

Cross-Lingual and Cross-Domain Crisis Classification for Low-Resource Scenarios

Cinthia Sánchez^{1,2}, Hernan Sarmiento^{1,2}, Andres Abeliuk^{1,3}, Jorge Pérez⁴, Barbara Poblete^{1,2,3}

¹Department of Computer Science, University of Chile, Santiago, Chile

²Millennium Institute for Foundational Research on Data (IMFD), Santiago, Chile

³National Center for Artificial Intelligence (CENIA), Santiago, Chile

⁴Cero.ai, Santiago, Chile

{cisanche, hsarmien, aabeliuk, bpoblete}@dcc.uchile.cl, perez@cero.ai

Abstract

Social media data has emerged as a useful source of timely information about real-world crisis events. One of the main tasks related to the use of social media for disaster management is the automatic identification of crisis-related messages. Most of the studies on this topic have focused on the analysis of data for a particular type of event in a specific language. This limits the possibility of generalizing existing approaches because models cannot be directly applied to new types of events or other languages. In this work, we study the task of automatically classifying messages that are related to crisis events by leveraging cross-language and cross-domain labeled data. Our goal is to make use of labeled data from high-resource languages to classify messages from other (low-resource) languages and/or of new (previously unseen) types of crisis situations. For our study we consolidated from the literature a large unified dataset containing multiple crisis events and languages. Our empirical findings show that it is indeed possible to leverage data from crisis events in English to classify the same type of event in other languages, such as Spanish and Italian (80.0% F1-score). Furthermore, we achieve good performance for the cross-domain task (80.0% F1-score) in a cross-lingual setting. Overall, our work contributes to improving the data scarcity problem that is so important for multilingual crisis classification. In particular, mitigating *cold-start* situations in emergency events, when time is of essence.

1 Introduction

Social media has created new possibilities for outreach and communications with worldwide scope. It enables users to share information and opinions rapidly to millions of others in just seconds. The content published in these platforms is diverse in format (e.g., video, images, audio), but also diverse in terms of the users' backgrounds (e.g., language, geographic location, culture). This one-to-many broadcast of information has played a key role in the dissemination of breaking news and, in recent years, of emergency events (Palen and Anderson 2016). Across the different phases of crisis management, social media also plays a role in planning and training, collaborative problem solving and decision making, and information gathering (Chan 2013).

The microblogging platform Twitter has become one of the primary data sources for real-time crisis-event information. Its short-text messages, called *tweets*, have been widely studied to extract useful and timely knowledge for crisis management (Olteanu et al. 2014; Olteanu, Vieweg, and Castillo 2015; Cresci et al. 2015; Imran, Mitra, and Castillo 2016; Alam, Ofli, and Imran 2018). Accurate and well-timed information about a crisis allows emergency relief agencies to act quickly and effectively, thus reducing the negative impact on society (Graf et al. 2018). In this context, automatically identifying user messages related to crises becomes relevant. However, the noisy nature of social media and unreliable information quality are significant challenges in this field. In particular, online conversations are short, imprecise, and cover a wide range of topics, with little to no contextual information (Troudi et al. 2018). Furthermore, text message processing can be complex from a Natural Language Processing (NLP) perspective because they contain informal language, abbreviations, spelling mistakes, multilingual constructions, among others (Ghosh et al. 2018). All of these issues are further exacerbated when considering multilingual content (Torres and Vaca 2019).

As in many other supervised learning tasks, most research and models address English content, which affects the availability of resources across languages (Sánchez 2021; Khare et al. 2018; Alam et al. 2021b). In general, existing approaches require a significant amount of labeled data to learn effective models (Alam et al. 2021a). Often, labeled data from previous disasters is available, but in languages or crisis domains different from the target event (Li and Caragea 2020). This is problematic when we address multilingual and cross-domain scenarios where labeled data is scarce, leading to *cold-start* problems in early stages of a disaster (Lorini et al. 2019; Li and Caragea 2020). Furthermore, training individual models for all possible scenario combinations, in terms of languages and types of events, is computationally expensive (Khare et al. 2018). Hence, it becomes critical to create ways to leverage existing knowledge about crisis situations to address new or diverse types of events.

We address this problem by studying the task of classifying crisis messages across languages and types of crisis domains. Specifically, we focus on the binary classification task of tweets that are *related* and *unrelated* to crises. We consider this particular task to be relevant because of the

high volume of noisy data that is found in Twitter during crisis events. This noise consists mostly of unrelated content, such as spam and advertising, that can be found even when message retrieval is done using event-specific keywords (such as hashtags or locations associated with the crisis) (Olteanu, Vieweg, and Castillo 2015; Graf et al. 2018). Identifying related and unrelated messages is one of the first and foremost tasks required by emergency practitioners when using social media data.

We center our work in two main objectives: 1) transferring knowledge from one or more crisis domains (e.g., *earthquakes*, *hurricanes*, etc.) to other crisis domains, and 2) transferring knowledge from crises in one (high-resource) language to another (low-resource) language. To achieve this, we perform a comprehensive evaluation of transfer learning performance, using different data representations and transfer learning scenarios.

Our results show that it is possible to leverage crisis data from high-resource languages, such as English, to classify crisis events in other languages, like Spanish or Italian (80.0% F1-score). In addition, we observe that information from one or more crisis domains can contribute to the classification of events from a new domain. These results indicate that some characteristics of crises cross-cut different domains and languages. This will allow researchers and emergency practitioners to improve the reach of automated social network analysis during crises at a low cost.

Contributions. In summary, our contributions are:

1. We introduce a new multilingual and multi-domain crisis dataset, containing 53 crisis events and more than 160,000 messages. This dataset is the result of our work to unify 7 publicly available crisis datasets, achieved by using a systematic approach to consistently merge labels, and to add specific event and message metadata.
2. We show empirically that it is possible to perform transfer learning, using crisis data from high-resource languages, to classify data from other (low-resource) languages and from new crisis domains. This addresses the important *cold-start* problem in crisis classification for low-resource languages and new crisis domains.
3. We perform a novel comparative study of 7 state-of-the-art data models for cross-domain and cross-language transfer learning scenarios.

Note on reproducibility. All of the source code and data used in this article are publicly available at <https://github.com/cinthiasanchez/Crisis-Classification>.

Roadmap. In Section 2, we present the literature review focusing on crisis messages classification and transfer learning techniques. In Section 3, we detail our proposed experimental framework. In Section 4, we explain the construction of the unified dataset. In Section 5, we describe the experimental setup among different domains and languages, and in Section 6, we report the results. In Sections 7 and 8, we present the discussion and final comments of our work.

2 Related Work

In this section, we discuss relevant prior literature for our work. In particular, we focus on two areas: *classification of crisis-related messages* and *transfer learning*.

Classification of Crisis-Related Messages

One of the most prevalent tasks in this area is the binary classification of messages that are *related* vs. those that are *not related* to a crisis (Olteanu et al. 2014; Cobo, Parra, and Navón 2015; Alam, Joty, and Imran 2018; Torres and Vaca 2019). In the literature, the term *related* to a crisis event is usually interchangeable with the terms *relevant* or *informative*. First, a message is considered as *related* to a crisis when it refers implicitly or explicitly to a specific disaster (Kruspe, Kersten, and Klan 2020). A message considered *relevant* to a crisis is one that contains actual information pertaining to the event (Kruspe, Kersten, and Klan 2020; Cobo, Parra, and Navón 2015). It also conveys or reports information useful for crisis response (Alam, Joty, and Imran 2018). Finally, a message is considered *informative* when it contributes to a better understanding of the situation on the ground (Olteanu, Vieweg, and Castillo 2015). These terms have certain differences, especially in their degree of generalization and usefulness for humanitarian aid.

Olteanu et al. (2014) proposed a crisis lexicon for sampling and filtering crisis-related messages in English during different emergency events. Cobo, Parra, and Navón (2015) studied user and content based features to classify relevant tweets to an earthquake in Spanish (73.4% F1-score) using Random Forest. Alam, Joty, and Imran (2018) proposed a deep learning framework based on semi-supervised learning to classify relevant messages in English. They used two Twitter datasets, one of the Nepal earthquake (65.11% F1) and another of the Queensland floods (93.54% F1). Li et al. (2018) proposed a feature-based adaptation framework, which considers pre-trained and crisis-specific word embeddings, as well as sentence embeddings and different supervised classifiers. They evaluated two classification tasks in English reaching an average accuracy of 88.5%. They noted that crisis-specific embeddings were more suitable for more specific crisis-related tasks (*informative* vs *non-informative*), while the pre-trained embeddings were more suitable for more general tasks (*relevant* vs *non-relevant*). Firoj, Imran, and Ofli (2019) performed a comparative study among various algorithms used to classify crisis-related messages. Their results show competitive results between Support Vector Machines, Random Forest and Convolutional Neural Networks.

In view of the difficulty of comparing results, models and techniques in this research area, Alam et al. (2021b) developed a standard dataset based on existing data and provided train/dev/test partitions. In addition, the authors provided benchmark results on English messages for informative (binary) and humanitarian (multi-class) classification tasks using deep learning algorithms. Alam et al. (2021a) created a large-scale dataset of English-language tweets, which is composed of 19 disaster events that occurred between 2016 and 2019. The authors report the results of the classification

of humanitarian information using classical and deep learning algorithms. They achieved an average weighted F1 of 78.1% with the RoBERTa model.

Transfer Learning

Knowledge transfer or transfer learning involves two main concepts, *domain* and *task* (Pan and Yang 2009). A domain consists of a feature space and a marginal probability distribution, while a task consists of a label space and an objective predictive function. It aims to use knowledge from a source domain and learning task to improve the learning of the predictive function in the target domain, where source and target domains or tasks differ. The authors categorize it as inductive, transductive, and unsupervised, based on the settings between source and target. In transductive transfer learning, tasks are the same, but domains are different. It can have two variants (Pan and Yang 2009; Ruder 2019):

- Cross-lingual adaptation: domains have different feature spaces (e.g., documents written in two different languages).
- Domain adaptation: domains have different marginal probability distributions (e.g., documents discuss different topics).

Cross-Lingual Adaptation. This variant has been studied in different Twitter message classification tasks in a multi-lingual manner, such as election analysis, emotion recognition, crisis detection, among others. In crisis message classification, Torres and Vaca (2019) compared traditional and deep learning models using sparse representations and word embeddings to classify earthquake-related conversations in English and Spanish. The best cross-lingual results were using a Long Short Term Memory model including multilingual stacked embeddings, reporting a macro F1-score of 85.88% from Spanish to English and 77.49% from English to Spanish. Lorini et al. (2019) evaluated pre-trained language-agnostic and language-aligned word embeddings with Convolutional Neural Networks for the classification of flood-related messages in German, English, Spanish, and French. They compared a mono-lingual classifier, a cross-lingual classifier with *cold start* (using no training data in the target language), and a cross-lingual classifier with *warm start* (using 300 labeled instances in the target language). They showed that both types of word embeddings could be used to classify a new language for which few or no labels are available. However, including a small set of data from the same target language improved the cross-lingual classification. Another approach was proposed by Khare et al. (2018), who considered messages in English, Italian and Spanish from 30 crisis events of different types. They proposed a statistical-semantic crisis representation, extracting semantic relationships from BabelNet and DBpedia knowledge bases. They achieved a cross-lingual classification F1-score of 59.9% on average. In addition, Kayi et al. (2020) trained cross-lingual models based on contextual embeddings such as BERT, RoBERTa and XLM-R, to detect urgency messages in low-resource languages, specifically in Sinhala and Odia.

Domain Adaptation. Regarding domain adaptation using Twitter data, Agrawal and Awekar (2018) analyzed cyberbullying detection across Twitter, Wikipedia and Formspring. They reported poor performance when transferring knowledge from Twitter to the other two datasets reaching around 10% F1-score. Similarly, Arango, Pérez, and Poblete (2020) evaluated the cross-dataset classification of hateful tweets, showing that the models do not generalize well over different Twitter datasets obtaining an F1-score of just 21% for the positive class. In the crisis classification problem, Imran, Mitra, and Srivastava (2016) validated the impact of adding training examples from a different domain than the target. They classified messages related to earthquakes and floods published in several languages. Their experiments showed that in scenarios where there is not enough data, increasing training examples with tweets from other languages can be useful if both are very similar (e.g., Italian and Spanish). In the case of domain adaptation, they concluded that using tweets from a different domain did not appear to improve performance. Li and Caragea (2020) proposed a domain adaptation with reconstruction to classify disaster tweets. This adaptation based on semi-supervised learning aims to reduce the shift between source and target data distributions. For this, the classifier is trained with labeled source data, together with unlabeled target data.

Differentiation from prior work. Our work differs in that we perform a systematical study of transfer learning for crisis message classification for scenarios in which little to no data is available. We focus on the case of how to leverage labeled data from high-resource languages to low-resource languages, as well as from well-known crisis domains to new domains. We do not focus on the classification algorithms, but on the data and experimental methodology. Moreover, we investigate which document representations and models work best for each particular target, putting together the most comprehensive dataset to date for this goal.

3 Proposed Approach

We propose an experimental analysis of transfer learning approaches across several crisis domains and different languages. In detail, we address the following specific objectives: 1) determine if it is possible to transfer knowledge from one language to another, 2) determine if it is possible to transfer knowledge from one crisis domain to another domain, and 3) explore how to achieve (1) and (2) effectively.

We focus on the task of classifying microblog messages into the binary categories *related* and *not related* to crisis events¹, within several transfer learning scenarios. We adopt the definition used by Kruspe, Kersten, and Klan (2020) that considers a message as *related* to a crisis when it *refers implicitly or explicitly to the specific disaster event for which it was retrieved*. However, we generalize this definition to consider messages as *related* if they meet the above criteria for *any disaster*, not only a particular specified event. In addition, we define a *domain* as a *type of disaster (or hazard)* such as earthquakes or floods.

¹We consider an event as a unique occurrence in time and place (McMinn, Moshfeghi, and Jose 2013).

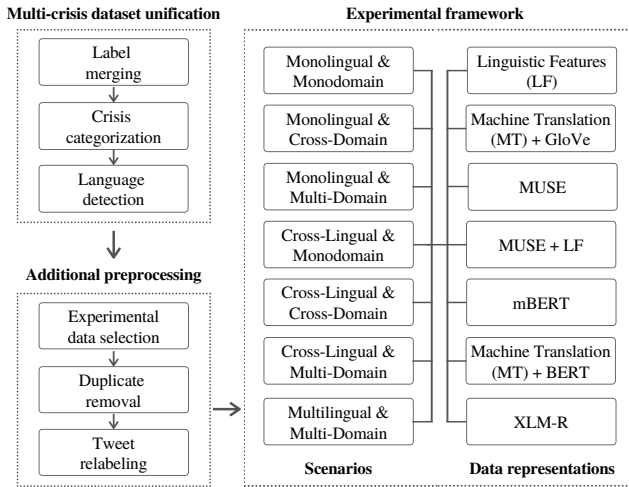


Figure 1: Overview of our proposed approach.

Figure 1 provides an overview of the main steps followed in this work, both for dataset construction and experimentation. We explain these steps more in detail next.

Transfer Learning Scenarios

We study 7 different classification scenarios that span across several transfer learning configurations. These scenarios aim to measure the impact on learning—of both domains and languages—for the classification of new crisis events. We include as baseline the scenario where there is no domain or language transfer learning. To avoid overfitting, we split the training and testing data at event level, ensuring that events that are used for training are not used for testing. More in detail:

- **Monolingual & Monodomain (MonoL & MonoD).** This is our baseline, where we train a classification model with messages from events of a specific crisis domain (e.g., earthquakes) all in the same language (e.g., English). We evaluate the model on messages from the same original domain and language, but from new events.
- **Monolingual & Cross-Domain (MonoL & CrossD).** In this scenario, we only perform domain transfer learning within the same language. The objective is to evaluate the effect of training with data from past events of a specific domain and language (e.g., earthquakes in English) to then classify messages from a different domain (e.g., explosions in English). This scenario evaluates the cold-start problem at domain level. This is, when a new type of disaster occurs for which there is no domain-specific labeled data.
- **Monolingual & Multi-Domain (MonoL & MultiD).** In this scenario we increase the amount of training data for a specific (low-resource) domain by using additional, same language, labeled data from other domains. For instance, we train a model using English earthquake, flood and explosion data. We then evaluate this model on English data from a new flood event.

- **Cross-Lingual & Monodomain (CrossL & MonoD).** In this scenario, we perform cross-lingual adaptation within the same domain. The objective is to evaluate the effect of training a classification model with data from the target domain (e.g., floods) in a language (e.g., English) that is different from the target language (e.g., Italian). This simulates the real-world scenario in which there is only labeled data from a high-resource language for one domain and a new event occurs in the same domain, but in a different language.
- **Cross-Lingual & Cross-Domain (CrossL & CrossD).** In this scenario we perform cross-lingual and domain adaptation. The objective is to evaluate the effect of training a classification model with data from a set of crisis domains (e.g., earthquakes and floods) in one language (e.g., English). Then, use this model to classify data of a different domain (e.g., explosions) in another language (e.g., Spanish). This simulates the case in which there is a need to classify messages of a new crisis domain in a language for which there is no labeled data.
- **Cross-Lingual & Multi-Domain (CrossL & MultiD).** In this scenario we perform cross-lingual adaptation and increase training data using other domains. The objective is to evaluate the effect of training a model using data from multiple domains in one language (e.g., floods, earthquakes, and hurricanes in English), to then classify a new event in another language (e.g., flood in Italian). This simulates the case in which there is a need to classify a well known type of event that occurs in a new language for which there is no labeled data.
- **Multilingual & Multi-Domain (MultiL & MultiD).** In this scenario we perform cross-lingual and domain data enrichment. This case is similar to the *Cross-lingual & Cross-domain* scenario with the addition of data from the target domain and language. For example, we train a model with English and Italian data about floods, earthquakes, and hurricanes, to then classify flood-related messages in Italian.

Data Representations

We focus on message representations and state-of-the-art models that allow us to convey multilingual data in a single feature space. Our representations correspond to 7 approaches described below.

- **Linguistic Features (LF).** This representation models each message as a set of linguistic features. We consider 48 features² represented in numerical and binary form. Some of these features have been previously used for classification of crisis messages (Graf et al. 2018; Khare et al. 2018). These features describe traditional message characteristics such as the number of *characters*, *words*, *links*, *mentions*, *hashtags*, *question marks*, among others. Furthermore, we consider attributes that are specific to each language, including *sentiment polarity*, *Named Entities* and *Part-of-Speech (POS)* (Al-Rfou et al. 2015). We

²The full description can be found in our repository.

also include binary features such as *has mention* and *has location*.

- **Machine Translation (MT) + GloVe.** This representation is based on modeling messages using English as a pivot language. Hence it implies translating all messages that are not written in the pivot language into English³. We choose English as our starting point for two reasons: 1) the high-quality of pre-trained embeddings in English, and 2) the trade-off of translating other languages to English, which is predominant in our dataset, as opposed to the other way around. Text in the pivot language is then tokenized and vectorized using the pre-trained GloVe⁴ model with 100-dimensions, trained on tweets (Pennington, Socher, and Manning 2014).
- **MUSE.** This representation uses MUSE⁵. MUSE consists of multilingual language-aligned word embeddings of 300-dimensions, based on fastText embeddings trained on Wikipedia (Lample et al. 2017).
- **MUSE + Linguistic Features (MUSE+LF).** This representation is a combination of MUSE and LF features. We use this to evaluate if the combination of semantic and statistical features improves model performance.
- **mBERT.** This representation models each message using the BERT-Base Multilingual Cased model⁶ trained with the top-104 languages on Wikipedia (Devlin et al. 2019).
- **Machine Translation (MT) + BERT.** This representation is similar to (MT) + GloVe, but instead of GloVe, we use BERT⁷ which has been trained for English using Wikipedia and BookCorpus (Devlin et al. 2019).
- **XLM-R.** This representation uses XLM-RoBERTa⁸, a multilingual version of RoBERTa trained on 100 languages of CommonCrawl (Conneau et al. 2020).

For Glove and MUSE, we combine single word embedding using mean aggregation, and for words Out-Of-Vocabulary, we create vectors with zero. On the other hand, when using mBERT, BERT and XLM-R (all 768-dimension models, with no fine-tuning), we combine the contextualized word embeddings using mean aggregation, ignoring the padding of zeros.

4 Unified Dataset Construction

We created a new dataset with 164,625 unique messages that is the result of collecting 7 crisis datasets from the literature (see Table 1 for details). These datasets met the criteria of having labeled data and being publicly available. The consolidated dataset, which we refer to as **Multi-Crisis Dataset** (available in our repository), contains Twitter messages from several crisis domains and events in different languages. To unify labels and to achieve an enriched, consistent dataset,

³We use the Google Translate API: <https://cloud.google.com/translate/>

⁴<https://nlp.stanford.edu/projects/glove/>

⁵<https://github.com/facebookresearch/MUSE>

⁶<https://huggingface.co/bert-base-multilingual-cased>

⁷<https://huggingface.co/bert-base-uncased>

⁸<https://huggingface.co/xlm-roberta-base>

Dataset	Lang.	Tweets	Events
CrisisLexT6 (Olteanu et al. 2014)	English	60,082	6
CrisisLexT26 (Olteanu, Vieweg, and Castillo 2015)	Multiple	27,933	26
CrisisNLP_R1 (Imran, Mitra, and Castillo 2016)	English, Spanish, French	49,596	14
Ecuador-Earthquake (Torres and Vaca 2019)	English, Spanish	8,360	1
SoSIItalyT4 (Cresci et al. 2015)	Italian	5,642	4
ChileEarthquakeT1 (Cobo, Parra, and Navón 2015)	Spanish	2,187	1
CrisisMMD (Alam, Ofli, and Imran 2018)	English	11,400	7
Multi-Crisis Dataset (Ours)	Multiple	164,625	53

Table 1: Datasets used to create our Unified Dataset.

we performed a process that consisted of several data curation steps. These mainly consisted of *label merging*, *crisis categorization*, and *language detection*. Next, we provide relevant details of this process⁹.

Label Merging. We based our dataset unification process on prior work in the area of crisis informatics that has addressed the label merging problem. In particular, in order to merge the 7 different datasets used in our study we followed the systematic label mapping approach used by Khare et al. (2018); Alam et al. (2021b); Sarmiento and Poblete (2021). This methodology maintains data consistency by unifying labels into a more general class according to their semantic meaning (Alam et al. 2021b). Specifically, the datasets used in our study had different types of labels, both binary and categorical at diverse granularities. We therefore merged these labels at a higher generalization level, which in this case was their relatedness to crisis events (i.e., *related* and *not-related* messages). These new labels correspond with the binary task required for our current work, and are similar to those already used in the *CrisisLexT6*, *ChileEarthquakeT1* and *Ecuador-Earthquake* datasets.

For instance, examples of original labels that we mapped to the *not-related* category in our dataset were: “*not related*”, “*not relevant*”, “*not related or irrelevant*”, “*off-topic*”, “*not informative*”, “*not applicable*”, and “*not physical landslide*”. On the other hand, examples of labels mapped to the *related* category were: “*related*”, “*relevant*”, “*damage*”, “*no damage*”, “*on-topic*”, “*informative*”, “*related and informative*”, “*related but not informative*”, among others. Furthermore, we manually relabeled messages from the CrisisMMD and CrisisLexT26 datasets originally labeled as “*not informative*” and “*Not applicable*” to accurately match our current task definition. The detailed mappings and relabeled data can be found in our repository.

⁹The in-depth procedure can be found in our repository.

Lang.	Count	%	Domain	Count	%
English	137,743	83.67	Earthquake	41,931	25.47
Spanish	12,025	7.30	Flood	31,923	19.39
Italian	7,002	4.25	Hurricane	19,578	11.89
French	1,144	0.70	Typhoon	13,674	8.31
Portuguese	771	0.47	Explosion	12,004	7.29
Tagalog	502	0.31	Bombings	11,012	6.69
Russian	238	0.15	Tornado	9,992	6.07
German	124	0.08	Landslide	4,492	2.73
Indonesian	111	0.07	Wildfires	3,533	2.15
Dutch	101	0.06	Viral disease	3,512	2.13
Others	4,864	2.96	Others	12,974	7.88
Total	164,625	100.00	Total	164,625	100.00

(a) Messages by language.

(b) Messages by domain.

Table 2: Number and percentage of messages in the Unified Multi-Crisis Dataset detailed by (a) language and (b) crisis domain separately.

Crisis Categorization. We annotated each message according to the crisis dimensions of the event that it belongs to. We used a similar pipeline to Olteanu, Vieweg, and Castillo (2015). We categorized crises by hazard types (such as “*earthquake*” or “*explosion*”), hazard categories (“*natural*” or “*human-induced*”), sub-categories (e.g., “*geophysical*”, “*hydrological*”, “*accidental*”, etc.), temporal development (“*instantaneous*” or “*progressive*”), and geographic spread (“*focalized*” or “*diffused*”). We decided to include this information because it will allow the study of communication patterns along the different dimensions. For example, to study the similarities and differences in communication during natural events versus human induced events. Furthermore, we aggregated information about the crisis, such as country and year.

Unified Dataset Description. Our dataset is mainly composed of messages in English (83.67%), followed by Spanish (7.30%) and Italian (4.25%). With regard to the domain, it is mainly composed of earthquake (25.47%), flood (19.39%) and hurricane (11.89%) domains. Table 2 shows more detailed information per language and domain.

5 Experimental Setup

In this section we explain our experimental setup for evaluating the proposed transfer learning scenarios and data representations (see Section 3). For our experiments we used a portion of the Multi-Crisis Dataset, selecting messages from the top-3 most represented languages in the dataset (English, Spanish and Italian). In addition, we discarded messages corresponding to events that 1) contained very little data, or 2) were from hazard domains not available for more than one language. This resulted in a dataset with 67,001 tweets from various regions in 3 languages and from 3 hazard domains: earthquakes (46.9%), floods (38.4%) and explosions (14.7%). Additionally, our experimental dataset contains 80.0% of English messages, 11.3% in Spanish, and

8.7% in Italian. Regarding the label distribution, 36.0% were categorized as *not-related* to crisis while 64.0% were labeled as *related*.

For evaluation, each event was distributed into a training or test set. To provide more representative examples for generalization, training sets prioritized events with the highest and most balanced number of instances in our dataset. We did not use English as a target language for cross-language evaluation because it is a high-resource language. The detailed training/testing partitions for each evaluation scenario can be found in our repository.

Table 3 shows the number of instances used for model training across scenarios (after balancing the classes), grouped by target language and domain. This data displayed an important imbalance between the positive (*related*) and negative (*not related*) classes, where the positive was approximately 27% more represented than the negative class. To handle this we applied random subsampling and oversampling of the positive and negative classes, respectively.

We trained all models using a Random Forest classifier from the Scikit-Learn implementation¹⁰ by considering the following default parameters: $n_estimators = 100$ that represents the number of trees in the forest, *gini* criterion, and *nolimit* for the maximum tree depth, so nodes expand until all leaves are pure or until all leaves contain less than two samples. We chose Random Forest because previous works have shown competitive results for tweet classification tasks and practical applications (Firoj, Imran, and Offi 2019; Li et al. 2018; Cobo, Parra, and Navón 2015; Arango, Pérez, and Poblete 2020; Alam et al. 2021a). Nevertheless, the scope of our current work is not on comparing classification algorithms, but on classifier initialization approaches for low-resource languages, which could help train more robust classifiers.

Table 4 shows the number of instances related to the events that remained for model testing. As in our training data, testing data also showed an important class imbalance. Due to the smaller amount of messages in our negative class for testing, oversampling to improve class balance was not a feasible solution. To deal with this situation we opted to augment the negative class by including a random sample of negative instances from the complete testing set (i.e., not only from the test event for that scenario). As negative instances are by definition messages not related to any crises, this does not affect evaluation. In the cases where this type of augmentation was not possible due to language restrictions for the scenario, we augmented with translated negative instances from English (translated using Google Translate).

6 Results

We present a comprehensive summary of the results of our evaluation following our experimental setup. There are many additional results based on the combinations of scenarios, domains and languages, which we cannot include due to space constraints. However, the exhaustive evaluation results can be found in our repository, including an evaluation

¹⁰<https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>

Lang.	Domain	MonoL & MonoD	MonoL & CrossD	MonoL & MultiD	CrossL & MonoD	CrossL & CrossD	CrossL & MultiD	MultiL & MultiD
English	Earthquake	11,214	35,720	46,238	-	-	-	-
	Explosion	-	51,662	-	-	-	-	-
	Flood	10,346	34,418	46,916	-	-	-	-
Spanish	Earthquake	2,822	150	-	18,324	35,720	56,346	67,260
	Explosion	-	4,182	-	9,282	51,662	63,460	69,768
Italian	Earthquake	1,520	414	-	20,426	35,720	63,122	67,520
	Flood	-	2,114	-	26,674	34,418	63,244	67,686

Table 3: Number of training instances by target (language and domain) used in each scenario, considering the balanced sets. The symbol “-” means that no experiments were performed.

Lang.	Domain	Related	Not related
English	Earthquake	8,611	1,225
	Explosion	4,415	4,641
	Flood	8,272	5,747
Spanish	Earthquake	2,507	453
	Explosion	747	50
Italian	Earthquake	698	198
	Flood	1,759	138

Table 4: Data for testing each target (before augmentation).

of XLM-T¹¹ (model fine-tuned to multilingual Twitter data), which was outperformed by XLM-R. We begin this section by describing the results obtained by target language and domain, running each model five times and averaging their values. Finally, we present an exhaustive analysis to determine the best performance for each classification scenario.

English Classification

We evaluated the classification of crisis related messages in English. Specifically, monolingual scenarios for the same domain, cross-domain (domain adaptation) and multiple domains (data enrichment). We did not evaluate cross lingual scenarios with English as a target language as we consider it as a high-resource language in our setup. However, we performed cross-lingual adaptation for low-resource languages using English as a source, detailed in the following sections.

Table 5 details the results obtained per scenario for each data representation and each of the three domains: earthquakes, explosions and floods, respectively. The overall results show that models achieved their best performance by using XLM-R, followed by MT+GloVe and MT+BERT, while using LF, MUSE and mBERT did not work as well in this case. However, mBERT’s performance would probably improve if we not only use its representations (as with the other pre-trained models), but also use its capabilities to adapt the model to the crisis domain.

The best performing model for earthquake message classification (evaluation included messages from *Chile 2014*, *California*, *Pakistan* and *Ecuador* events) was 89% F1-score

in the cross-domain scenario. In this scenario, we developed domain adaptation from explosion and flood domains. In particular, this scenario obtained an average improvement of 5% over the baseline scenario (monodomain). For explosion message classification, our dataset only contained one event (*West Texas* event). Using this event as target we evaluated the only possible scenario, i.e., cross-domain (training was done with earthquakes and floods), yielding an F1-score of 92% with the best data representation (GloVe). Recall that this scenario simulates the case in which we need to classify events from a new type of crisis event.

For flood message classification (evaluation included events of *Queensland*, *Pakistan* and *India*), the best scenario was the baseline with F1-score of 90% using MUSE+LF. Although, cross-domain (87% F1) showed good performance, similar to that of earthquake classification. On the other hand, multi-domain (89% F1) or domain enrichment achieved competitive results.

Overall, we observe that domain adaptation appears to work well within a high-resource language. This would allow us to use past knowledge to classify new and unexpected events in the same language. Therefore, it indicates one could use a pre-trained classifier to detect new types of crises that emerge. In addition, data augmentation (multi-domain) by including data from other domains can potentially improve model performance, or in the worst-case perform similar to the baseline.

Spanish Classification

We consider this language as low-resource, since the amount of labeled data is significantly less than for English (see Table 2). Our experiments included monolingual, cross-lingual and multilingual scenarios. Table 6 presents the results for earthquake and explosion domains, respectively, broken down by the data representations. Overall, MT+BERT and XLM-R obtained the best results, followed by MT+GloVe and mBERT, while LF and MUSE did not perform as well.

In the case of earthquake messages classification, the best performance scenarios with an 86% F1-score (on events from *Ecuador*, *Guatemala* and *Costa Rica*) used MT+BERT features, improving over the (*Monolingual & Monodomain*) baseline in average a 3%. We observed this improvement for the cross-lingual and multilingual scenarios that used some sort of cross-language adaptation from English. Among

¹¹<https://huggingface.co/cardiffnlp/twitter-xlm-roberta-base>

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & MonoD	0.77	0.82	0.82	0.82	0.81	0.81	0.84
MonoL & CrossD	0.84	0.87	0.86	0.87	0.88	0.89	0.89
MonoL & MultiD	0.83	0.87	0.85	0.86	0.87	0.89	0.88

(a) F1-score (Earthquake)

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & CrossD	0.89	0.92	0.86	0.89	0.87	0.86	0.90

(b) F1-score (Explosion)

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & MonoD	0.84	0.88	0.88	0.90	0.85	0.88	0.89
MonoL & CrossD	0.82	0.86	0.84	0.86	0.84	0.87	0.86
MonoL & MultiD	0.83	0.88	0.87	0.88	0.86	0.89	0.89

(c) F1-score (Flood)

Table 5: Comparison of the data representations’ performance in (a) earthquake, (b) explosion, and (c) flood domains for English message classification. For each domain, the best data representation is highlighted in bold and the best scenario in a gray cell.

these best scenarios, there is the *Cross-Lingual & Cross-Domain* scenario, which simulates classifying messages of a new type of crisis domain in a new language. Hence, there does not appear to be additional improvement when including data from the target domain and from the target language. In addition, the worst performance was for the *Monolingual & Cross-Domain* scenario, which simulates the case when we attempt to classify messages from a new domain in the same language. This most likely occurs due to the small amount of training data in Spanish, which limits cross-domain learning within that language.

For the explosion message classification in Spanish, we only had one event’s worth of data (an event in *Venezuela*). Therefore, we were not able to evaluate the baseline scenario (*Monolingual & Monodomain*). We observed that the best performance scenario, with 85% F1-score, was the *Multilingual & Multi-Domain* scenario using XLM-R as features, and data in English and Spanish as the source. As with earthquakes, the worst performance scenario was training with another domain in the same language as the target (i.e., *Monolingual & Cross-Domain*). The worst performance feature along all scenarios was MUSE. In general,

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & MonoD	0.73	0.80	0.81	0.80	0.80	0.83	0.79
MonoL & CrossD	0.66	0.72	0.75	0.75	0.78	0.79	0.75
CrossL & MonoD	0.68	0.79	0.75	0.75	0.76	0.83	0.80
CrossL & CrossD	0.78	0.84	0.81	0.81	0.84	0.86	0.85
CrossL & MultiD	0.77	0.84	0.80	0.81	0.84	0.86	0.84
MultiL & MultiD	0.75	0.84	0.74	0.77	0.79	0.86	0.83

(a) F1-score (Earthquake)

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & CrossD	0.73	0.77	0.69	0.72	0.76	0.69	0.77
CrossL & MonoD	0.75	0.84	0.64	0.74	0.72	0.79	0.79
CrossL & CrossD	0.73	0.81	0.74	0.77	0.82	0.77	0.75
CrossL & MultiD	0.74	0.82	0.74	0.78	0.82	0.82	0.80
MultiL & MultiD	0.76	0.83	0.75	0.79	0.83	0.83	0.85

(b) F1-score (Explosion)

Table 6: Comparison of the data representations’ performance in (a) earthquake and (b) explosion domains for Spanish message classification. For each domain, the best data representation is highlighted in bold and the best scenario in a gray cell.

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & MonoD	0.67	0.72	0.75	0.72	0.70	0.74	0.72
MonoL & CrossD	0.67	0.73	0.77	0.76	0.78	0.76	0.77
CrossL & MonoD	0.56	0.77	0.78	0.77	0.75	0.79	0.70
CrossL & CrossD	0.63	0.78	0.79	0.76	0.68	0.77	0.77
CrossL & MultiD	0.62	0.78	0.80	0.77	0.71	0.78	0.76
MultiL & MultiD	0.67	0.79	0.76	0.79	0.73	0.82	0.75

(a) F1-score (Earthquake)

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & CrossD	0.72	0.72	0.74	0.78	0.76	0.77	0.81
CrossL & MonoD	0.79	0.81	0.73	0.82	0.53	0.68	0.83
CrossL & CrossD	0.74	0.75	0.73	0.76	0.65	0.74	0.82
CrossL & MultiD	0.77	0.79	0.75	0.80	0.63	0.73	0.84
MultiL & MultiD	0.75	0.81	0.79	0.83	0.78	0.80	0.79

(b) F1-score (Flood)

Table 7: Comparison of the data representations’ performance in (a) earthquake and (b) flood domains for Italian message classification. For each domain, the best data representation is highlighted in bold and the best scenario in a gray cell.

our results indicate that baseline performance can be improved by augmenting a low-resource language with high-resource language data, including multiple domains.

Italian Classification

For Italian, we performed a similar evaluation to Spanish. Table 7 presents the classification results for earthquake and flood domains, respectively. We observed that the MT+BERT representation performed better in the classification of earthquake messages, while XLM-R performed better with the flood domain. The other representations showed behavior that varies according to the target domain. For example, for classifying earthquakes, LF features obtained the worst performance, and MUSE features obtained the best performance in most scenarios. However, such behavior was different when classifying floods.

For earthquakes the best performance (on *L’Aquila* event) was achieved in the *Multilingual & Multi-Domain* scenario (82% F1) using MT+BERT and for floods (on *Sardinia* and *Genova* events) it was in the *Cross-Lingual & Multi-Domain* scenario (84% F1) using XLM-R. In general, the cross-lingual adaptation and augmentation using English (a high-resource language) improved performance. Also, adding multiple domains from English and the target language increased this improvement.

Results by Classification Scenario

We present the aggregated results for Spanish and Italian languages, grouped by scenario and detailed by features in Table 8. We exclude the baseline scenario since not all targets were evaluated in that scenario due to the scarcity of data. The best performance was achieved in the *Multilingual & Multi-Domain* scenario, i.e., when we trained our model with all available data in English and the target language

from multiple domains, followed by the *Cross-Lingual & Multi-Domain* scenario. On the contrary, the lowest performance was obtained in the *Monolingual & Cross-Domain* scenario. As for the data representations, we observed that the *baseline feature (LF)* presents the lowest score in all scenarios, while *MT+BERT* achieved the highest score. However, XLM-R and MT+GloVe performed better in most scenarios. Overall, we found that increased training data with multiple domains in English language contributed positively to the construction of multilingual and multi-domain models. Given the constraints posed by the diversity of data sources and the disparities in the size of the data, this improved our results. However, with better quality data we may observe larger effects on the improvements.

7 Discussion

Our findings indicate that *Multilingual & Multi-Domain* adaptation is an effective way to improve the classification of low-resource languages. When no labeled data is available for a target language, a good option is to perform cross-lingual adaptation from a high-resource language using all available domain data. Most importantly, we show that it is possible to classify messages that correspond to a new, previously unseen, crisis event when they are in an unknown language. This can be very useful in identifying unexpected emerging crisis situations for early response. However, the similarity of the studied languages, selected according to the availability of data, can aid in classification results. Even though English belongs to the Germanic language family, it shares a significant number of cognate words with Spanish and Italian, which are Romance languages (Comrie 1987). This similarity is attributed to the influence of Latin, which has a shared origin between these languages (Whitley 2002; Maiden and Robustelli 2014). Thus, the suitability of En-

Scenario	LF	MT+GloVe	MUSE	MUSE+LF	mBERT	MT+BERT	XLM-R
MonoL & CrossD	0.69	0.73	0.74	0.75	0.77	0.75	0.78
CrossL & MonoD	0.69	0.80	0.73	0.77	0.69	0.78	0.78
CrossL & CrossD	0.72	0.80	0.77	0.78	0.75	0.79	0.80
CrossL & MultiD	0.72	0.81	0.77	0.79	0.75	0.80	0.81
MultiL & MultiD	0.73	0.82	0.76	0.79	0.78	0.83	0.81

Table 8: Average F1-score by classification scenario and feature for the low-resource languages Spanish and Italian, including their different crisis domains. The test examples are the same in all scenarios. The best result per scenario is highlighted in bold and the best overall result is in a gray cell.

lish as a high-resource language is uncertain for other languages with significant grammatical and syntactical differences.

We propose and evaluate data augmentation (or enrichment) mechanisms that encompass a suite of techniques that enhance the size and quality of training datasets such that better models can be built using them. Data augmentation relies on creating a larger input data size to reduce the shift between the source and target distributions. However, incorporating new data on training sets could decrease the accuracy of the classification models because of the lack of the ability to generalize adequately. In this sense, rather than adding noise when we included new content, our results indicated that data augmentation from multiple domains and languages could provide an opportunity in scenarios when data is insufficient for specific types of events.

A comparison of cross-domain results for Spanish and Italian reveals that transferring knowledge from one or more domains to another is useful from high-resource languages, such as English. Nevertheless, it does not help if we use as a source low-resource languages (e.g., training with Spanish explosion to classify earthquakes in the same language).

Regarding the most effective ways to represent data for knowledge transfer, translating the content into English and using BERT or GloVe models provides the most accurate results for Spanish messages in most scenarios. However, for Italian, XLM-R and MUSE+LF provide the best results. This can be due to variability in the quality of machine translation, for example. In practice, translating messages to English may not be cost-effective. However, this analysis allows us to understand the available alternatives. In terms of F1-score, the observed difference between MT+GloVe, MT+BERT and XLM-R models is not statistically significant at a 95% confidence level. In addition, translating the content into English and then using the BERT model (trained in English corpus), obtained statistically superior results to using the mBERT model (trained on multilingual corpus).

Aligned embeddings, such as MUSE, are a promising less expensive approach for low-resource language classification. We also observe that the LF representation was not as competitive across scenarios. However, when this feature is combined with MUSE, it improves results, specifically for Italian. Regarding mBERT, we show that it provides competitive results for some classification scenarios, but this is not consistent. For future work, we will apply fine-tuning to

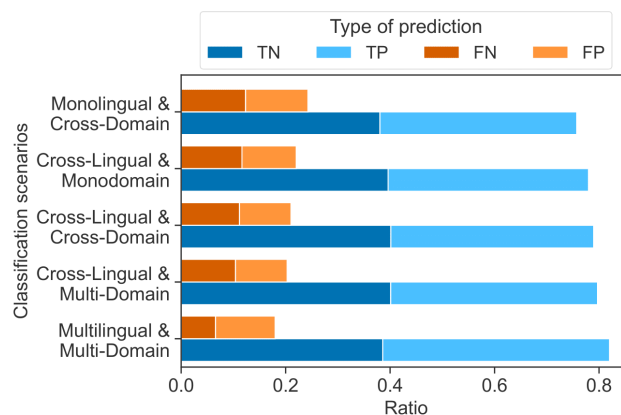


Figure 2: Predictions on Spanish and Italian test sets using the MT+BERT model. The test examples are the same in all scenarios. TP: True Positive, TN: True Negative, FP: False Positive, FN: False Negative.

this model (as well as to XLM-R) to explore whether the performance increases.

For explosion in Spanish and flood in Italian, we observed a decrease in performance for the *Cross-Lingual & Monodomain* scenario, using the MUSE and mBERT compared to LF. This could be due to: 1) the dependence of specific words to the crisis domain and 2) the similarity of the representations of those words in both languages (English and the target language). The latter could be explained by the general corpus used to train MUSE and mBERT (Wikipedia).

Analysis of Predictions

We consider the analysis of the predictions to focus on the errors and suggest future improvements. Figure 2 shows the rate of correct and incorrect predictions for the Spanish and Italian languages using the MT+BERT model. We chose this model because it achieved the best overall performance for these languages (see Table 8). We observe that the last scenario (*Multilingual & Multi-Domain*) presents the lowest error rate (FN + FP), where the rate of false negatives is lower than that of false positives. These rates vary depending on the feature used. For example, with the MUSE+LF classifier, this scenario has a higher false positive-rate.

We performed a manual analysis of some classification

Tweet text	Crisis	Lang.
(A) <i>Idk why these things are happening but I know god is with us !!!</i>	West Texas explosion	EN
(B) <i>Svegliato dal terremoto, non sar� un bel giorno per molti.</i>	L'Aquila earthquake	IT
(C) <i>Cifra de muertos en Jap�n aumenta a 41 tras terremoto: Equipos apuran el rescate https://t.co/uQJVU3ijDf</i>	Ecuador earthquake	ES
(D) <i>15 Germania: sparatoria in tribunale, almeno due morti e diversi feriti: Il conflitto a fuoco nell'aula di Lan.. http://tinyurl.com/cdb73r</i>	L'Aquila earthquake	IT

Table 9: Examples of misclassified messages in their original languages. (A) and (B) represent false negatives and (C) and (D) false positives.

errors. This analysis can provide insights into what might be the errors we need to address for future research. Table 9 shows some misclassified examples for all three languages. Examples (A) and (B) are false negatives mainly belonging to the original label: *related - but not informative* (CrisisLexT26), *on-topic* (CrisisLexT6) and *related* (Ecuador-Earthquake). This reveals that false negatives are sometimes expressions of personal opinions, prayers, sympathy and emotional support. Examples (C) and (D) show false positives mainly belonging to: *off-topic* (CrisisLexT6), *not related* (CrisisLexT26) and *not related or irrelevant* (CrisisNLP_R1). In the first example, the message contains explicit information about an earthquake in Japan, but since it had been collected for the Ecuadorian earthquake, it was considered as *not related*, which could be considered as a mislabeled example for our task definition. Likewise, (D) is related to a shooting event, but not to L'Aquila earthquake. In both cases, the classification model correctly marked them as being *related* to crises. However, these messages were (incorrectly) labeled in the dataset as *not related*, because in their original datasets they were not related to the crisis that was being collected at the time.

We recognize within the limitations of our work, the unification of datasets from different sources. We have done our best to achieve coherent labels, but there may be inconsistencies. Moreover, it is worth noting that automatic translation tools, such as Google translate, may not yield completely accurate results in all instances, particularly due to the non-structural and informal nature of social media content. While alternative translation tools may provide superior accuracy, we did not have access to them due to the high costs associated with their use for a significant volume of data. Finally, we evaluated our approach only on the binary task by relatedness to crises. We expect that our methodology allows to generalize other fine-grained tasks, such as the classification of types of humanitarian information.

8 Conclusion

The main goal of this work was to assess the impact of transfer learning in the multilingual and multi-domain classification of messages relating to crises. In particular, data from

a high-resource language such as English can contribute to the classification of messages from low-resource languages such as Spanish and Italian. Furthermore, adding messages from the target language also helps in some cases. Our findings indicate that there exist patterns in crisis communications that expand across crisis domains and languages. As a result we can increase our ability to classify data in languages and domains for which we have little to no labeled data. However, the most efficient data representations may vary depending on the target language.

For future work we want to explore in-depth different deep learning classifiers and techniques for cross-lingual and domain adaptation. In addition, we would like to further improve our dataset by manually relabeling certain mislabeled messages that are effectively related to crises according to our task definition. This will allow us to have a data collection that is oriented towards crises in general, rather than event oriented. Moreover, we will explore this transfer learning approach in a more fine-grained task, such as categorizing actionable humanitarian information.

Ethical Statement

We use existing social network data, clearly indicating the source and respecting the policies on data sharing and anonymity. Our overarching goal is to expand crisis informatics applications to low-resource languages. However, there exist inequalities in access to the Internet, especially prevalent in low-income countries. These geographical and socio-demographic biases are present in the source of the data and create important challenges for crisis informatics.

Acknowledgments

C.S. was supported by SENESCYT-Ecuador Scholarship and ANID / Scholarship Program / DOCTORADO BECAS CHILE/2022 - 21222229. H.S. was supported by Scholarship Program / DOCTORADO BECAS CHILE/2020 - 21201101. This work was partially funded by Millennium Institute for Foundational Research on Data and the National Center for Artificial Intelligence CENIA FB210017, Basal ANID. In addition, we are very grateful to Aym  Arango, James Caverlee and Aidan Hogan for their useful comments.

References

- Agrawal, S.; and Awekar, A. 2018. Deep learning for detecting cyberbullying across multiple social media platforms. In *Proceedings of the 40th European Conference on Information Retrieval*, 141–153. Springer.
- Al-Rfou, R.; Kulkarni, V.; Perozzi, B.; and Skiena, S. 2015. Polyglot-NER: Massive Multilingual Named Entity Recognition. *Proceedings of the 2015 SIAM International Conference on Data Mining*, 586–594.
- Alam, F.; Joty, S.; and Imran, M. 2018. Domain Adaptation with Adversarial Training and Graph Embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 1077–1087. ACL.
- Alam, F.; Ofli, F.; and Imran, M. 2018. CrisisMMD: Multimodal Twitter Datasets from Natural Disasters. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*, 465–473. AAAI.

- Alam, F.; Qazi, U.; Imran, M.; and Ofli, F. 2021a. HumAID: Human-Annotated Disaster Incidents Data from Twitter with Deep Learning Benchmarks. *Proceedings of the International AAAI Conference on Web and Social Media*, 15(1): 933–942.
- Alam, F.; Sajjad, H.; Imran, M.; and Ofli, F. 2021b. CrisisBench: Benchmarking Crisis-related Social Media Datasets for Humanitarian Information Processing. *Proceedings of the International AAAI Conference on Web and Social Media*, 15(1): 923–932.
- Arango, A.; Pérez, J.; and Poblete, B. 2020. Hate speech detection is not as easy as you may think: A closer look at model validation (extended version). *Information Systems*.
- Chan, J. C. 2013. The role of social media in crisis preparedness, response and recovery. *Citeseer*.
- Cobo, A.; Parra, D.; and Navón, J. 2015. Identifying Relevant Messages in a Twitter-Based Citizen Channel for Natural Disaster Situations. In *Proceedings of the 24th International Conference on World Wide Web*, 1189–1194. ACM.
- Comrie, B. 1987. *The World's Major Languages*. Croom Helm Ltd.
- Conneau, A.; Khandelwal, K.; Goyal, N.; Chaudhary, V.; Wenzek, G.; Guzmán, F.; Grave, E.; Ott, M.; Zettlemoyer, L.; and Stoyanov, V. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8440–8451. ACL.
- Cresci, S.; Tesconi, M.; Cimino, A.; and Dell'Orletta, F. 2015. A Linguistically-Driven Approach to Cross-Event Damage Assessment of Natural Disasters from Social Media Messages. In *Proceedings of the 24th International Conference on World Wide Web*, 1195–1200. ACM.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, 4171–4186. ACL.
- Firoj, A.; Imran, M.; and Ofli, F. 2019. CrisisDPS: Crisis Data Processing Services. In *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management*, 719–733. ISCRAM.
- Ghosh, S.; Ghosh, K.; Ganguly, D.; Chakraborty, T.; Jones, G. J. F.; Moens, M.-F.; and Imran, M. 2018. Exploitation of Social Media for Emergency Relief and Preparedness: Recent Research and Trends. *Information Systems Frontiers*, 20(5): 901–907.
- Graf, D.; Retschitzegger, W.; Schwinger, W.; Pröll, B.; and Kapsammer, E. 2018. Cross-Domain Informativeness Classification for Disaster Situations. In *Proceedings of the 10th International Conference on Management of Digital EcoSystems*, 183–190. ACM.
- Imran, M.; Mitra, P.; and Castillo, C. 2016. Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages. In *Proceedings of the 10th International Conference on Language Resources and Evaluation*, 1638–1643. ELRA.
- Imran, M.; Mitra, P.; and Srivastava, J. 2016. Cross-Language Domain Adaptation for Classifying Crisis-Related Short Messages. In *Proceedings of the 13th International Conference on Information Systems for Crisis Response and Management*. ISCRAM.
- Kayi, E.; Nan, L.; Qu, B.; Diab, M.; and McKeown, K. 2020. Detecting Urgency Status of Crisis Tweets: A Transfer Learning Approach for Low Resource Languages. In *Proceedings of the 28th International Conference on Computational Linguistics*, 4693–4703. ICCL.
- Khare, P.; Burel, G.; Maynard, D.; and Alani, H. 2018. Cross-Lingual Classification of Crisis Data. In *Proceedings of the 17th International Semantic Web Conference (Part I)*, 617–633. Springer.
- Kruspe, A.; Kersten, J.; and Klan, F. 2020. Review article: Detection of informative tweets in crisis events. *Natural Hazards and Earth System Sciences Discussions*, 1–18.
- Lample, G.; Conneau, A.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2017. Word Translation Without Parallel Data. In *6th International Conference on Learning Representations*.
- Li, H.; Li, X.; Caragea, D.; and Caragea, C. 2018. Comparison of word embeddings and sentence encodings as generalized representations for crisis tweet classification tasks. In *Proceedings of ISCRAM Asia Pacific 2018: Innovating for Resilience – 1st International Conference on Information Systems for Crisis Response and Management Asia Pacific*, 480–493. Massey University.
- Li, X.; and Caragea, D. 2020. *Domain Adaptation with Reconstruction for Disaster Tweet Classification*, 1561–1564. ACM.
- Lorini, V.; Castillo, C.; Dottori, F.; Kalas, M.; Nappo, D.; and Salamon, P. 2019. Integrating Social Media into a Pan-European Flood Awareness System: A Multilingual Approach. In *Proceedings of the 16th International Conference on Information Systems for Crisis Response and Management*, 646–659. ISCRAM.
- Maiden, M.; and Robustelli, C. 2014. *A reference grammar of modern Italian*. Routledge.
- McMinn, A. J.; Moshfeghi, Y.; and Jose, J. M. 2013. Building a Large-scale Corpus for Evaluating Event Detection on Twitter. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, 409–418. ACM.
- Olteanu, A.; Castillo, C.; Diaz, F.; and Vieweg, S. 2014. CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises. In *Proceedings of the 8th International Conference on Weblogs and Social Media*, 376–385. AAAI.
- Olteanu, A.; Vieweg, S.; and Castillo, C. 2015. What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 994–1009. ACM.
- Palen, L.; and Anderson, K. 2016. Crisis informatics—New data for extraordinary times. *Science*, 353(6296): 224–225.
- Pan, S. J.; and Yang, Q. 2009. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10): 1345–1359.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532–1543. ACL.
- Ruder, S. 2019. *Neural transfer learning for natural language processing*. Ph.D. thesis, National University of Ireland.
- Sánchez, C. 2021. Transfer Learning for the Multilingual and Multi-Domain Classification of Messages Relating to Crises. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2708. ACM.
- Sarmiento, H.; and Poblete, B. 2021. Crisis Communication: A Comparative Study of Communication Patterns across Crisis Events in Social Media. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, 1711–1720. ACM.
- Torres, J.; and Vaca, C. 2019. Cross-lingual Perspectives About Crisis-Related Conversations on Twitter. In *Companion Proceedings of The 2019 World Wide Web Conference*, 255–261. ACM.
- Troudi, A.; Zayani, C.; Jamoussi, S.; and Amous, I. 2018. A New Mashup Based Method for Event Detection from Social Media. *Information Systems Frontiers*, 20(5): 981–992.
- Whitley, M. S. 2002. *Spanish/English contrasts: A course in Spanish linguistics*. Georgetown University Press.