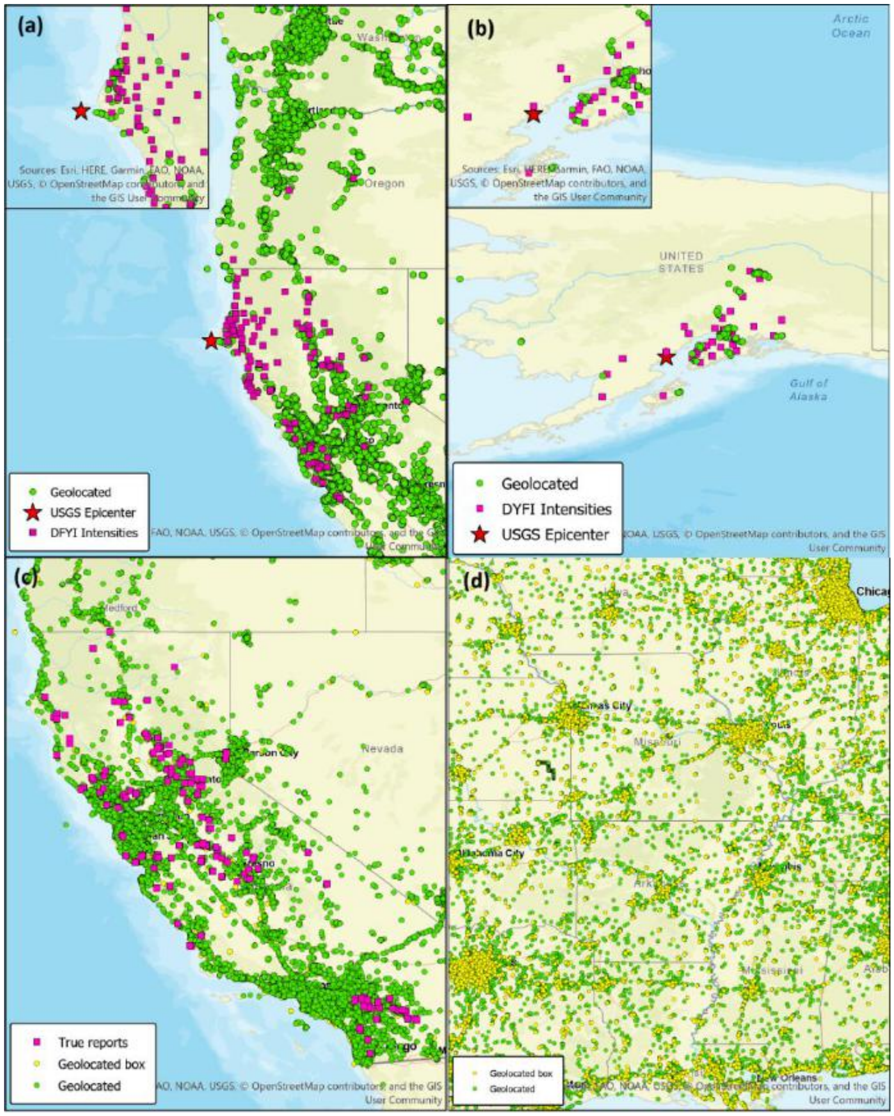


**Figure 2:** The research area of the inspected test cases in USA. The distributions of tweets that are Geolocated (using GPS meta coordinates), Geolocated box (a polygon bounding box meta coordinates, we use the upper left coordinates) and Georeferenced (manipulated location extraction using a gazetteer) are marked with yellow, green, and blue circles, respectively. Here the tweets are yet to be filtered by keyword/tokens: (a) The earthquake occurred in California at 28.01.2015 21:08:53 (UTC). Seismic intensities (source: USGS) are noted by purple squares; (b) The earthquake occurred in California of 29.07.2015 02:35:59 (UTC); (c) Two major fire events prevailed between 12-18 September 2015. Reports of the CAL FIRE (Office of the State Fire Marshal, USA) is noted by purple squares; and (d) Floods in 23-29 December 2015.[2]

[2] Some of the listed datasets will processed during the 2nd year of the study: (1) Georeferenced tweets using gazetteers; (2) The spread of fires in Haifa and its surrounding in November 2016; (3) data associated with the flash flood at Tzafit river in April 2018; and (4) data associated with the USA floods in December 2015

## 3. RELATED WORK

In this section we provide a brief survey of the main elements in our study.

***Text Classification:*** Classification is an important functional ability of categorizing data instances into different predetermined class labels. This is done using supervised learning, where labeled data samples are exploited in order to create a mechanism that can learn patterns and generalize to new ones. Classic machine learning algorithms such as Perceptron, Naive Bayes, Nearest Neighbor, etc., had lead this field for many years. However, when dealing with raw textual data, these methods might face problems due to the difficulty of creating meaningful features without a relational structure to rely on. During the last decade, deep learning (DL) has gained a significant role in this game, obviating the necessity of manual feature engineering. The emergence of word embedding (Mikolov et al., 2013) has allowed to capture complex semantic representation of sentences and deal with more complex problems.

***Neural Networks:*** Neural network is a proven tool in the world of machine learning, which has gained acceleration for the last few years. The most rudimentary structure of neural network is fully connected layer. In this component, input vector $\mathrm{x}$ is multiplied by weights vector $\mathrm{w}$ and the result is added to a bias variable $\mathrm{b}$ . Then, the result can pass through non-linear activation function $\sigma$. This linearity breaking enables neural networks to have better perception of complex tasks and datasets, which classical machine learning methods usually would not be able to obtain. The idea of fully connected layer can be extended to multi-layer neural network by piling up sequentially several layers. The first layer receives the initial input, and from there on each successive layer gets the output from the previous one. This is done until the last one transmits its output as the final decision of the network.

***Word Embeddings:*** Word Embeddings, presented by (Mikolov et al., 2013) has an integral role in most NLP tasks. Its main concept is creating a vector-space representation of words, based on their context. It means that words which tend to co-occur in similar environment will probably have similar vectors. Usually the vectors are trained beforehand and combined with other neural networks, depends on the task.

***Topic Modeling:*** Topic modeling (Blei et al., 2009) refers to statistical techniques in natural language processing (NLP), aiming to mine patterns out of documents collections and cluster them according to a shared semantic. Latent Dirichlet Allocation (Blei et al,. 2003) is a useful generative probabilistic topic modeling approach, where each word of a documented is modeled over a topics distribution vector. Another core principle of this approach is that document can include multiple topics, therefore it can also be modeled over a topics distribution vector.

## 4. METHODOLOGY

In this section we present the data used in our study and elaborate on our approach to extract meaningful information out of it. Given an event, its corresponding dataset was created according to temporal and geographical conditions. The tweets were collected from slightly before the beginning of the event until little after its termination. As one could guess, a large portion of the data is actually not related to the examined events. Hence, we developed a complete pipeline, ending with a classifier, for detecting whether or not a tweet is relevant (i.e., related to the event). Due to the different characteristics of the datasets in terms of language, jargon and type of event, we performed this pipeline separately. This statement holds also within the same dataset when it includes multiple languages (the datasets from Israel). The pipeline is illustrated in Figure .

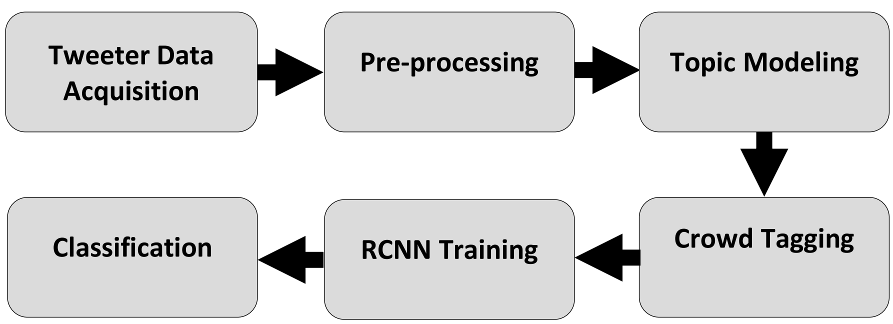


**Figure 3:** Work pipeline architecture

### 4.1 Tweeter data acquisition

Twitter data associate with Israel and its close neighbors was acquired from Twitter© in accordance to pre-defined acquisition rules. The prime languages were Hebrew, Arabic and English. Tweets from USA were acquired in the curtesy of the Infomedia MAGNET project which include a sample of 5% from all tweets sent in English from the USA by users who posted at least 20 statuses/tweets during 2015 and have at least 10 followers and 10 friends (Table 1).

| | Event | Tweets availability | Total tweets | Geolocated Tweets |
|---|---|---|---|---|
| **Israel** | The fire in Haifa, November 2016 | 2016-11-22 00:00:00 – 2016-11-27 23:58:59 | 3975325 | 6249 |
| **Israel** | Earthquakes in Israel, July 2018 | 2016-07-03 00:00:00 – 2016-07-29 23:58:57 | 3255699 | 87000 |
| **USA** | The fires in California, September 2015 | 2015-09-12 01:00:00 – 2015-09-18 23:59:59 | 12591787 | 1140314 |

**Table 1**: General information of the inspected tweets.

Raw data acquired from Twitter© is organized in Java Script Object Notation (JSON) files and classified into intervals of 10 minutes. That is, if the tweets corpus includes records along the entire day intervals then it is bundled within 144 files starting with the 00:00:00 – 00:09:59 AM bundle while ending with the 23:50:00 – 23:59:59 PM bundle. The raw tweets files were parsed using python and uploaded to relational tables in Microsoft© Sql Server database, as part of an Azure Microsoft© cloud environment. We have used two parsers for this process whereas the JSON structure of the 2015 tweets is significantly different from that of 2016; up to 2016 a tweet length was limited to only 140 characters while afterwards the length was extended to 280 characters.

### 4.2 Preprocessing

Processing the content of tweets deserves specific attention for their short length, alongside their tendency to include more slang and noise than formal texts (Sriram et al., 2010; Nguyen et al., 2015; Steiger et al. 2015, Saif et al., 2014). Because of these reasons, twitter data requires preliminary steps of making the data more consistent and coherent. It is important to mention that this step might be suitable for some classification tasks, but not for others. Since our pipeline involves two stages of classification, namely topic modeling and relevance classification, different steps are performed for each one.

***Tokenization:*** Token is a textual unit (string, number or sequence of special character) separated from others by white space. Tokenization of words is the process of getting the tokens out of the text content, i.e., breaking the sentences into list of words. This is a crucial step in many machine learning algorithms, allowing to models to learn by frequency of words in the data.

***Lowercasing:*** Substitute all the capitalized letters with small ones. This step is executed because we want to treat the identical words with different capitalization as the same one.

***Short Words Removal:*** Remove all the words that include less than 3 characters. Those words are usually less indicative of the content semantic (Pak et al., 2010).

***Stopwords Removal:*** Stopwords are the most common words in a language, usually do not deliver an indicative semantic information as stand-alone tokens. We used the NLTK (Loper et al., 2002) stopwords list to eliminate them. The effectiveness of removing those words as a preliminary step for classification tasks is unclear (Saif et al., 2014), yet it is necessary for reducing the computational effort and to avoiding the dominance of frequent words in tasks relied on words distribution. However, when we deal with more sophisticated learning model, such as the deep learning based classifier, stopwords are important.

***Stemming:*** Maintaining the base root of words by removing their prefixes and suffixes of words. This step allows us to combine together words with the same root but different form. We used Porter Stemmer, implemented as part of NLTK (Loper et al., 2002).

***Others:*** We used regular expressions to remove multiple spaces, special characters and words preceded by '\'.

### 4.3 Topic Modeling

Topic modeling (TM) allows us to extract the "essence" of a tweet using the distribution of words over the entire collection. Traditionally, TM is used for a collections of documents, but applying it to tweets requires different approach. The short length of a tweet makes the mining of its underlying topic much harder. (Hong et al., 2010) suggested to aggregate tweets that were tweeted by the same user into a single document. Following this line, we filtered all the tweets that were tweeted by a user with less than 3 documented tweets in the dataset. For these ones, we concatenated the tweets of the same user and performed Latent Dirichlet Allocation (LDA) according to their vocabulary. Our underlying assumption is that user is likely to tweet about the same topics in a short period of time. The described mechanism resulted in a distribution vector of users over the topics, where topics are characterized by a set of representative words. We assigned each tweet its topic with respect to element with the largest value in the topics distribution vector of the tweet's user. In this stage, we surveyed the topics and decided which ones appear to be relevant to the examined event. Then, we kept only potentially relevant tweets.

### 4.4 Crowd

After the execution of the topic modeling, only the selected candidate relevant tweets have remained. In order to train the classifier, annotated tweets are required. From each filtered dataset we randomly sampled a collection of tweets and sent it to crowd-sourcing. In other words, human annotators were hired to label a train-set according to pre-defined guidelines. For the data from Israel, the annotators were asked to assign each tweet one of the two labels: *relevant*, or *irrelevant*. Relevant tweets are ones that contain reaction related to the examined event. Tweets which are not related at all to the events are considered as irrelevant.

Some tweets are actually just sharing of others' tweets (retweet). In this case, the exact same content appear in multiple tweets. We do not want to ignore it when we analyze the reaction to the events, as reflected in Twitter. However, we do not want to use the same content in many training samples, as it biases the model and prevents it from learning new patterns. To make sure that the same content does not appear multiple times before delivering the data to the annotators, we created a set of strings such that each element is a content of tweet. We did not allow new tweets to enter the final set of candidates if its content was already included in it.

### 4.5 RCNN Training

Lai et al., (2015) presented a model called Recurrent Convolutional Neural Networks (RCNN). This model can capture contextual information relying on word representations and use it to classify textual input to its underlying category. Unlike topic modeling, stopwords should not be ignored here since they can provide a contextual information of other words. In this subsection we introduce the model architecture and explain how we trained it using the annotated tweets.

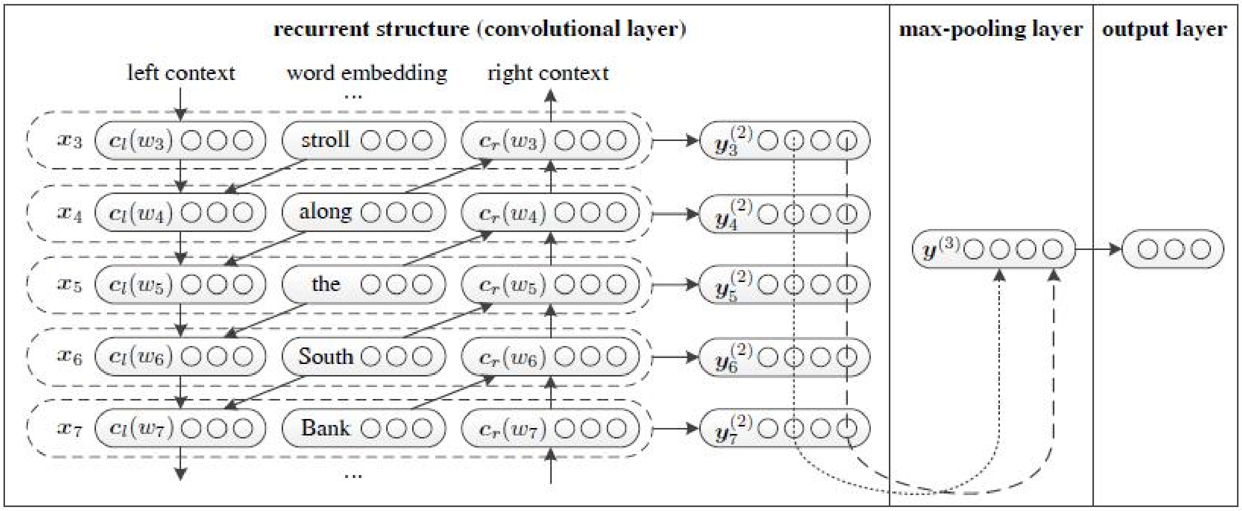


**Figure 4**: RCNN architecture (Lei et al., 2015)

### 4.5.1 Architecture

***Text Representation***: Given a text input, $T = w_1, w_2, ..., w_{|T|}$ we associate each word $w_i \in T$ with a pre-trained embedding vector $e(w_i)$ of size $|E|$. This allows us to represent a word according to its common context in the language, but still not as integral part of the current input. Therefore, we use the architecture of BILSTM (Huang et al., 2015) and create the vector $c_i = \text{BILSTM}(T_{1:|T|}, i)$ where $i$ is the index of the word in the sentence. Thanks to its structure, $c_i$ captures both left and right context of the word in the text, without window size limitation. Next step is creating a new representation for $w_i$ by concatenating the two described vectors such that $x_i = e(w_i) \circ c_i$, where $\circ$ is the concatenation operator. In this way we can define $X$ as the set of all representative vectors of the text words ($X = \{x_i\}, i = 1, 2, ..., |T|$).

***Breaking Linearity***: $\forall x \in X$ we run the non-linear activation function $\tanh$ applied on fully connected layer. Without breaking linearity, the model would not be able to contain the complexity of the input. If so, we define the latent semantic vector, $y_i^{(2)}$ where (2) implies the stage index in our process, as:

$$y_i^{(2)} = \tanh\left(W^{(2)} x_i + b^{(2)}\right)$$

***Pooling***: Afterwards, we apply max-pooling on the latent semantic vectors, $y_i^{(2)}$. The goal of max-pooling layer is to extract the most important features observed in the entire text by converting texts with variety of lengths to a fixed-length vector. The operation is executed by extracting the largest value over each $y_i^{(2)}$, for all the entries in the vector $y^{(2)}$, such that:

$$y^{(3)} = \max_{i=1}^{|T|} \left( y_i^{(2)} \right)$$

***Output Layer***: Just before we transmit the result, we apply fully connected layer which maps the latent vectors to the output vectors, compatible with the labels space size.

$$y_I^{(4)} = w^{(4)} y_i^{(3)} + b^{(4)}$$

Finally, we transform the output values to probabilities by a softmax function.

$$P_I = \frac{\exp\left(y_i^{(4)}\right)}{\sum_{i=1}^{n} \exp\left(y_i^{(4)}\right)}$$

Where n is the number of possible labels (categories). The reported label is chosen as the one with the highest probability to be the true label:

$$\text{label} = \max_{i=1}^{n} (p_i)$$

We apply this architecture and create a unique model for each examined event and language.

### 4.5.2 Settings

The hyper-parameters set includes the following variables: LSTM hidden dimension, linear layer dimension, epochs number, learning rate, batch size, momentum, loss function and optimizer.

***Loss Function***: In every single configuration we use negative log likelihood. The function is defined as follows:

$$\min_{\theta}(\mathrm{NLL}) = \min_{\theta}\left(\sum_{(\mathrm{T},\ell)\in \mathrm{D}} -\log\left(\frac{\exp(\mathrm{score}(\mathrm{T}_{\ell}))}{\sum_{\ell'\in \mathrm{L}}\exp(\mathrm{score}(\mathrm{T}_{\ell'}))}\right)\right)$$

where $\theta$ denotes the network parameters, $\mathrm{D}$ is all the tuples of the form ($text$, true $label$ ) and $\mathrm{score}(\mathrm{T}_{\ell'})$ is the output of the network for label $\ell'$, before applying the softmax operation.

***Optimizer***: As for the optimization, we use stochastic gradient descent (Bottou, 1991) aiming at finding $\theta$ which maximizes the log-likelihood of the network.

***Weights Initialization:*** All the inner weights of the network are randomly initialized using Xavier (Glorot et al., 2010), sampled from $\mathrm{U}\left[-\frac{\sqrt{6}}{\sqrt{\mathrm{n}_i+\mathrm{n}_{i+1}}}, \frac{\sqrt{6}}{\sqrt{\mathrm{n}_i+\mathrm{n}_{i+1}}}\right]$, where $\mathrm{n}_i$ is the number of incoming connections to the layer and $\mathrm{n}_{i+1}$ is the number of outgoing connections of it.

***Word Embeddings*****:** In our work we used different pre-trained word embeddings. For English we used glove (Pennington et al., 2014) with dimension of 200, For Hebrew we used (Shemesh 2018) with dimension 100 and for Arabic (Soliman et al., 2017) with dimension 100. Despite the vectors were already trained in advanced, we allow our model to update their weights during the backpropagation process.

***Train/Validation*****:** For the fire datasets (both for Israel and USA) we allocate 0.8 of the annotated data for training, and 0.2 for validation in order to evaluate the performance of the trained models. As for the EQ dataset (Israel) we do not use validation set due to the small number of tweets in Hebrew.

### 4.6 Classification

Table 2 shows the general information of the used data. Among the tweets that passed the topic modeling filtering, a small share was used for training and validation (after it was annotated) and the rest for classification for further analysis (target tweets). As explained in Section 4.5.1, the output value of each label is transformed into a probability, such that the label with the highest value is chosen. We decided to consider tweets as relevant only if it was the classifier decision, and its confidence surpassed a threshold of 0.6.

| | Event | Language | # Annotated Tweets | # Target Tweets | F1 Validation |
|---|---|---|---|---|---|
| **Israel** | The fire in Haifa, November 2016 | Hebrew<br>English<br>Arabic | 4439<br>4189<br>3231 | 32765<br>26051<br>219363 | 0.691<br>0.941<br>0.847 |
| **Israel** | Earthquakes in Israel, July 2018 | Hebrew | 1377 | 2333 | - |
| **USA** | The fires in California, September 2015 | English | 3690 | 120741 | 0.927 |

**Table 2**: Characteristics of the data

The evaluation of the model performance over the validation set is done using F1 measure, which is calculated as:

$$2 \cdot \frac{precision \cdot recall}{precision + recall}$$

Where:

$$precision = \frac{\left|\{annotated\ as\ relevant\ tweets\} \cap \{classified\ as\ relevant\ tweets\}\right|}{\left|\{classified\ as\ relevant\ tweets\}\right|}$$

$$recall = \frac{\left|\{annotated\ as\ relevant\ tweets\} \cap \{classified\ as\ relevant\ tweets\}\right|}{\left|\{annotated\ as\ relevant\ tweets\}\right|}$$

This score is calculated after every single epoch of training, and the final model is the one that achieves the best results. Since the dataset of earthquakes does not have a validation set, the final model is the one that reaches the highest score over the training set.

### 4.7 georeferencing none-geolocated tweets

Geolocated tweets are those associated with meta-info tags containing exact location (GPS coordinates in case a user has activated location services in his mobile phone), or a bounding box extracted using other geolocation capabilities such as the IP of the sending device and presenting an approximate location. For each of the tweets we have extracted the geolocation or bounding box, if existed. That is, each of the tweets that contain any reference of a geographic location in its content (text), user location or user description tags, was cross-corelated with a gazetteer in order to find a match (s). The latter are location-based reference of Points of Interest (POI) such as cities, streets, hospitals, shops, restaurants, bus stations and other known localities. Along with their name, gazetteers also contain synonyms and spatial reference i.e., geographic coordinates. Twitter has been used before for georeferencing in USA but rarely in Israel, in particular for content in Hebrew and Arabic languages. For georeferencing tweets associated with events occurred in Israel, we have used the OpenStreetMap (OSM, https://www.openstreetmap.org, last accessed July 2019) which, although scares in remoted regions, is fairly accurate (Haklay, 2010; Neis et al., 2012) and is being utilized in several studies (Clemens, 2015; Helbich et al., 2012; Liu & Long, 2016).[3] For georeferencing tweets from California we have used the Geonames (https://www.geonames.org/, last accessed July 2019) gazetteer that has been tested before (Inkpen et al., 2017; Jurgens et al., 2015) and is recognized as one of the largest and most frequently used with more than 10 million records about geographic entities in different languages. When georeferencing using text matching to a gazetteer, one should consider two main ambiguity difficulties, namely the geo/non-geo ambiguity and geo/geo ambiguity (Brunner & Purves, 2008). The first can be demonstrated in location names that are also associated with other none-geographic terms. For example, the settlement of Lapid (in Hebrew: לפיד) is also the surname of an Israeli politician named Yair Lapid (in Hebrew: יאיר לפיד). The Geo-Geo ambiguity occurs when several places have the same name. For instance, the Allenby Street is located within the cities of Tel Aviv, Jerusalem, and Haifa. That is, upon appearance of the token 'Alenby' one cannot resolve which of the cities it refers to. Accordingly, we have georeferenced the tweets using the message (text), user location and user description tags.

[3] For the land use layer, we have used only the polygon features tagged as 'residential' to form an output layer of geographic zones. The streets layer was intersected with the newly created zones layer to avoid street duplication between different zones and create a unique identifier to each street (e.g., the 'Alenbi' street in Tel Aviv, Haifa and Shderot). The streets were then dissolved by zone and street name to result in a single feature for each street name in each zone. Then, the mid-point of each street was extracted. The same process was implemented also for the point POI layer and the two leyers were merged into a unified gazetteer output layer. Altogether, the zones and gazetteer layers total 1594 and 60584 features, respectively. We kept a tri-language interface (English, Hebrew and Arabic) and added translations of major zones/gazetteer locations where required (e.g., 'חיפה', 'Haifa' and 'حيفا').

## 5. PRELIMINARY RESULTS

The classification results are presented in Table 3. As one can notice, the labels prediction distribution is different for every single dataset. While the dataset of the fires in Haifa is more uniform, the other two datasets extremely tend to a majority class (relevant tweets for the earthquakes and irrelevant for the fires in USA). Besides, the tweets characteristics vary between the datasets, even if we only consider the relevant ones. For example, plenty of the relevant English-written tweets in fires datasets are news-related content, while the Arabic-written mostly contain gloater and jokes about the ongoing event. Another example is the connection between the tweet writer and the event. While most of the writers in the fire datasets are not directly affected by the events, some of the earthquakes tweeters immediately reported on the event when they felt it.

| | Event | Language | Relevance<br>Relevant | <br>Irrelevant |
|---|---|---|---|---|
| **Israel** | The fire in Haifa, November 2016 | Hebrew<br>Arabic<br>English | 14341<br>8048<br>92371 | 18424<br>18003<br>126992 |
| **Israel** | Earthquakes in Israel, July 2018 | Hebrew | 2233 | 100 |
| **USA** | The fires in California, September 2015 | English | 3731 | 117010 |

**Table 3**: The classification results

Figure 5 presents the preliminary results of the temporal distribution of tweets surrounding the Butte and Valley fires in September 2015 at California. Peaks of associated tweets is recorded (Figure 5b) in 12/09 during 13:00-15:00; 14/09 during 10:00-12:00 and 14:00-16:00; 15/09 during 06:00 – 12:00 and 18:00-20:00; 16/09 during 19:00-21:00; and 17/09 during 22:00-24:00. Superposition of the associated tweets and all tweets (Figure 5c) presents exceeded associated tweets in 12/09 10:00-14:00; 13/09 06:00-10:00 and 14:00-18:00; 14/09 06:00-16:00; 15/09 05:00-20:00; and 16/09 18:00-22:00. That is, between 12/09 13:00 (the burst of the fire) and 15/09 afternoon (the contain of the fire) (Pilmott et al., 2015) an alleged matching is found between the tweets trend and the burst of the fire as well as additional four peaks. However, these are only preliminary results that should be further investigated and verified (spatially and temporally).

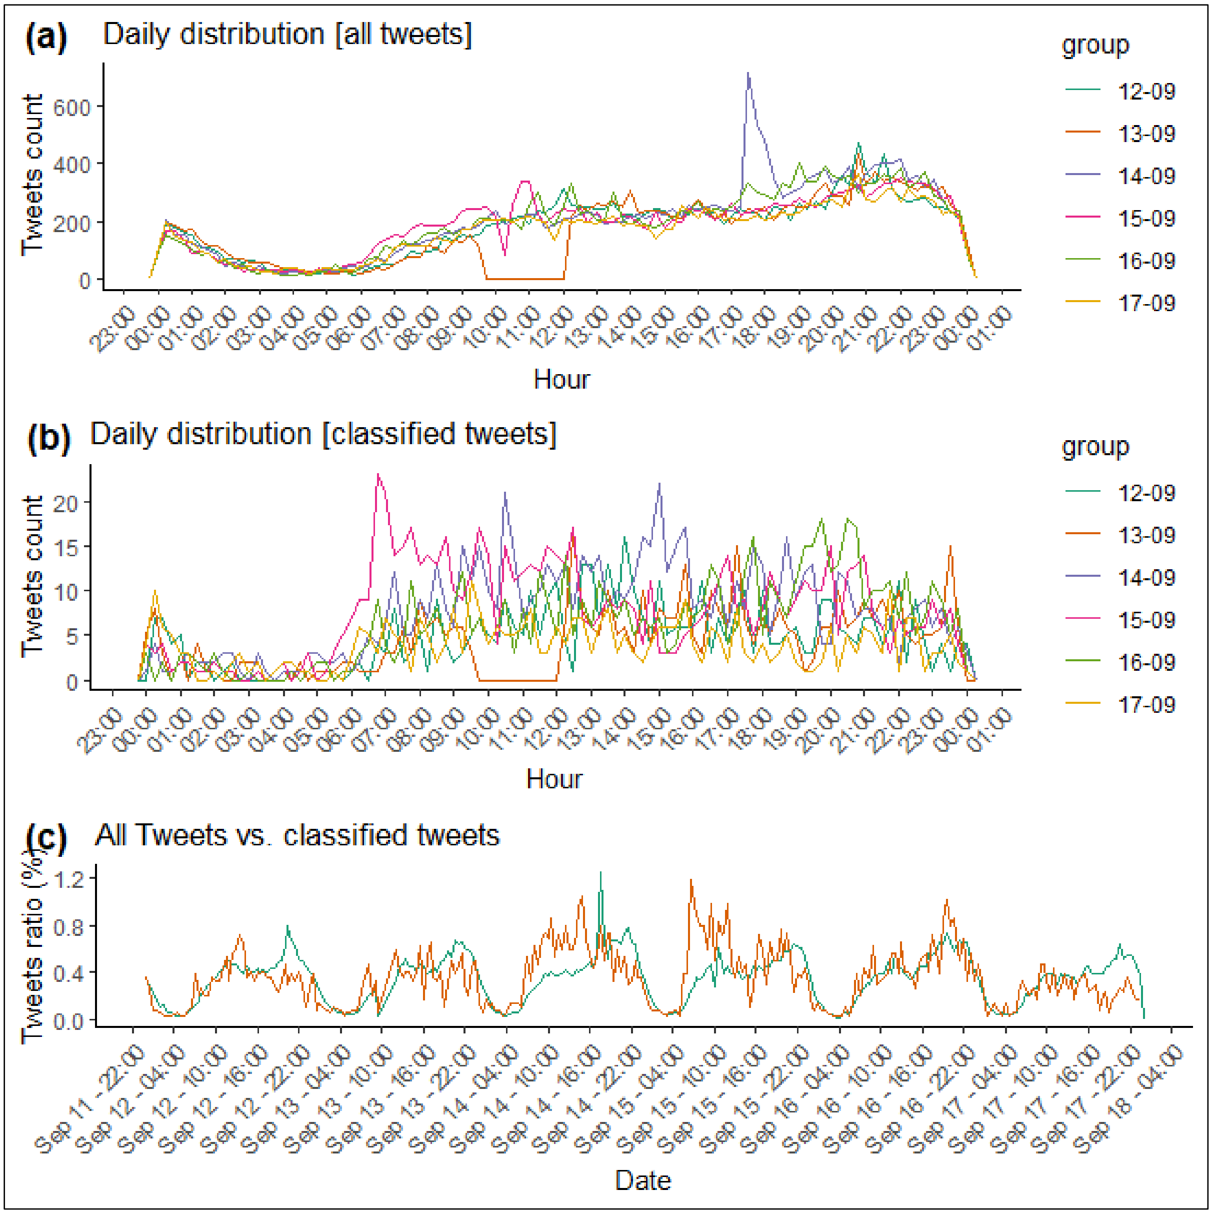


**Figure 5**: Temporal distribution of tweets associated with the Butte and Valley fires in California between 12.09.2015 and 18.09.2015 (a) daily distribution of all tweets (associated and not associated). Note the peak in 14/09 between 17:00 and 19:00; (b) daily distributions of the associated tweets. Note the peaks in 14/09 15:00, 15/09 07:00 and 16/09 20:00; (c) temporal distribution of the whole period presenting all tweets (green) vs. associated tweets (brown). Note the differences in 12/09 10:00 – 16:00, 14/09 10:00 – 16:00, 15/09 04:00 – 10:00 and 16/09 16:00 – 22:00

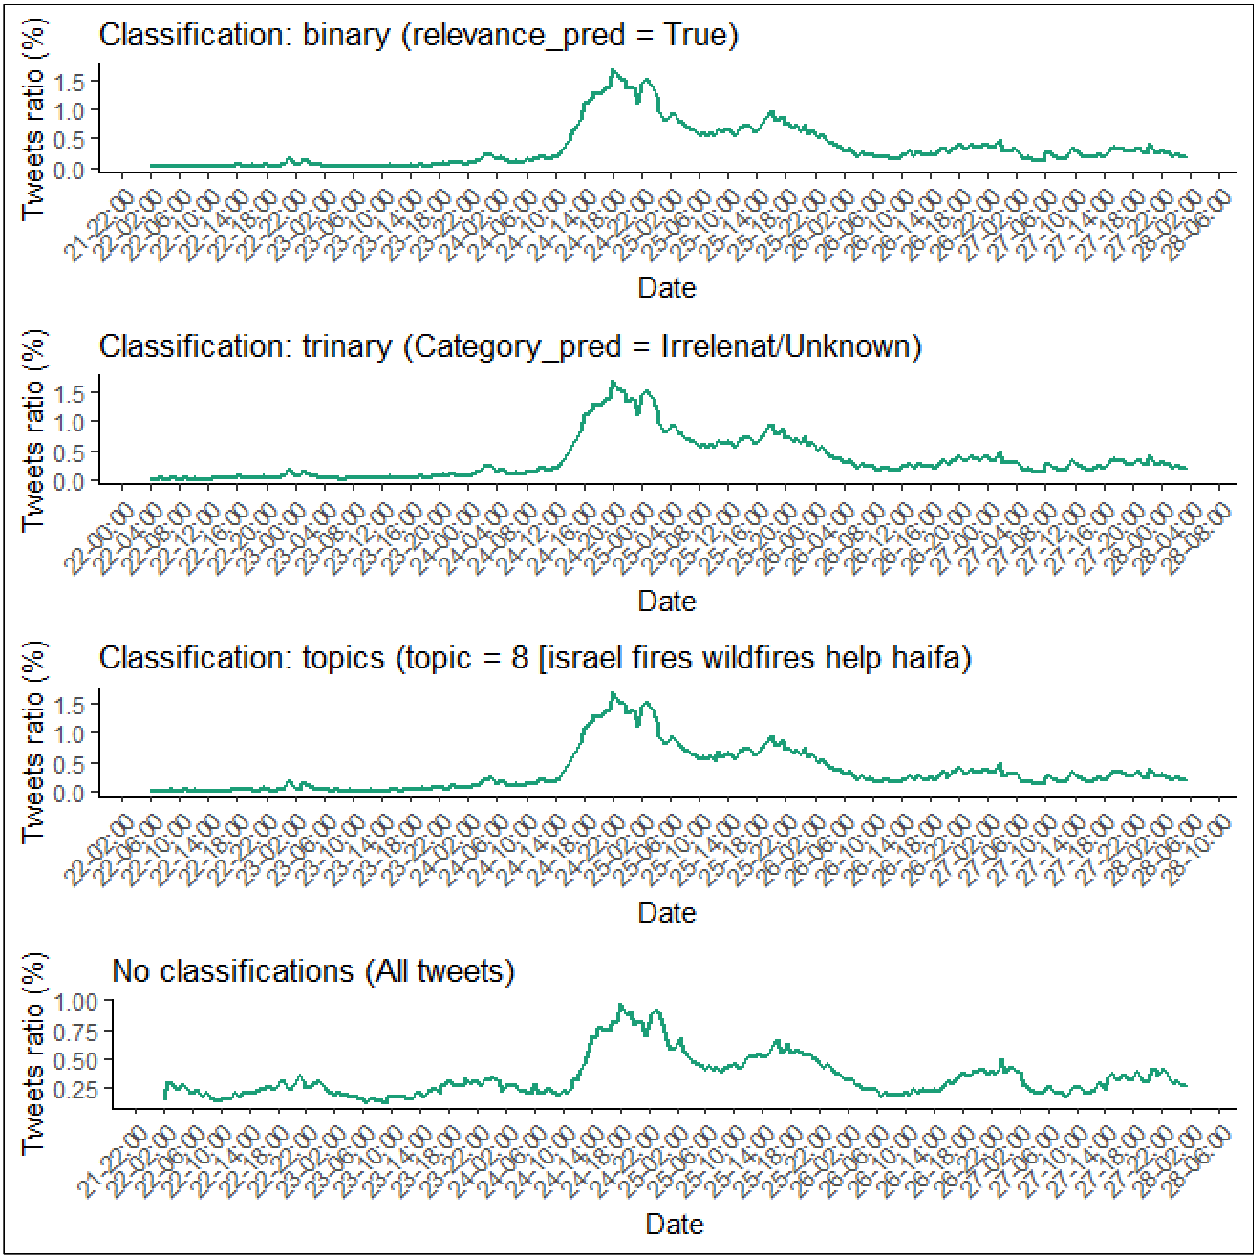


**Figure 6:** Temporal distribution of tweets associated with the Nov 2016 fire in Haifa. Note the resemblance of the different classification methods (top 3 images). The distribution of all tweets is presented in the lower image.

Figure 6 presents the comparison of three different classification methodologies (the binary, trinary and topic modeling) to derive the relevant tweets of the November 2016 fire in Haifa (Tessler et al., 2019). The top three figures demonstrate similar tendencies implying of cross-verification of the screened relevant tweets. Accordingly, the burst of the fire is well trended beginning the 24/11

12:00 with three consecutive daily peaks in 25/11, 26/11 and 27/11 between ~10:00-22:00. The first (25/11) is relatively notable but is lower that the one accompanying the initiating burst (24/11) while the second and third are relatively moderate. Yet, like the case of the California fire, these are preliminary results and should be inspected spatially in order to verify whether they accord with 'true' data that describe the cascading fire events

**Presentation in academic conferences and planned papers**

During 2021 we intend to complete the spatial and temporal examinations and move forward towards complementing the research. Accordingly, we intend to complete the following activities:

- The Israeli geographic association, Haifa, 2019
- Planned paper & participation in the GIScience conference, Poznan Poland 2021. Topic: using the free geographical databases of GeoNames and Open Street Map (OSM) to infer the location of tweets in Israel?
- Planned paper: The social media of Twitter during the fire in Haifa, Nov 2016, and during the sequence of earthquakes in the Lake of Galilee, July 2018

**References**


Blei, D. M., & Lafferty, J. D. (2009). Topic models. Text mining: classification, clustering, and applications, 10(71), 34.

Bouillot, F., Poncelet, P., & Roche, M. (2012). *How and why exploit tweet's location information?* In, N-A

Brunner, T. J., & Purves, R. S. (2008). *Spatial autocorrelation and toponym ambiguity.* In, Citeseer, 25-26

Bottou, L. (1991). Stochastic gradient learning in neural networks. Proceedings of Neuro-Nımes, 91(8), 12.

Cheng, Z., Caverlee, J., & Lee, K. (2010). *You are where you tweet: a content-based approach to geo-locating twitter users.* In Proceedings of the 19th ACM international conference on Information and knowledge management, ACM, 759-768

Clemens, K. (2015). Geocoding with openstreetmap data. *GEOProcessing 2015*, 10

Criss, R. E., & Luo, M. (2016). River management and flooding: The lesson of December 2015–January 2016, central USA. *Journal of Earth Science, 27*(1), 117-122

De Longueville, B., Smith, R. S., & Luraschi, G. (2009). *Omg, from here, i can see the flames!: a use case of mining location based social networks to acquire spatio-temporal data on forest fires.* In, ACM, 73-80

Gao, S., Li, L., Li, W., Janowicz, K., & Zhang, Y. (2017). Constructing gazetteers from volunteered big geo-data based on Hadoop. *Computers, Environment and Urban Systems, 61*, 172-186

Gelernter, J., & Mushegian, N. (2011). Geo-parsing messages from microtext. *Transactions in GIS, 15*(6), 753-773

Giva'ti, A., Zigl, E., Cundinoov, M., Paz, D., Galili, U., Wulin, L., Swartz, Z., Levi, Y., Chalfon, N., Wadislavski, E., & Munbaz, Y. (2018). Flash floods in the southern rivers 25 - 27/04/2018 (pp. 11). Jerusalem: The Water Authority.

Glorot, X., & Bengio, Y. (2010, March). Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the thirteenth international conference on artificial intelligence and statistics (pp. 249-256).

Hahmann, S., Purves, R., & Burghardt, D. (2014). Twitter location (sometimes) matters: Exploring the relationship between georeferenced tweet content and nearby feature classes. *Journal of Spatial Information Science, 2014*(9), 1-36

Haklay, M. (2010). How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and planning B: Planning and design, 37*(4), 682-703

Hong, L., & Davison, B. D. (2010, July). Empirical study of topic modeling in twitter. In Proceedings of the first workshop on social media analytics (pp. 80-88).

Huang, Z., Xu, W., & Yu, K. (2015). Bidirectional LSTM-CRF models for sequence tagging. arXiv preprint arXiv:1508.01991.

Han, B., Cook, P., & Baldwin, T. (2014). Text-based twitter user geolocation prediction. *Journal of Artificial Intelligence Research, 49*, 451-500

Helbich, M., Amelunxen, C., Neis, P., & Zipf, A. (2012). Comparative spatial analysis of positional accuracy of OpenStreetMap and proprietary geodata. *Proceedings of GI_Forum*, 24-33

Inkpen, D., Liu, J., Farzindar, A., Kazemi, F., & Ghazi, D. (2017). Location detection and disambiguation from twitter messages. *Journal of Intelligent Information Systems, 49*(2), 237-253

Java, A., Song, X., Finin, T., & Tseng, B. (2007). *Why we twitter: understanding microblogging usage and communities.* In Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis, ACM Press, 56-65

Jurgens, D., Finethy, T., McCorriston, J., Xu, Y. T., & Ruths, D. (2015). *Geolocation prediction in twitter using social networks: A critical analysis and review of current practice.* In Ninth International AAAI Conference on Web and Social Media, Oxford, England, Kent, J. D., & Capello Jr, H. T. (2013). Spatial patterns and demographic indicators of effective social media content during theHorsethief Canyon fire of 2012. *Cartography and Geographic Information Science*, *40*(2), 78-89

Leetaru, K., Wang, S., Cao, G., Padmanabhan, A., & Shook, E. (2013). Mapping the global Twitter heartbeat: The geography of Twitter. *First Monday*, *18*(5)

Lai, S., Xu, L., Liu, K., & Zhao, J. (2015, February). Recurrent convolutional neural networks for text classification. In Twenty-ninth AAAI conference on artificial intelligence.

Li, L., Goodchild, M. F., & Xu, B. (2013). Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartography and geographic information science*, *40*(2), 61-77

Liu, X., & Long, Y. (2016). Automated identification and characterization of parcels with OpenStreetMap and points of interest. *Environment and Planning B: Planning and Design*, *43*(2), 341-360

Loper, E., & Bird, S. (2002). NLTK: the natural language toolkit. arXiv preprint cs/0205028.

Marwick, A. E., & Boyd, D. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society*, *13*(1), 114-133

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Advances in neural information processing systems (pp. 3111-3119).

Naaman, M., Becker, H., & Gravano, L. (2011). Hip and trendy: Characterizing emerging trends on Twitter. *Journal of the Association for Information Science and Technology*, *62*(5), 902-918

Neis, P., Zielstra, D., & Zipf, A. (2012). The street network evolution of crowdsourced maps: OpenStreetMap in Germany 2007–2011. *Future Internet*, *4*(1), 1-21

Nguyen, M. T., Kitamoto, A., & Nguyen, T. T. (2015, May). TSum4act: A framework for retrieving and summarizing actionable tweets during a disaster for reaction. In Pacific-Asia Conference on Knowledge Discovery and Data Mining (pp. 64-75). Springer, Cham.

Pak, A., & Paroubek, P. (2010, May). Twitter as a corpus for sentiment analysis and opinion mining. In LREc (Vol. 10, No. 2010, pp. 1320-1326).

Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP) (pp. 1532-1543).

Pimlott, K., Laird, J. and J., B. E., 2015. Wildfire Activity Statistics. Sacramento, California: California Department of Forestry and Fire Protection. pp. 45

Ribeiro, S., Davis, C. A., Oliveira, D. R., Meira, W., Gonçalves, T. S., & Pappa, G. L. (2012). *Traffic observatory: a system to detect and locate traffic events and conditions using Twitter*. In 5th ACM SIGSPATIAL International Workshop on Location-Based Social Networks, ACM, 5-11

Ribeiro, S., & Pappa, G. L. (2018). Strategies for combining Twitter users geo-location methods. *Geoinformatica*, *22*(3), 563-587

Saif, H., Fernández, M., He, Y., & Alani, H. (2014). On stopwords, filtering and data sparsity for sentiment analysis of twitter.

Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). *Earthquake shakes Twitter users: real-time event detection by social sensors.* In, ACM, 851-860

Saravanou, A., Valkanas, G., Gunopulos, D., & Andrienko, G. (2015). *Twitter Floods when it Rains: A Case Study of the UK Floods in early 2014.* In, ACM, 1233-1238

Singh, P., Singh, P., Park, I., & Lee, J. (2009). Information sharing: a study of information attributes and their relative significance during catastrophic events *Cyber Security and Global Information Assurance: Threat Analysis and Response Solutions* (pp. 283-305): IGI Global

Smith, D. A., & Crane, G. (2001). *Disambiguating geographic names in a historical digital library.* In, Springer, 127-136

Soliman, A. B., Eissa, K., & El-Beltagy, S. R. (2017). Aravec: A set of arabic word embedding models for use in arabic nlp. Procedia Computer Science, 117, 256-265.

Sriram, B., Fuhry, D., Demir, E., Ferhatosmanoglu, H., & Demirbas, M. (2010, July). Short text classification in twitter to improve information filtering. In Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval (pp. 841-842).

Steiger, E., Westerholt, R., Resch, B., & Zipf, A. (2015). Twitter as an indicator for whereabouts of people? Correlating Twitter with UK census data. Computers, environment and urban systems, 54, 255-265.

Takayasu, M., Sato, K., Sano, Y., Yamada, K., Miura, W., & Takayasu, H. (2015). Rumor diffusion and convergence during the 3.11 earthquake: a Twitter case study. *PLoS one, 10*(4), e0121443

Tessler, N., Borger, H., Rave, E., Argaman, E., Kopel, D., Brook, A., Elkabets, E., & Wittenberg, L. (2019). Haifa fire restoration project–urban forest management: a case study. *International Journal of Wildland Fire, 28*(7), 485-494

Vieweg, S., Hughes, A. L., & Starbird, K. P. (2010). *Microblogging during two natural hazards events: What Twitter may contribute to situational awareness.* In Proceedings of the 2010 Annual Conference on Human Factors in Computing Systems Atlanta, Georgia, 1079–1088

White, C. (2010). *Social media and meta-networks for crisis mapping: Collaboratively building spatial data for situation awareness in disaster response and recovery management.* In Specialist Meeting Spatio-Temporal Constraints on Social Networks, University of California, Santa Barbara, Center for Spatial Studies. Retrieved from http://www.ncgia.ucsb.edu/projects/spatio-temporal/docs/White-position.,

Zhang, Q., Jin, P., Lin, S., & Yue, L. (2011). *Extracting focused locations for web pages.* In, Springer, 76-89