# NARMADA: Need and Available Resource Managing Assistant for Disasters and Adversities


**Kaustubh Hiware*** Mercari, Inc Tokyo, Japan hiwarekaustubh@gmail.com

**Ritam Dutt*** Indian Institute of Technology Kharagpur, India ritam.dutt@gmail.com

**Sayan Sinha** Indian Institute of Technology Kharagpur, India sayan.sinha@iitkgp.ac.in

**Sohan Patro** Microsoft Corporation Redmond, USA sopatr@microsoft.com

**Kripabandhu Ghosh** Tata Consultancy Services Pune, India kripa.ghosh@gmail.com

**Saptarshi Ghosh** Indian Institute of Technology Kharagpur, India saptarshi@cse.iitkgp.ac.in



## Abstract

Although a lot of research has been done on utilising Online Social Media during disasters, there exists no system for a specific task that is critical in a post-disaster scenario – identifying resource-needs and resource-availabilities in the disaster-affected region, coupled with their subsequent matching. To this end, we present NARMADA, a semi-automated platform which leverages the crowd-sourced information from social media posts for assisting post-disaster relief coordination efforts. The system employs Natural Language Processing and Information Retrieval techniques for identifying resource-needs and resource-availabilities from microblogs, extracting resources from the posts, and also matching the needs to suitable availabilities. The system is thus capable of facilitating the judicious management of resources during post-disaster relief operations.


## 1 Introduction

In recent years, microblogging sites like Twitter and Weibo have played a pivotal role in gathering situational information during disasters or emergency scenarios such as earthquakes, epidemic outbreaks, floods, hurricanes, and so on (Imran et al., 2015; Nazer et al., 2017; Li et al., 2017). Specifically, there are two types of information which are considered useful (or 'actionable') by rescue workers for assisting post-disaster relief operations.[1] These include (i) **Resource needs** that talk about the requirement of a specific resource (such as food, water, shelter) and (ii) **Resource availabilities** that talk about the availability of a specific resource in the region. Some examples of tweets that inform about resource-needs and resource-availabilities, taken from a dataset of tweets related to the 2015 Nepal earthquake, are shown in Table 1. We refer to such tweets as 'needs' and 'availabilities' henceforth.

| Needs (excerpts) | Availabilities (excerpts) |
|---|---|
| Mobile phones are not working, no electricity, no **water** in #Thamel, #Nepalquake | Please contact for drinking free service **water** specially for Earthquake Victim. Sanjay Limbu [mobile num] |
| Over 1400 killed. Many Trapped. **Medical Supplies** Requested. | 20,000 RSS personnel with **medical supplies** and other help the first to reach earthquake damaged zones in #Nepal |
| Nepal earthquake: thousands in need of **shelter** in country little able to cope [url] | can anyone we know pick the 2000 second hand **tents** from Sunauli and distribute it to the people in need in Nepal? |

Table 1: Examples of tweets stating resource-needs and tweets stating corresponding matching resource-availabilities, from a dataset of tweets on 2015 Nepal earthquake. The common resources for each pair are shown in boldface (table reproduced from our prior work (Dutt et al., 2019)).

The two major practical challenges faced in this regard include (i) automated *identification* of need and availability posts from social media sites such as Twitter and (ii) automated *matching* of the appropriate needs and availabilities. There have been prior works which have tried to address each of these challenges separately. However, to the best of our knowledge, there exists no system that integrates the two tasks of identifying needs and availabilities and their subsequent matching.

In this work, we present NARMADA (**N**eed and **A**vailable **R**esource **M**anaging **A**ssistant for **D**isasters and **A**dversities), a unified platform for the coordination of relief efforts during disasters by managing the resources that are needed and/or available in the disaster-affected region. NARMADA is designed to be a **semi-automated sys-**

*Equal Contribution

[1]We discussed with relief workers from 'Doctors For You' (http://doctorsforyou.org/) and SPADE (http://www.spadeindia.org/).




*Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 15–24
Vancouver, Canada, July 30 - August 4, 2017. ©2017 Association for Computational Linguistics
https://doi.org/10.18653/v1/P17

**tem** to ensure supervision and accountability.

In this paper, we describe the Natural Language Processing and Information Retrieval techniques used in NARMADA for the following tasks – (i) identifying resource-needs and resource-availabilities from microblogs, (ii) extracting resource names and other critical information from the posts (e.g., where the resource is needed, the quantity that is needed/available), and (iii) matching the needs to suitable availabilities. The system can be accessed from `https://osm-dm-kgp.github.io/Narmada/`. Although the system is currently applied over tweets only, NARMADA can also seamlessly integrate information from other sources, as well as enable users to add new information as they deem fit. We believe that the use of this system during a real-time disaster event will help in expediting relief operations.

Our work makes the following contributions.

1) We leverage contextual word embeddings to develop supervised models for automated classification of tweets that inform about need or availability of a resource.

2) We automate the process of categorising the type of resource present in needs and availabilities into food, health, shelter or logistics. This helps us to identify covert information present in tweets.

3) We deploy NARMADA that leverages NLP and IR techniques to identify resource needs and availabilities from microblogs, extract relevant information, and subsequently match needs to suitable availabilities. We believe that such a system would assist in post-disaster relief operations.

## 2 Related Work

There has been a lot of recent work on utilising Online Social Media (OSM) to facilitate post-disaster relief operations – see (Imran et al., 2015; Nazer et al., 2017; Li et al., 2017) for some recent surveys on this topic. For instance, there have been works on classifying situational and non-situational information (Rudra et al., 2015, 2018), location inferencing from social media posts during disasters (Karimzadeh et al., 2013; Lingad et al., 2013; Paule et al., 2018; Dutt et al., 2018; Kumar and Singh, 2019), early detection of rumours from social media posts (Mondal et al., 2018), emergency information diffusion on social media during crises (Kim et al., 2018), event detection (Hasan et al., 2018), extraction of event-specific informative tweets during disaster (Laylavi et al., 2017) and so on. Tweets specific to particular disasters have been studied in (Gautam et al., 2019), along with their categorisation. Certain other works have focused on the classification of such tweets by determining the probability of them being re-shared in Twitter (Neppalli et al., 2019). A comparison of various learning-based methods has also been recently conducted in (Assery et al., 2019).

Automated retrieval of needs and availabilities have been attempted by employing regular expressions (Purohit et al., 2013), pattern-matching techniques (Temnikova et al., 2015), language models (Basu et al., 2017), and neural IR methods such as word and character embeddings (Basu et al., 2017; Khosla et al., 2017). Likewise, there has been prior research on the automated matching of the needs and availabilities using tf-idf similarity (Purohit et al., 2013) and our prior works (Basu et al., 2018; Dutt et al., 2019) that used word-embeddings for the task. However, no prior work has attempted end-to-end identification and matching of needs and availabilities, which we attempt in this work.

Some information systems have also been implemented for disaster situations such as AIDR (AID, 2015) and Ushahidi (Ush, 2008) which employs crowd-sourced information using social media to assist disaster operations. To our knowledge, none of the existing systems have attempted the specific tasks in this work – identification and matching of resource-needs and resource-availabilities.

## 3 Dataset

We reuse the dataset made available by our prior works (Khosla et al., 2017; Basu et al., 2018; Dutt et al., 2019) which comprises tweets posted during two disaster events i.e. (i) the earthquake in Nepal in April, 2015 [2], and (ii) the earthquake in central Italy in August, 2016 [3]. Henceforth, we refer to the scenarios as Nepal-quake and Italy-quake.

The tweets were collected using the Twitter Search API[4] with the queries 'nepal quake' and 'italy quake'. The dataset consists of only English tweets since it was observed that most tweets are posted in English to enable rapid communication between international agencies and the local population.

[2] `https://en.wikipedia.org/wiki/April_2015_Nepal_earthquake`

[3] `https://en.wikipedia.org/wiki/August_2016_Central_Italy_earthquake`

[4] `https://dev.twitter.com/rest/public/search`

Removing duplicates and near-duplicates yielded a corpus of *50,068* tweets for Nepal-quake and *70,487* tweets for Italy-quake. However, the number of tweets that inform about needs and availabilities is very low – there are 499 and 1333 need and availability tweets for the Nepal-quake dataset. Likewise, the Italy-quake had only 177 needs and 233 availabilities (see (Dutt et al., 2019) for more details).

## 4 Methodology

In this section, we describe the methodologies that are incorporated within the NARMADA system. The overarching goal of the system is to facilitate post-disaster relief coordination efforts using the vast information available on social media. To that end, it performs three essential tasks – (i) identifying needs and availabilities, (ii) extracting actionable information from the need and availability tweets, and (iii) matching appropriate needs and availabilities. NARMADA is designed to execute each of the above three tasks in an automated fashion. We elaborate on the specific methodology involved for each of these sub-tasks in the ensuing subsection. However, prior to each of these tasks, we perform pre-processing on the tweet text as follows.

**Pre-processing tweets:** We employed standard pre-processing techniques on the tweet text to remove URLs (but not email ids), mentions, characters like brackets, 'RT', and other non-ASCII characters like #, &, ellipses and Unicode characters corresponding to emojis. We also segmented CamelCase words and joint alphanumeric terms like 'Nepal2015' into distinct terms ('Nepal' and '2015'). However, we did *not* perform case-folding or stemming on the tweet-text to enable subsequent detection of proper nouns (explained below).

### 4.1 Identifying needs and availabilities

Identifying needs and availabilities is challenging since they account for only $\approx 3.64\%$ and $\approx 0.58\%$ of the entire Nepal-quake and Italy-quake datasets, respectively. Prior works have approached this problem as a retrieval task using a wide array of techniques such as regular-expressions (Purohit et al., 2013), pattern-matching (Temnikova et al., 2015), language models (Basu et al., 2017), and recently neural IR techniques such as word and character embeddings (Basu et al., 2017; Khosla et al., 2017; Basu et al., 2019).

To enable the real-time deployment, a system needs to filter out tweets on an individual basis. To that end, we decided to adopt a supervised approach for classifying a tweet as 'need', or as 'availability' or as 'others' (i.e., a three-class classification problem). We experimented with different neural architectures for both in-domain and cross-domain classification. In-domain classification implies that the model is trained and tested on tweets related to the same disaster event. On the other hand, cross-domain classification involves training on tweets related to one event (say 'Nepal-quake') and evaluating on tweets related to another event ('Italy-quake') (Basu et al., 2019).

**Baseline methods:** Convolutional Neural Networks (CNN) have been found to work well in the classification of disaster-related tweets (Caragea et al., 2016; Nguyen et al., 2017). Hence we use the CNN of (Kim, 2014) as a baseline model. We operate on 300-dimensional word-embeddings and fix the feature maps to 100 dimensions. We implement convolutional filters with kernel-size 3, 4, and 5 respectively, with stride 1, and non-linear ReLU activation units. Finally, we apply max-pooling before passing it through a fully-connected layer and softmax with negative log-likelihood (NLL) loss. We experiment with randomly initialized embeddings as well as different kinds of pre-trained embeddings – Glove(Pennington et al., 2014)[5], word2vec (Mikolov et al., 2013) [6], fasttetxt embeddings (Bojanowski et al., 2017)[7] and CrisisNLP embeddings (Imran et al., 2016) trained on tweets posted during many disaster events.

**Proposed model:** We propose to use a pre-trained BERT model (Devlin et al., 2018) (bert-base-uncased) to represent a tweet as a 768-dimensional embedding. We pass the represented tweet through a fully connected layer which classifies it into the aforementioned three categories. Using BERT pre-trained embeddings helps us in two ways. Firstly, the BERT model itself remains a part of the entire end-to-end system; hence it gets fine-tuned while training. Moreover, BERT uses multiple bidirectional self-attention modules which helps capture contextual information.

**In-domain classification:** Table 2 notes the per-

[5]https://nlp.stanford.edu/projects/glove/

[6]https://code.google.com/p/word2vec/

[7]https://fasttext.cc/docs/en/english-vectors.html

| Methodology | Nepal-quake Prec | Nepal-quake Rec | Nepal-quake F1 | Italy-quake Prec | Italy-quake Rec | Italy-quake F1 |
|---|---|---|---|---|---|---|
| CNN + random | 0.803 | 0.612 | 0.681 | 0.926 | 0.552 | 0.637 |
| CNN + Glove | 0.790 | 0.668 | 0.716 | 0.846 | 0.680 | 0.727 |
| CNN + Word2vec | 0.796 | 0.660 | 0.712 | 0.847 | 0.644 | 0.709 |
| CNN + Fasttext | 0.771 | 0.628 | 0.683 | 0.870 | 0.640 | 0.703 |
| CNN + CrisisNLP | 0.767 | 0.634 | 0.682 | 0.734 | 0.585 | 0.635 |
| **BERT** (proposed for F1) | 0.786 | 0.866 | 0.823 | 0.856 | 0.722 | 0.779 |
| **BERT** (proposed for Rec) | 0.791 | 0.872 | **0.828** | 0.843 | 0.810 | **0.826** |

Table 2: Performance of the neural architectures for in-domain classification of tweets into three classes – needs, availabilities, and others. Best F1-scores in boldface.

formances of the various classification models in *in-domain settings*, averaged over both the classes *needs* and *availabilities*. For each of the two datasets, we consider 20% (randomly sampled) of the labelled data as the test set, 70% of the rest as the training set, and the rest 10% was used as the validation set. We report the Precision, Recall and F1-score on the test set as the evaluation measures. We consider F1-score as the primary score since it incorporates both Precision and Recall. The proposed BERT model outperforms all other models in terms of F1-score.

We trained two versions of our proposed BERT model – (i) one version was trained to optimise the F1-score (our primary measure) on the validation set, and (ii) the second version was trained to optimise the *Recall* on the validation set. We specifically tried one version to optimise Recall, since it is usually considered important to identify all needs and availabilities in a disaster situation. As seen in Table 2, both versions of the model achieved comparable performance on the Nepal-quake dataset (F1-scores of 0.823 and 0.828). But the version trained for optimizing Recall achieved substantially higher performance on the Italy-quake dataset where needs and availabilities are much sparser. This improved performance justifies our decision of focusing on improving Recall.

**Cross-domain classification:** In a cross-domain setting, the model is trained on tweets of one event and then evaluated on tweets of the other event. We compare the performance of the BERT model against the best-supervised model ('Best-SM') of (Basu et al., 2019), which is a CNN classifier initialised with CrisisNLP embeddings. Table 3 shows results when the models are trained on Italy-quake tweets and tested on Nepal-quake tweets. Similarly, Table 4 shows the opposite setting, i.e., the models are trained on Nepal-quake tweets and tested on Italy-quake tweets. In both the cases, we use the BERT model optimised for F1-score, as described above. Even for cross-domain performance, we see that the BERT model outperforms the CNN-based baseline of (Basu et al., 2019).

| Method | P@100 | R@100 | F1@100 |
|---|---|---|---|
| **Needs** | | | |
| Best-SM (Basu et al., 2019) | 0.443 | 0.044 | 0.080 |
| **BERT** (proposed) | 0.320 | 0.066 | **0.110** |
| **Availabilities** | | | |
| Best-SM (Basu et al., 2019) | 0.533 | 0.019 | 0.037 |
| **BERT** (proposed) | 0.500 | 0.038 | **0.070** |

Table 3: Performance of the neural architectures when trained on Italy-quake and tested on Nepal-quake. Best F1-scores in boldface.

| Method | P@100 | R@100 | F1@100 |
|---|---|---|---|
| **Needs** | | | |
| Best-SM (Basu et al., 2019) | 0.198 | 0.056 | 0.087 |
| **BERT** (proposed) | 0.32 | 0.184 | **0.234** |
| **Availabilities** | | | |
| Best-SM (Basu et al., 2019) | 0.216 | 0.046 | 0.076 |
| **BERT** (proposed) | 0.28 | 0.121 | **0.168** |

Table 4: Performance of the neural architectures when trained on Nepal-quake and tested on Italy-quake. Best F1-scores in boldface.

### 4.2 Extracting relevant fields from needs and availabilities

Prior discussions with relief workers helped us identify the following five fields that are deemed relevant in coordinating the relief efforts, namely: (i) resource – which items are needed/available, (ii) quantity – how much of each resource is needed/available, (iii) location – where is the resource needed/available, (iv) source – who needs the resource or who is offering, and (v) contact – how to contact the said source.

We adapt the unsupervised methodology of our prior work (Dutt et al., 2019) to extract the relevant fields from needs and availabilities. We sought to incorporate this technique due to the paucity of labelled instances which discourages a supervised machine learning approach (and because gathering many labelled instances is difficult in a disaster scenario). Moreover, the unsupervised approach was shown to be generalizable across several datasets (Dutt et al., 2019). We describe the adapted methodology in this section.

**Unsupervised resource extraction:** We start by giving a brief description of the methodology in (Dutt et al., 2019). We perform dependency parsing on the text to obtain a Directed Acyclic

| Tweet Text | Resource |
|---|---|
| villagers in the remote community of ghyangphedi fear *hunger* and *#starvation* | food |
| earthquake victims *sleeping outside* in nepal | shelter |
| people are *shivering in the cold* | shelter |
| free calls to italy in the wake of earthquake | logistics |

Table 5: Examples of covert tweets and the corresponding resource class assigned to the tweet by our BERT-based resource classifier.

| Dataset | Precision | Recall | F1-score |
|---|---|---|---|
| Nepal-quake | 0.838 | 0.882 | 0.843 |
| Italy-quake | 0.825 | 0.858 | 0.823 |

Table 6: Performance of the multi-label BERT-based resource classifier on in-domain classification.

| **Method** | **P@100** | **R@100** | **F1@100** |
|---|---|---|---|
| **Nepal-quake** | | | |
| USM (Dutt et al., 2019) | 0.623 | 0.833 | **0.685** |
| **BERT** (trained on Italy) | 0.484 | 0.670 | 0.522 |
| **BERT** (trained on Italy + 5% Nepal) | 0.636 | 0.834 | 0.680 |
| **Italy-quake** | | | |
| USM (Dutt et al., 2019) | 0.487 | 0.595 | 0.516 |
| **BERT** (trained on Nepal) | 0.798 | 0.862 | **0.808** |

Table 7: Comparing the BERT-based resource classifier with the unsupervised methodology (USM) of (Dutt et al., 2019) in cross-domain setting. Best F1-scores in boldface.

Graph (DAG). We compile an initial list of headwords (*send, need, donate*, etc.) which consists of the verbs in the query-set and the ROOT word of the DAG. We have identified specific characteristics of the child nodes of the headwords that enable us to label the node as a potential resource.

For example, if a word $w$ is tagged as a NOUN and is the direct object of the 'donates', $w$ can be expected to be a potential resource. We have also identified dependency rules, that increases the list of head-words to improve our recall. We thus obtain a list of potential resources after dependency parsing. We then verify these potential resources by checking for the semantic similarity of the extracted words with a pre-compiled list of resources commonly used during disasters. The resource list is obtained from several reputed sources like UNOCHA[8], UNHCR[9] and WHO[10]. This pre-compiled list also enables us to categorise the resources into four classes namely *food* (bottled water, biscuits, rice), *health* (blood, medicine, latrines), *shelter* (tents, blankets, tarpaulins), and *logistics* (electricity, helicopters, cash).

**Adapting the method to deal with covert tweets:** One of the limitations of the unsupervised methodology in (Dutt et al., 2019) is the inability to glean relevant information from *covert tweets* where the resource needed/available is not mentioned explicitly. We illustrate instances of such covert tweets in Table 5. Since the resource name is not explicitly stated in the tweet-text, the methodology in (Dutt et al., 2019) cannot identify the resources for such tweets.

To circumvent this problem, we again use the pre-trained BERT model (Devlin et al., 2018) to encode a tweet. We pass this representation through a linear layer and perform multi-label classification into the aforementioned four categories, i,e. food, health, shelter and logistics. We use multi-label classification since a particular tweet can mention multiple resources. This adaptation helps the methodology to correctly classify many of the covert tweets, as demonstrated in Table 5 (the second column shows the resource-class that is assigned by our methodology).

We report the *in-domain* classification performance of our BERT-based resource classifier for the Nepal-quake and Italy-quake datasets in Table 6. We test on 20% of the data (sampled randomly) and train on the remaining 70% while using 10% for validation. We optimise the model with the highest macro F1-score on the validation set.

Next, we compare the performance of the proposed BERT-based resource classifier with that of the unsupervised methodology of (Dutt et al., 2019) (which we refer to as 'USM'). To ensure a fair comparison, we perform this comparison in a *cross-domain* setting wherein we train the supervised model on one dataset (e.g., Nepal-quake) and evaluate on another (e.g., Italy-quake). We present the results of this comparison in Table 7.

We observe from Table 7 that the BERT resource classifier trained on Nepal-quake significantly outperforms USM over the Italy-quake dataset (F1-score of 0.808 for the BERT method and 0.516 for USM). In contrast, the BERT resource classifier when trained on Italy-quake yielded significantly poorer results on Nepal-quake dataset than USM. However, training only on an additional 5% of labelled instances of the Nepal-quake dataset, demonstrated comparative performance (F1-score of 0.680 for the BERT method and 0.685 for USM). The reason for these performances is as follows. The Italy-quake dataset does not contain mention of several amenities that are heavily prevalent in

[8]https://www.unocha.org/
[9]https://www.unhcr.org/
[10]https://www.who.int/

| Tweet text (excerpts) | Resource | Location | Quantity | Source | Contact |
|---|---|---|---|---|---|
| Urgent need of analgesic,antibiotics, betadiene, swabs in kathmandu!! Call for help 98XXX-XXXXX #earthquake #Nepal #KTM (N) | analgesic, antibiotics, betadiene, swabs | kathmandu, ktm, nepal | | | 98XXX-XXXXX |
| India sends 39 #NDRF team, 2 dogs and 3 tonnes equipment to Nepal Army for rescue operations: Indian Embassy in #Nepal (A) | NDRF team, dogs, | nepal | dogs - 2, NDRF team - 39 | India | |
| Visiting Sindhupalchok devastating earthquake highly affected district . Delivery Women in a tent . No water no toilet (N) | tent, delivery women, water | Sindhupalchok | | | |
| Rajasthan Seva Samiti donates more than 800 tents to Nepal Earthquake victims (A) | tents | | tents-800 | Rajasthan Seva Samiti | |

Table 8: Examples of information extracted from need (N) and availability (A) tweets by the methodologies proposed in this work. Red colour indicates wrongly extracted information.

the Nepal-quake dataset, but *not* vice-versa. This difference is mainly because the Italy earthquake was a comparatively mild one in a developed region, and hence not many resources were needed; in contrast, the Nepal earthquake was a severe one in a developing region, and a lot of resources were needed in Nepal. Hence the Nepal-quake dataset contains mention of far more varied resources, as compared to the Italy-quake dataset.

Thus, including the BERT-based resource classifier in addition to the unsupervised methodology improves resource extraction performance, and also lends generalisability across different datasets.

**Extracting Locations:** We extract geographical locations from the tweet text using the methodology in our prior work (Dutt et al., 2018). First, we apply several unsupervised techniques to extract a set of potential locations. These techniques include (i) segmentating hashtags, (ii) disambiguating proper nouns from parse trees, (iii) identifying phrases with regex matches, (iv) dependency parsing to locate nouns close from words in query-set in the DAG, and (v) employing pre-trained Named Entity Recognizers [11] to identify words tagged as geographical location. Next, we verify these potential locations using a gazetteer. We consider those locations to be valid only if their geospatial coordinates lie within the boundary of the affected region (e.g., Nepal or Italy). We used two gazetteers namely Geonames [12] and Open Street Map [13] to identify locations with varying levels of granularity (as detailed in (Dutt et al., 2018)).

**Extracting the source:** We consider as viable sources two types of words – (i) proper nouns that are tagged as organisations, persons or geographical locations by a Named Entity Recognizer, and (ii) proper nouns that are child nodes of dependency parsing – provided they have not been identified previously as 'location' or 'resources' during the verification phase. See our prior work (Dutt et al., 2019) for details of the methodology.

**Extracting Quantity:** For each resource extracted, we identify whether it is preceded by a numeric token. The numeric token may be the orthographic notation of a number (e.g., '100') or may semantically represent a number (e.g., 'hundred'). We assign the numeric token as the quantity of the particular resource.

**Extracting Contact:** We use regular expressions to identify contacts corresponding to email-ids and phone numbers.

The performance of our information extraction methods (in terms of precision, recall and F1-score) was similar to what is presented in (Dutt et al., 2019). In our experiments, we obtained F1-scores of 0.89, 0.91, 0.76, 0.58 and 1.00 for identifying Resources, Location, Quantity, Source and Contact respectively, for need-tweets. Likewise, the F1-scores for availability-tweets were 0.85, 0.85, 0.84, 0.65 and 1.00 respectively. Table 8 shows some examples of the fields extracted by our methods from some need-tweets and availability-tweets.

### 4.3 Matching needs and availabilities

We propose a fast and real-time algorithm for matching needs and availabilities based on **proportion of common resources**. Specifically, for a given need-tweet, we compute the match with a particular availability-tweet as the fraction of the resources extracted from the need-tweet, that are also present in the availability-tweet. For the given

[11] We use the inbuilt NER tool of SpaCy (https://spacy.io/)

[12] http://www.geonames.org/

[13] http:420//geocoder.readthedocs.io/providers/OpenStreetMap.html

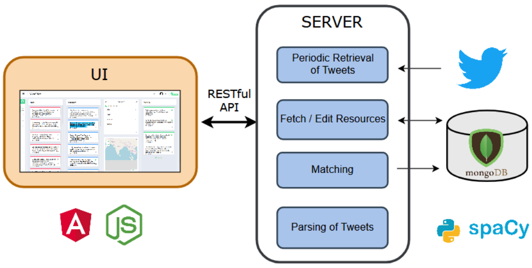


Figure 1: NARMADA's architecture overview

need tweet, availability-tweets are ranked in decreasing order of the fraction of common resources (ties resolved arbitrarily).

We also experiment with some baseline methodologies, namely using common nouns (Basu et al., 2018), tf-idf vectors of the tweet text (Purohit et al., 2013) and local word embeddings of the tweet (Basu et al., 2018). Our methodology (based on the proportion of common resources) obtains an F1-score of 0.84 for Nepal-quake and an F1-score of 0.87 for Italy-quake dataset respectively, which is competitive with the performance of the baselines.

This section described the NLP and IR techniques used in NARMADA. The next section describes the system architecture.

## 5 System Architecture

The high-level system architecture for NARMADA is shown in Figure 1. The system can be accessed from `https://osm-dm-kgp.github.io/Narmada/`, where further details and a demonstration video are also provided. NARMADA is designed and built for the Web, thus not restricting it to any particular operating system or browser type, allowing cross-platform (desktop/mobile) functionality.

### 5.1 User Interface

The user interface has been designed in Typescript using Angular, a popular web-application framework. ngx-admin[14] was used as a boilerplate for front-end components. The interface has been designed to be intuitive, yet presenting as much information as possible without overcrowding. A detailed note is available at `https://osm-dm-kgp.github.io/Narmada/`.

The user interface comprises a dashboard (shown in Figure 2) that acts as a landing page. Besides providing an initial view of active needs and availabilities (at the present point of time), it displays matched resources. The user is provided with various options to make it easy to search and locate resources as well as highlight items as deemed necessary.

An alternate section is available where users can enter new needs/availabilities manually. The class labels of the information are detected automatically, but the user is allowed to modify the same. Another section for "Completed matches" is to be used for logging completed or exhausted needs and resources. A user manual is also attached to the UI.

### 5.2 Server

The major services provided by the backend server include classification and categorisation of the tweets in the system. It also provides support for the addition of new information and their automatic categorisation. Facilities have been provided for marking resources once their need is fulfilled or the availability gets exhausted.

The server side uses NodeJS framework and is written in Javascript. Nginx is used as an HTTP server to make the frontend accessible to the public. However, the NLP-related extraction tasks are handled better in Python. The server partly uses a Flask-based Python backend, a micro web framework. The Flask server makes API calls to the deep learning classifiers, featuring BERT, which returns the output. The output is further reflected in the frontend. The server sends information requested by the user interface via *RESTful API*, which supports cached responses on the frontend and enables the system to be scalable, thus allowing more users to use this service. API endpoints are publicly available, which would allow programmatic access to the server's functionalities (see `https://osm-dm-kgp.github.io/Narmada/`).

## 6 Discussion

NARMADA intends to assist in crossing the initial barrier in identifying and matching needs and availabilities from social media during the occurrence of a disaster. In practice, it becomes necessary for other service providers to be triggered in order to make sure that the needs are addressed, by proper collection, transportation and provisioning of the matched resources deemed to be available. For instance, the needs and availabilities could be marked on a map, with each type of resource be-

[14] `https://github.com/akveo/ngx-admin`

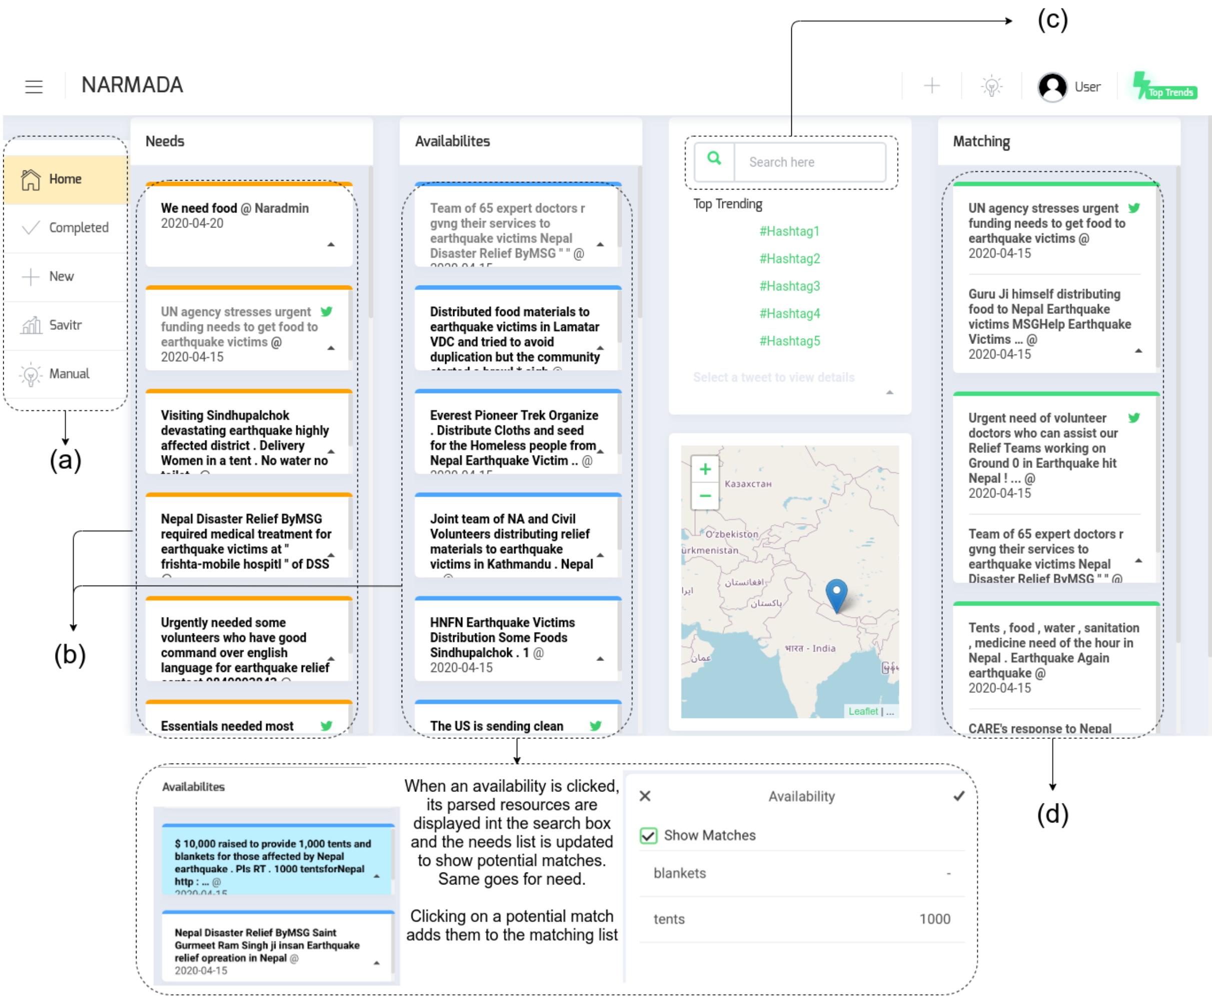


Figure 2: Dashboard of NARMADA – (a) **Navigation Buttons**. (b) **Needs and Availabilities List**: tweets are displayed in reverse chronological order; gray tweet: already matched; black tweet: unmatched; each tweet contains a notch at the bottom-right corner, clicking on which reveals more details. (c) **Search Box**: when a query is entered, the needs and availabilities containing the query-phrase are displayed. (d) **Matching List**: displays the matched needs and availabilities; clicking a matching displays its resources, and gives the user an option to mark it as completed.

ing represented with a different symbol, making it easy to physically locate them. Local volunteers might be provided with a mobile app to help them find nearby needs and availabilities. Misinformation in twitter is common (Bal et al., 2020). The volunteers would also need a facility to confirm that the posted needs and availabilities are indeed genuine, concerning various parameters such as quantity (since at times of disasters, needs may be exaggerated).

## 7 Conclusion and Future Work

We proposed a system NARMADA for resource management during a disaster situation. Though the system is developed to work across posts from various social media platform, this research focused on data from Twitter. The real-time nature and easy access to large volumes of information provided by Twitter have made it a lucrative choice for disaster analytics.

Currently, the system allows all users to perform any action on the system. One future task would be to implement a login system that would allow different access-levels to different users. For instance, a visitor would be able to only view and query information, a volunteer would be able to add new resources, mark a need as matched, etc., while a system administrator would have rights to undo all actions of all users, etc. The current system does not allow multiple volunteers to communicate within the platform over a resource, which we wish to incorporate in the future. We also plan to incorporate support for vernacular languages, provided the requisite tools are available.

## References

2008. Ushahidi. `https://www.ushahidi.com/`.

2015. Aidr (artificial intelligence for disaster response). `http://aidr.qcri.org/`.

Nasser Assery, Yuan Xiaohong, Sultan Almalki, Roy Kaushik, and Qu Xiuli. 2019. Comparing learning-based methods for identifying disaster-related tweets. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 1829–1836. IEEE.

Rakesh Bal, Sayan Sinha, Swastika Dutta, Risabh Joshi, Sayan Ghosh, and Ritam Dutt. 2020. Analysing the extent of misinformation in cancer related tweets. *arXiv preprint arXiv:2003.13657*.

Moumita Basu, Kripabandhu Ghosh, Somenath Das, Ratnadeep Dey, Somprakash Bandyopadhyay, and Saptarshi Ghosh. 2017. Identifying post-disaster resource needs and availabilities from microblogs. In *Proc. ASONAM*.

Moumita Basu, Anurag Shandilya, Kripabandhu Ghosh, and Saptarshi Ghosh. 2018. Automatic matching of resource needs and availabilities in microblogs for post-disaster relief. In *Comp. Proc. WWW 2018 2018*, pages 25–26.

Moumita Basu, Anurag Shandilya, Prannay Khosla, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. Extracting resource needs and availabilities from microblogs for aiding post-disaster relief operations. *IEEE Transactions on Computational Social Systems*, 6(3):604–618.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

Cornelia Caragea, Adrian Silvescu, and Andrea H Tapia. 2016. Identifying informative messages in disaster events using convolutional neural networks. In *International Conference on Information Systems for Crisis Response and Management*, pages 137–147.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Ritam Dutt, Moumita Basu, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. Utilizing microblogs for assisting post-disaster relief operations via matching resource needs and availabilities. *Information Processing & Management*, 56(5):1680–1697.

Ritam Dutt, Kaustubh Hiware, Avijit Ghosh, and Rameshwar Bhaskaran. 2018. Savitr: A system for real-time location extraction from microblogs during emergencies. In *Proc. WWW Workshop SMERP*.

Akash Kumar Gautam, Luv Misra, Ajit Kumar, Kush Misra, Shashwat Aggarwal, and Rajiv Ratn Shah. 2019. Multimodal analysis of disaster tweets. In *2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM)*, pages 94–103. IEEE.

Mahmud Hasan, Mehmet A. Orgun, and Rolf Schwitter. 2018. Real-time event detection from the Twitter data stream using the TwitterNews+ Framework. *Information Processing & Management*. Online: https://doi.org/10.1016/j.ipm.2018.03.001.

Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. Processing Social Media Messages in Mass Emergency: A Survey. *ACM Computing Surveys*, 47(4):67:1–67:38.

Muhammad Imran, Prasenjit Mitra, and Carlos Castillo. 2016. Twitter as a lifeline: Human-annotated twitter corpora for nlp of crisis-related messages. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).

Morteza Karimzadeh, Wenyi Huang, Siddhartha Banerjee, Jan Oliver Wallgrün, Frank Hardisty, Scott Pezanowski, Prasenjit Mitra, and Alan M. MacEachren. 2013. Geotxt: A web api to leverage place references in text. In *Proceedings of the 7th Workshop on Geographic Information Retrieval*, pages 72–73.

Prannay Khosla, Moumita Basu, Kripabandhu Ghosh, and Saptarshi Ghosh. 2017. Microblog retrieval for post-disaster relief: Applying and comparing neural IR models. *CoRR*, abs/1707.06112.

Jooho Kim, Juhee Bae, and Makarand Hastak. 2018. Emergency information diffusion on online social media during storm Cindy in U.S. *International Journal of Information Management*, 40:153 – 165.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.

Abhinav Kumar and Jyoti Prakash Singh. 2019. Location reference identification from tweets during emergencies: A deep learning approach. *International journal of disaster risk reduction*, 33:365–375.

Farhad Laylavi, Abbas Rajabifard, and Mohsen Kalantari. 2017. Event relatedness assessment of Twitter messages for emergency response. *Information Processing & Management*, 53:266–280.

Tao Li, Ning Xie, Chunqiu Zeng, Wubai Zhou, Li Zheng, Yexi Jiang, Yimin Yang, Hsin-Yu Ha, Wei Xue, Yue Huang, Shu-Ching Chen, Jainendra Navlakha, and S. S. Iyengar. 2017. Data-Driven Techniques in Disaster Information Management. *ACM Comput. Surv.*, 50(1):1:1–1:45.

John Lingad, Sarvnaz Karimi, and Jie Yin. 2013. Location extraction from disaster-related microblogs. In *Proceedings of the 22nd international conference on world wide web*, pages 1017–1020. ACM.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.

Tamal Mondal, Prithviraj Pramanik, Indrajit Bhattacharya, Naiwrita Boral, and Saptarshi Ghosh. 2018. Analysis and Early Detection of Rumors in a Post Disaster Scenario. *Information Systems Frontiers*, 20(5).

Tahora H. Nazer, Guoliang Xue, Yusheng Ji, and Huan Liu. 2017. Intelligent disaster response via social media analysis a survey. *SIGKDD Explor. Newsl.*, 19(1):46–59.

Venkata Kishore Neppalli, Cornelia Caragea, Doina Caragea, Murilo Cerqueira Medeiros, Andrea H Tapia, and Shane E Halse. 2019. Predicting tweet retweetability during hurricane disasters. In *Emergency and Disaster Management: Concepts, Methodologies, Tools, and Applications*, pages 1277–1298. IGI Global.

Dat Tien Nguyen, Kamla Al-Mannai, Shafiq R Joty, Hassan Sajjad, Muhammad Imran, and Prasenjit Mitra. 2017. Robust classification of crisis-related data on social networks using convolutional neural networks. In *ICWSM*, pages 632–635.

Jorge David Gonzalez Paule, Yeran Sun, and Yashar Moshfeghi. 2018. On fine-grained geolocalisation of tweets and real-time traffic incident detection. *Information Processing & Management*.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proc. EMNLP*.

Hemant Purohit, Carlos Castillo, Fernando Diaz, Amit Sheth, and Patrick Meier. 2013. Emergency-relief coordination on social media: Automatically matching resource requests and offers. *First Monday*, 19(1).

Koustav Rudra, Niloy Ganguly, Pawan Goyal, and Saptarshi Ghosh. 2018. Extracting and summarizing situational information from the twitter social media during disasters. *ACM Trans. Web*, 12(3):17:1–17:35.

Koustav Rudra, Subham Ghosh, Pawan Goyal, Niloy Ganguly, and Saptarshi Ghosh. 2015. Extracting situational information from microblogs during disaster events: A classification-summarization approach. In *Proc. CIKM*.

Irina Temnikova, Carlos Castillo, and Sarah Vieweg. 2015. EMTerms 1.0: A Terminological Resource for Crisis Tweets. In *Proc. ISCRAM*.