

CrisisBERT: A Robust Transformer for Crisis Classification and Contextual Crisis Embedding

Junhua Liu*
Information Systems Technology and
Design Pillar
Singapore University of Technology
and Design
Singapore
junhua_liu@mymail.sutd.edu.sg

Trisha Singhal
SUTD-MIT International Design
Centre
Singapore University of Technology
and Design
Singapore
trisha_singhal@sutd.edu.sg

Lucienne T.M. Blessing
Engineering Product Development
Pillar
Singapore University of Technology
and Design
Singapore
lucienne_blessing@sutd.edu.sg

Kristin L. Wood†
Engineering Product Development
Pillar
Singapore University of Technology
and Design
Singapore
kristinwood@sutd.edu.sg

Kwan Hui Lim
Information Systems Technology and
Design Pillar
Singapore University of Technology
and Design
Singapore
kwanhui_lim@sutd.edu.sg

ABSTRACT

Detecting crisis events accurately is an important task, as it allows the relevant authorities to implement necessary actions to mitigate damages. For this purpose, social media serve as a timely information source due to its prevalence and high volume of first-hand accounts. While there are prior works on crises detection, many of them do not perform crisis embedding and classification using state-of-the-art attention-based deep neural networks models, such as Transformers and document-level contextual embeddings. In contrast, we propose *CrisisBERT*, an end-to-end transformer-based model for two crisis classification tasks, namely crisis detection and crisis recognition, which shows promising results across accuracy and F1 scores. The proposed *CrisisBERT* model demonstrates superior robustness over various benchmarks, and it includes only marginal performance compromise while extending from 6 to 36 events with a mere 51.4% additional data points. We also propose *Crisis2Vec*, an attention-based, document-level contextual embedding architecture, for crisis embedding, which achieves better performance than conventional crisis embedding methods such as *Word2Vec* and *GloVe*.

* Also with Forth AI, j@forth.ai.

† Also with University of Colorado Denver, kristin.wood@ucdenver.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HT '21, August 30-September 2, 2021, Virtual Event, Ireland

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8551-0/21/08...\$15.00

<https://doi.org/10.1145/3465336.3475117>

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**; **Neural networks**; • **Information systems** → **Clustering and classification**; **Data encoding and canonicalization**.

KEYWORDS

Social Computing, Crisis Classification, Word Embedding, Social Networks, Twitter

ACM Reference Format:

Junhua Liu, Trisha Singhal, Lucienne T.M. Blessing, Kristin L. Wood, and Kwan Hui Lim. 2021. CrisisBERT: A Robust Transformer for Crisis Classification and Contextual Crisis Embedding. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media (HT '21)*, August 30-September 2, 2021, Virtual Event, Ireland. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3465336.3475117>

1 INTRODUCTION

Crisis-related events, such as earthquakes, hurricanes and train or airliner accidents, often generate a sudden surge of attention and actions from both media and the general public. Despite the fact that crises, such as natural disasters, are often reported by mainstream media and professional institutions, certain events are first signaled by everyday citizens, i.e., civilians, on social media and other web platforms [9, 12, 48]. For instance, the recent COVID-19 pandemic was first highlighted by the general public in China via *Weibo*, a popular social media site, before pronouncements by government officials. This example highlights the potential of social media for highlighting such crisis-related events.

Social media sites have become centralized hubs that facilitate timely information exchange across government agencies, enterprises, working professionals and the general public [27, 35, 46]. As one of the most popular social media sites, Twitter enables users to asynchronously communicate and exchange information with *tweets*, which are mini-blog posts limited to 280 characters. Twitter

is also frequently used for various research purposes and applications [7, 38, 40, 42]. There are on average over half a billion tweets posted daily [43]. Therefore, one can leverage on such high volume and frequent information exchange to discover firsthand signals on crisis-related events for early detection and warning systems to reduce overall damage and negative impacts [3, 10, 38].

Event detection from tweets has received significant attention in research due to its ability for analyzing crisis-related messages to improve disaster management and increase situational awareness. These works mainly focus on either detecting general events or specific types of events such as political discourse or crisis events. For example, several recent works studied various natural crisis events such as hurricanes and earthquakes [39, 40], artificial disasters such as terrorist attacks and explosions [14, 41], and disease outbreaks such as epidemics and pandemics [19, 23].

These works focus on binary classification tasks for the various attributes of crisis, such as classifying source type, predicting relatedness between tweets and the crises, and assessing informativeness and applicability [30, 44, 49]. Several works also propose multi-label classifiers on affected individuals, infrastructure, casualties, donations, caution, advice, etc. [13, 15]. Crisis recognition tasks are likewise conducted such as identifying crisis types, i.e. hurricanes, floods and fires [6, 7].

Machine Learning-based models are commonly introduced in performing the above mentioned tasks. Conventional linear models such as Logistic Regression, Naive Bayes and Support Vector Machine (SVM) are reported for automatic binary classification on informativeness [32] and relevancy [44], among others. These models were implemented with pre-trained word2vec embeddings [26]. Several unsupervised approaches are also proposed for classifying crisis-related events, such as the ClusTop algorithm utilizing Community Detection for automatic topic modelling [21]. A transfer-learning approach is also proposed [33], though its classification is only limited to two classes. The ability for cross-crisis evaluation remains unaddressed in these earlier works.

More recently, numerous works proposed Neural Networks (NN) models for crisis-related data detection and classification. For instance, ALRashdi and O’Keefe investigated two deep learning architectures, namely Bidirectional Long Short-Term Memory (BiLSTM) and Convolutional Neural Networks (CNN) using domain-specific and GloVe embeddings [2]. Nguyen *et al.* propose a CNN-based classifier with Word2Vec embedding pretrained on Google News [26] and domain-specific embeddings [28]. Lastly, parallel CNN architecture was proposed to detect disaster-related events using tweets [16, 17].

While prior works report remarkable performance on various crisis classification tasks using NN models and word embeddings, limited studies are found to leverage the recent Natural Language Understanding (NLU) techniques, such as attention-based deep classification models [45] and document-level contextual embeddings [37]. These techniques reportedly improve state-of-the-art performance for many challenging natural language problems from upstream tasks such as Named Entity Recognition and Part of Speech Tagging, to downstream tasks such as Machine Translation and Neural Conversation. In this work, we explore how such techniques can be utilized to improve the tasks of crisis detection and recognition.

1.1 Main Contributions

This work utilizes deep attention-based classification models and document-level contextual representation models to address two important crisis classification tasks. We study recent NLU models and techniques that reportedly demonstrate drastic improvement on state-of-the-art and adapt them for domain-specific crisis related tasks.

Overall, our main contribution in this work includes:

- Proposing *CrisisBERT*, an attention-based classifier that improves state-of-the-art performance for both crisis detection and recognition tasks;
- Demonstrating superior robustness over various benchmarks, where extending *CrisisBERT* from 6 to 36 events with 51.4% of additional data points only results in marginal performance decline, despite an increase in the number of unique crises by 500%;
- Proposing *Crisis2Vec*, a document-level contextual embedding approach for crisis representation, and showing substantial improvement over conventional crisis embedding methods such as *Word2Vec* and *GloVe*.

The remaining of this paper is organized as follows. Section 2 provides background information on transformer networks and describes our proposed CrisisBERT model for crisis classification. Section 3 introduces our crisis classification problem, including two sub-problems of crisis detection and crisis recognition. Section 4 shows the experimental results of our proposed CrisisBERT model against various baselines. Section 5 discusses related works in the area of crisis classification and embedding techniques. Section 6 summarizes the main findings of our paper.

2 ATTENTION-BASED NEURAL CRISIS CLASSIFIER

In this section, we discuss recent works that propose various machine learning approaches for crisis classification tasks. While these works report substantial improvement in performance over prior works, none of them uses the state-of-the-art attention-based models, i.e., Transformers [45], to perform crisis classification tasks. In contrast, we propose *CrisisBERT*, a transformer-based architecture that builds upon a Distilled BERT model, fine-tuned by large-scale hyper-parameter search.

2.1 Conventional Crisis classifiers

Various works propose linear classifiers for crisis-related events. For instance, Parilla-Ferrer *et al.* proposed an automatic binary classification, based on informative and uninformative tweets using Naive Bayes and Support Vector Machine (SVM) [32]. A SVM with pretrained word2vec embeddings approach was also proposed in [26].

Besides linear models, recent works also propose deep learning based methods with different neural network architectures. For instance, ALRashdi and O’Keefe investigated Bidirectional Long Short-Term Memory (BiLSTM) and Convolutional Neural Network (CNN) models using domain-specific and GloVe embeddings [34]. Nguyen *et al.* proposed a CNN model to classify tweets to get

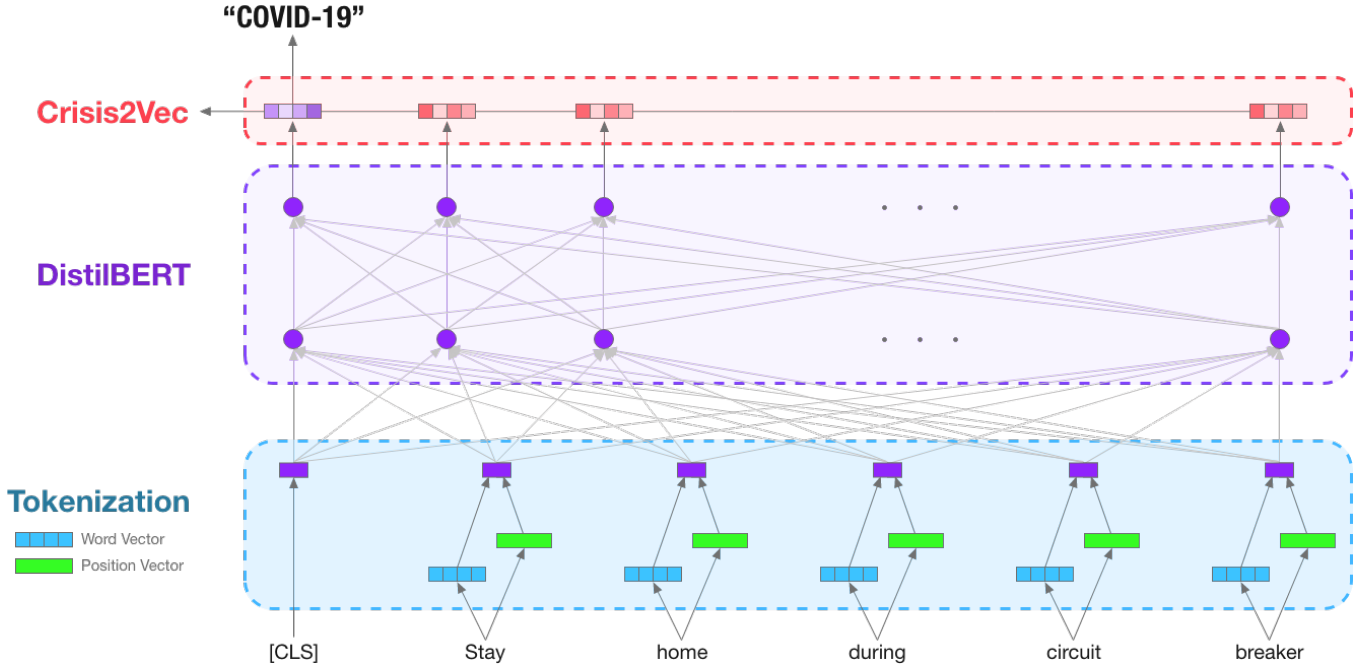


Figure 1: Illustration of *CrisisBERT* and *Crisis2Vec* models. During the tokenization phase, word-level and subword-level embeddings are created to produce context embedding for each input word. The vectors are then prefixed with a class header (shown as *[CLS]* in the figure) and passed into a DistilBERT model. Since we are performing classification task, the CLS token vector, i.e. the first output vector, is then passed into a linear classifier for detection or recognition task, whereas the remainder of the output vectors are average-pooled to create *Crisis2Vec* embeddings.

information-types using Google News and domain-specific embeddings [28].

2.2 Transformer

In 2017, Vaswani *et al.* from Google introduced Transformer [45], a new category of deep learning models which are solely attention-based and without convolution and recurrent mechanisms. Later, Google proposed the Bidirectional Encoder Representations from Transformers (BERT) model [8] which drastically improved state-of-the-art performance for multiple challenging Natural Language Processing (NLP) tasks. Since then, multiple transformer-based models have been introduced, such as GPT [36] and XLNet [47], among others. Transformer-based models were also deployed to solve domain specific tasks, such as medical text inferencing [20] and occupational title embedding [22], and demonstrated remarkable performance.

The Bidirectional Encoder Representation of Transformer (BERT), for instance, is a multi-layer bidirectional Transformer encoder with attention mechanism [8]. The proposed BERT model has two variants, namely (a) BERT Base, which has 12 transformer layers, a hidden size of 768, 12 attention heads, and 110M total parameters; and (b) BERT Large, which has 24 transformer layers, a hidden size of 1024, 16 attention heads, and 340M total parameters. BERT is pre-trained with self-supervised approaches, i.e., Masked Language Modeling (MLM) and Next Sentence Prediction (NSP).

While Transformers such as BERT are reported to perform well in natural language processing, understanding and inference tasks, no prior works propose and examine the performance of transformer-based models for crisis classification.

2.3 CrisisBERT

In this work, we investigate the transformer approach for crisis classification tasks and propose *CrisisBERT*, a transformer-based classification model that surpasses conventional linear and deep learning models in performance and robustness. Figure 1 illustrates an overview of the proposed architecture. Overall, the architecture of *CrisisBERT* includes three layers, namely (1) Contextual Tokenization, (2) Transformer Language Model, and (3) Logistic Classifier.

Tokenization. DistilBERT inherits the embedding layer from BERT, where each word is tokenized contextually with subword-level word embeddings and positional encoders. Prefixed with a special token (*[CLS]*), the vector pairs are concatenated and passed into a DistilBERT LM for training.

Transformer LM. We performed a large-scale model selection and hyperparameter search to find the best performing transformer LM for the crisis classification tasks. Several transformer models are investigated, including *BERT*, *XLNet*, *GPT2* and *RoBERTa*. For each model, we conduct hyper-parameter search, where the search space includes variations of distillation [11], optimizers, learning rates, and batch sizes.

Hyper Parameter	Search Space	Final
Transformer	{BERT, XLNet, GPT2, RoBERTa}	BERT
Distillation	{True, False}	True
Optimizer	{AdamW, Adam, SGD}	ADAMW
Learning Rate	{5e-3, 2e-3, 5e-4, 2e-4, 5e-5, 2e-5}	5e-5
Batch Size	{16, 32, 64}	32

Table 1: Hyper-parameter search space and final values used

Table 1 shows the breakdown of the search space and the final hyper-parameters for CrisisBERT. Each set of parameters is randomly chosen and ran with 3 epochs and two trials. In total, we evaluate over 300 hyper-parameters sets using a Nvidia Titan-X (Pascal) for over 1,000 GPU hours.

Taking into consideration of performance and efficiency trade-off, we select the DistilBERT model for our Transformer LM layer. DistilBERT is a compressed version of BERT Base through Knowledge Distillation. With utilization of only 50% of the layers of BERT, DistilBERT performs 60% faster while preserving 97% of the capabilities in language understanding tasks. The optimal set of hyper-parameters for DistilBERT includes an AdamW [24] optimizer, and initial learning rate of 5e-5, and a batch size of 32.

Output Layer. The output layer of *DistilBERT* LM is a set of 768-d vectors led by the class header vector. Since we are conducting classification tasks, only the *[CLS]* token vector is used as the aggregate sequence representation for classification with a linear classifier. The remainder of the output vectors are processed into *Crisis2Vec* embeddings using Mean-Pooling operation.

2.4 Crisis2Vec

As discussed in Section 2.3, *Crisis2Vec* embedding is a byproduct of *CrisisBERT*, where the embeddings are constructed based on a pre-trained BERT model, and subsequently fine-tuned with three corpora of crisis-related tweets [30, 31, 50] to be domain-specific for crisis-related tweet representation.

Crisis2Vec leverages the advantages of Transformers, including (1) leveraging a self-attention mechanism to incorporate sentence-level context bidirectionally, (2) leveraging both word-level and positional information to create contextual representation of words, and (3) taking advantage of the pre-trained models on large relevant corpora.

Upon convergence, we construct the fixed-length tweet vector using a MEAN-Pooling strategy [37], where we compute the mean of all output vectors, as illustrated in Algorithm 1. The constructed *Crisis2Vec* is then used for our crisis classification tasks, which we elaborate next.

3 CRISIS CLASSIFICATION

In this work, we conduct two crisis classification tasks, namely *Crisis Detection* and *Crisis Recognition*. We formulate the *Crisis Detection* task as a binary classification model that identify if a tweet is relevant to a crisis-related event or not. The *Crisis Recognition*

Algorithm 1: *Crisis2Vec* with Mean Pooling

Input: *Tweet* as a sequence of *Words*
Output: *Tweet_Embedding*

```

1 Initialize Token_Vector;
2 for word in Tweet do
3   Word-Token  $\leftarrow$  Embed(Word);
4   AppendToken_Vector  $\leftarrow$  Word-Token;
5 end
6 Output_Vectors  $\leftarrow$  CrisisBERTLM(Token_Vector);
7 Tweet_Embedding  $\leftarrow$  PoolMean(Output_Vectors);
8 return Tweet_Embedding;
```

task on the other hand extends the problem into multi-class classification, where the output is a probability vector that indicate the likelihood of a tweet indicating specific events. Both tasks are modelled as Sequence Classification problems that are formally defined below.

3.1 Crisis Detection

We define the Crisis Detection task $D = (S, \Phi)$, which is specified by $S = \{s_1, \dots, s_n\}$ a finite sample space of tweets with size n . Each sample s_i is a sequence of tokens at T time steps, i.e., $s_i = \{s_i^1, \dots, s_i^T\}$. Φ denotes the set of labels that has the same sequence as the sample set, $\Phi = \{\phi_1, \dots, \phi_n\}$ and $\phi_i \in \{0, 1\}$ where $\phi_i = 1$ indicates that sample s_i is relevant to crisis, and $\phi_i = 0$ indicates otherwise. A deterministic classifier $C_D : S \rightarrow \phi$ specifies the mapping from sample tweets to their flags.

Our objective is to train a crisis detector using the provided tweets and labels that minimizes the differences between predicted labels and true labels, i.e.,

$$\min. \mathbb{J}_D(\Phi, C_D(S)) \quad (1)$$

where \mathbb{J}_D denotes a cost function.

3.2 Crisis Recognition

Similarly, we define a Crisis Recognition task $R = (S, L)$, where sample space S is identical to that in Crisis Detection. L denotes a sequence of multi-class labels that have the same sequence as S , i.e., $L = \{l_1, \dots, l_n\}$, and $l_i \in \mathbb{R}^m$ for m number of classes. A deterministic classifier $C_R : S \rightarrow L$ specifies the mapping from the sample tweets to the crisis classes.

The objective of the crisis classification tasks is to train a sequence classifier using the provided tweets and labels that minimizes the differences between predicted labels and true labels, i.e.,

$$\min. \mathbb{J}_R(L, C_R(S)) \quad (2)$$

where \mathbb{J}_R denotes a cost function for classifier C_R .

4 EXPERIMENTS AND RESULTS

In this section, we describe the various experiments performed and discuss the results of our proposed approach compared to various state-of-the-art baselines.

Label	Crisis event	# Data Points	Dataset	Label	Crisis event	# Data Points	Dataset
0	off-topic	30972	C36				
1	2012_Sandy_Hurricane	6318	C6	2	2013_Alberta_Floods	5189	C6
3	2013_Boston_Bombings	6577	C6	4	2013_Oklahoma_Tornado	4827	C6
5	2013_Queensland_Floods	6333	C6	6	2013_West_Texas_Explosion	6157	C6
7	sydneysiege	837	C8	8	charliehebdo	903	C8
9	ferguson	859	C8	10	germanwings-crash	248	C8
11	putinmissing	59	C8	12	ottawashooting	639	C8
13	ebola-essien	21	C8	14	prince-toronto	93	C8
15	2012_Colorado_wildfires	953	C26	16	2012_Costa_Rica_earthquake	909	C26
17	2012_Guatemala_earthquake	940	C26	18	2012_Italy_earthquakes	940	C26
19	2012_Philippines_floods	906	C26	20	2012_Typhoon_Pablo	907	C26
21	2012_Venezuela_refinery	939	C26	22	2013_Australia_bushfire	949	C26
23	2013_Bohol_earthquake	969	C26	24	2013_Brazil_nightclub_fire	952	C26
25	2013_Colorado_floods	925	C26	26	2013_Glasgow_helicopter_crash	918	C26
27	2013_LA_airport_shootings	912	C26	28	2013_Lac_Megantic_train_crash	966	C26
29	2013_Manila_floods	921	C26	30	2013_NY_train_crash	999	C26
31	2013_Russia_meteor	1133	C26	32	2013_Sardinia_floods	926	C26
33	2013_Savar_building_collapse	911	C26	34	2013_Singapore_haze	933	C26
35	2013_Spain_train_crash	991	C26	36	2013_Typhoon_Yolanda	940	C26

Table 2: Description of Crisis Classes

4.1 Datasets

We used three datasets of labelled crisis-related tweets [30, 31, 50] to conduct the crisis classification tasks and evaluate the proposed methods against the various benchmarks. In total, these datasets consist of close to 8 million tweets, where overall 91.6k are labelled. These data sets are in the form of: (1) 60k labelled tweets on 6 crises [30], (2) 3.6k labelled tweets for 8 crises [50], and (3) 27.9k labelled tweets for 26 crises [31]. Table 2 describes more detail about each dataset and their respective classes.

For our experimental evaluation, the 91.6k labelled crisis-related tweets are organized into two datasets, annotated as C6 and C36. In particular, C6 consists of 60k tweets from 6 classes of crises, whereas C36 comprises all 91.6k tweets in 36 classes. Both datasets are randomly split into training, validation and test sets that consist of 90%, 5% and 5% of the original sets, respectively. Table 3 describes the statistics of the split sets.

Dataset	# Data Points	# Classes	Train	Validation	Test
C6	60k	6 + 1	54,072	3,004	3,004
C36	91.6k	36 + 1	82,506	4,584	4,584

Table 3: Dataset Description

4.2 Proposed Models

Our proposed CrisisBERT and Crisis2Vec were previously described in Sections 2.3 and 2.4, respectively. In the subsequent sections, we describe how these algorithms will be evaluated against various baselines.

CrisisBERT. We evaluate the performance of *CrisisBERT* against multiple benchmarks, which comprise recently proposed crisis classification models in the literature. These works include linear classifiers, such as Logistic Regression (LR), Support Vector Machine (SVM) and Naive Bayes [25], and non-linear neural networks, such as Convolutional Neural Network (CNN) [16] and Long Short-Term Memory [18].

Furthermore, we investigate the robustness of *CrisisBERT* for both detection and recognition tasks. This aspect is studied by extending the experiments from C6 to C36, which comprise 6 and 36 classes respectively. The extension from 6 crises and 36 crises represents a 500% increase in crisis classes but with only 51.4% additional data points. We evaluate the robustness of the proposed models against benchmarks by observing the compromise in robustness performance, while realizing the drastically improved classification performance.

As described in Section 2.3, we use the optimal set of hyper-parameters for *CrisisBERT* in the experiments, which include the use of a BERT model with distillation (i.e. DistilBERT), an AdamW [24] optimizer, an initial learning rate of 5e-5, a batch size of 32, and a word dropout rate of 0.25.

Crisis2Vec. To evaluate *Crisis2Vec*, we choose the two classifiers with the aim to represent both traditional Machine Learning approaches and the NN approaches. The two selected models are: (1) a linear Logistic Regression model, denoted as LR_{c2v} , and (2) a non-linear LSTM model, denoted as $LSTM_{c2v}$. We evaluate the performance of *Crisis2Vec* with the two models by replacing the original embedding to *Crisis2Vec*, ceteris paribus.

4.3 Metrics

We use two common evaluation metrics, namely Accuracy and F1 score, which are functions of True-Positive (TP), False-Positive (FP), True-Negative (TN) and False-Negative (FN) predictions. Accuracy is calculated by:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (3)$$

For a F1-score of multiple classes, we calculate the unweighted mean for each label, i.e., for n classes of labels as:

$$F1 = \frac{2 * Recall * Precision}{Recall + Precision} \quad (4)$$

where

$$Precision = \frac{TP}{TP + FP} \quad (5)$$

and

$$Recall = \frac{TP}{TP + FN} \quad (6)$$

4.4 Benchmark Algorithms

We select and implement several crisis classifiers proposed in recent works to serve as benchmarks for evaluating our proposed methods. Concretely, we compare *CrisisBERT* with the following models:

- LR_{w2v} : Logistic regression model with *Word2Vec* embedding pre-trained on Google News Corpus [25].
- SVM_{w2v} : Support Vector Machine model with *Word2Vec* embedding pre-trained on Google News Corpus [25].
- NB_{w2v} : Naive Bayes model assuming Gaussian distribution for features with *Word2Vec* embedding pre-trained on Google News Corpus [25].
- CNN_{gv} : Convolutional Neural Network model with 2 convolutional layers of 128 hidden units, kernel size of 3, pool size of 2, 250 filters, and *GloVe* for word embedding [16].
- $LSTM_{w2v}$: Long Short-Term Memory model with 2 layers of 30 hidden states and a *Word2Vec*-based Crisis Embedding [18].

4.5 Results

Overall, the experimental results show that both proposed models achieve significant improvement on performance and robustness over the benchmarks across all tasks. The experimental results for *CrisisBERT* and *Crisis2Vec* are tabulated in Table 4.

Crisis Detection. For Crisis Detection tasks on C6 dataset, *CrisisBERT* achieves a 95.5% F1-score and 95.6% Accuracy, which exceeds that of the best performing benchmark model, namely *LSTM*

with pre-trained *Word2Vec* embeddings, by 3.8% and 3.8% respectively. In terms of embedding, *LSTM* with *Crisis2Vec* records 95.1% for both F1-score and Accuracy, which shows 3.4% improvement over *LSTM* with *Word2Vec*. Similarly, *LR* with *Crisis2Vec* records 93.2% for both F1-score and Accuracy, which shows 1.6% improvement over *LR* with *Word2Vec*.

For Crisis Detection tasks on C36 dataset, *CrisisBERT* achieves 94.2% F1-score and 94.7% Accuracy, which exceeds that of the best performing benchmark model, namely *CNN* with pre-trained *GloVe* embedding, by 2.7% and 2.1% respectively. In terms of embeddings, *LSTM* with *Crisis2Vec* records 92.8% F1-score and 93.5% Accuracy, which exceed *LSTM* with *Word2Vec* by 4.8% and 4.2%. Similarly, *LR* with *Crisis2Vec* records 89.0% F1-score and 90.1% Accuracy, which exceed *LR* with *Word2Vec* by 3.7% and 3.5% respectively.

Crisis Recognition. For Crisis Recognition tasks on C6 dataset, *CrisisBERT* achieves 98.7% F1-score and 98.6% Accuracy, which exceeds previous best model results, namely *CNN* with pre-trained *GloVe* embeddings, by 8.2% for both F1-score and Accuracy. In terms of embedding, *LSTM* with *Crisis2Vec* records 97.5% for both F1-score and Accuracy, which shows a significant improvement of 10.2% over *LSTM* with *Word2Vec*. Similarly, *LR* with *Crisis2Vec* records 93.6% F1-score and 93.7% Accuracy, which shows 6.3% and 5.2% improvement respectively over *LR* with *Word2Vec*.

For Crisis Detection tasks on C36 dataset, *CrisisBERT* achieves 97.1% F1-score and 97.9% Accuracy, which significantly exceeds previous best model results, namely *LR* with pre-trained *Word2Vec* embedding, by 25.0% and 15.6% respectively. In terms of embeddings, *LSTM* with *Crisis2Vec* records 88.0% F1-score and 95.6% Accuracy, which significantly exceed *LSTM* with *Word2Vec* by 29.1% and 23.3%. Similarly, *LR* with *Crisis2Vec* records 85.1% F1-score and 90.9% Accuracy, which significantly exceed *LR* with *Word2Vec* by 13.0% and 8.6% respectively.

Robustness. Comparing the Crisis Detection task between C6 and C36, *CrisisBERT* shows 1.7% and 1.3% decline for F1-score and Accuracy, which is much better than most benchmarks, i.e., from 1.7% to 6.3%, except *CNN*. However, when we compare the more challenging Crisis Recognition tasks between C6 and C36, the performance of *CrisisBERT* compromises marginally, i.e., 1.6% for F1-score and 0.7% for Accuracy. On the contrary, all benchmark models record significant decline, i.e. from 6.0% to 67.2%.

Discussion. Based on experimental results discussed above, we make the following observations:

- (1) *CrisisBERT*'s performance exceeds state-of-the-art performance for both detection and recognition tasks, by up to 8.2% and 25.0% respectively.
- (2) *CrisisBERT* demonstrates higher robustness with marginal decline for performance (i.e. less than 1.7% in F1-score and Accuracy).
- (3) *Crisis2Vec* shows superior performance as compared to conventional *Word2Vec* embeddings, for both *LR* and *LSTM* models across all experiments.

These results demonstrate the strong performance of *CrisisBERT* against the various baselines, and that of *Crisis2Vec* over the commonly used *Word2Vec* embeddings.

Models	Detection-C6		Detection-C36		Recognition-C6		Recognition-C36	
	F1	Accuracy	F1	Accuracy	F1	Accuracy	F1	Accuracy
Proposed								
<i>CrisisBERT</i>	95.5	95.6	94.2	94.7	98.7	98.6	97.1	97.9
<i>LSTM_{c2v}</i>	95.1	95.1	92.8	93.5	97.5	97.5	88.0	95.6
<i>LR_{c2v}</i>	93.2	93.2	89.0	90.1	93.6	93.7	85.1	90.9
Benchmark								
<i>LR_{w2v}</i> [25]	91.5	91.6	85.3	86.6	87.3	88.5	72.1	82.3
<i>SVM_{w2v}</i> [25]	91.3	91.4	85.0	86.3	86.6	87.9	71.6	81.9
<i>NB_{w2v}</i> [25]	86.8	86.8	82.4	83.4	78.8	80.5	47.6	63.5
<i>CNN_{gv}</i> [16]	91.2	91.3	91.1	91.2	90.5	90.4	23.3	64.4
<i>LSTM_{w2v}</i> [18]	91.7	91.7	88.0	89.3	87.3	87.3	58.9	72.3

Table 4: Experimental results of Crisis Classification tasks on C6 and C36 datasets for proposed models and benchmarks, where best performers are emphasized. Results show that *CrisisBERT* obtains highest performance across all tasks.

5 RELATED WORK

There is a wide range of work on general event detection on Twitter [3, 10, 38] and in this work, we focus our discussion of related work on those related to the two main areas of crisis classification and the use of embedding technique.

5.1 Crisis Classification

Event detection from tweets has received significant attention in research in order to analyze crisis-related messages for better disaster management and increasing situational awareness [14, 39–41]. Parilla-Ferrer *et al.* propose automatic binary classification of informativeness using Naive Bayes and Support Vector Machine (SVM) [32]. Stowe *et al.* present an annotation scheme for tweets to classify relevancy to various disaster events [44]. Furthermore, use of pre-trained word2vec reportedly improved SVM for Crisis classification [26]. Lim *et al.* propose ClusTop algorithm utilizing the community detection approach for automatic topic modelling [21]. Pedrood *et al.* propose to transfer-learn classification of one event to the other using a sparse coding model [33], though the scope was only limited to only two events, i.e. Hurricane Sandy (2012) and Supertyphoon Yolanda (2013).

A substantial number of works focuses on using Neural Networks (NN) with word embeddings for crisis-related data classification. Manna *et al.* [25] compared NN models with conventional ML classifiers [25]. ALRashdi and O’Keefe investigated and showed good performance for two deep learning architectures, namely Bidirectional Long Short-Term Memory (BiLSTM) and Convolutional Neural Networks (CNN) with domain-specific GloVe embeddings [2]. However, the study had yet to validate the relevance of model on a different crisis type. Nguyen *et al.* applied CNN to classify information types using Google News and domain-specific embeddings [28].

Kersten *et al.* [16] implemented a parallel CNN to detect two disasters, namely hurricanes and floods, which reported a F1-score of 0.83. The CNN architecture was proposed earlier by Kim *et al.* [17].

5.2 Crisis Embedding

Word-level Embeddings such as *Word2Vec* [4] and *GloVe* [34] are commonly used to form the basis of Crisis Embedding [28, 29] in various crisis classification works to improve model performance. For context, *Word2Vec* uses a Neural Network Language Model (NNLM) that is able to represent latent information on the word level. *GloVe* achieved better results with a simpler approach, constructing global vectors to represent contextual knowledge of the vocabulary.

More recently, a series of high quality embedding models, such as FastText [5] and Flair [1], are proposed and reported to have improved state of the art for multiple NLP tasks. Both word-level contextualization and character-level features are commonly used for these works. Pre-trained models on large corpora of news and tweets collections are also made publicly available to assist in downstream tasks. Furthermore, Transformer-based models are proposed to conduct sentence-level embedding tasks [37].

6 CONCLUSION

Social media such as Twitter has become a hub of crowd generated information for early crisis detection and recognition tasks. In this work, we present a transformer-based crisis classification model *CrisisBERT*, and a contextual crisis-related tweet embedding model *Crisis2Vec*. We examine the performance and robustness of the proposed models by conducting experiments with three datasets and two crisis classification tasks. Experimental results show that *CrisisBERT* improves state of the art for both detection and recognition class, and further demonstrates robustness by extending from 6

classes to 36 classes, with only 51.4% additional data points. Finally, our experiments with two classification models show that *Crisis2Vec* enhances classification performance as compared to *Word2Vec* embeddings, which is commonly used in prior works.

ACKNOWLEDGMENTS

This research is funded in part by the Singapore University of Technology and Design under grant SRG-ISTD-2018-140. The authors thank the anonymous reviewers for their useful comments.

REFERENCES

- [1] Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual String Embeddings for Sequence Labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*. 1638–1649.
- [2] Reem ALRashdi and Simon O’Keefe. 2019. Deep Learning and Word Embeddings for Tweet Classification for Crisis Response. *arXiv preprint arXiv:1903.11024* (2019).
- [3] Farzindar Atefeh and Wael Khreich. 2015. A survey of techniques for event detection in twitter. *Computational Intelligence* 31, 1 (2015), 132–164.
- [4] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research* 3, Feb (2003), 1137–1155.
- [5] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [6] Grégoire Burel, Hassan Saif, Miriam Fernandez, and Harith Alani. 2017. On semantics and deep learning for event detection in crisis situations. (2017).
- [7] Justin Michael Crow. 2020. Verifying baselines for crisis event information classification on Twitter. In *Proceedings of the 17th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2020)*. ISCRAM Digital Library.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [9] Yasmeen George, Shanika Karunasekera, Aaron Harwood, and Kwan Hui Lim. 2021. Real-time Spatio-temporal Event Detection on Geotagged Social Media. *Journal of Big Data* 8, 91 (Jun 2021), 1–28.
- [10] Mahmud Hasan, Mehmet A Orgun, and Rolf Schwitter. 2018. A survey on real-time event detection from the twitter data stream. *Journal of Information Science* 44, 4 (2018), 443–463.
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [12] Yuqian Huang, Yue Li, and Jie Shan. 2018. Spatial-Temporal Event Detection from Geo-Tagged Tweets. *ISPRS International Journal of Geo-Information* 7, 4 (2018), 150.
- [13] Amanda L Hughes, Lise AA St. Denis, Leysia Palen, and Kenneth M Anderson. 2014. Online public communications by police & fire services during the 2012 Hurricane Sandy. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1505–1514.
- [14] Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)* 47, 4 (2015), 1–38.
- [15] Muhammad Imran, Shady Elbassuoni, Carlos Castillo, Fernando Diaz, and Patrick Meier. 2013. Extracting information nuggets from disaster-related messages in social media. In *Proceedings of the 10th Proceedings of the International Conference on Information Systems for Crisis Response and Management (ISCRAM’13)*.
- [16] Jens Kersten, Anna Kruspe, Matti Wiegmann, and Friederike Klan. 2019. Robust filtering of crisis-related tweets. In *ISCRAM 2019 Conference Proceedings-16th International Conference on Information Systems for Crisis Response and Management*.
- [17] Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [18] Abhinav Kumar, Jyoti Prakash Singh, Yogesh K Dwivedi, and Nripendra P Rana. 2020. A deep multi-modal neural network for informative Twitter content classification during emergencies. *Annals of Operations Research* (2020), 1–32.
- [19] Jolin Shaynn-Ly Kwan and Kwan Hui Lim. 2020. Understanding Public Sentiments, Opinions and Topics about COVID-19 using Twitter. In *Proceedings of the 2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM’20)*. 623–626.
- [20] Lung-Hao Lee, Yi Lu, Po-Han Chen, Po-Lei Lee, and Kuo-Kai Shyu. 2019. NCUEE at MEDIQA 2019: Medical Text Inference Using Ensemble BERT-BiLSTM-Attention Model. In *Proceedings of the 18th BioNLP Workshop and Shared Task*. 528–532.
- [21] Kwan Hui Lim, Shanika Karunasekera, and Aaron Harwood. 2017. Clustop: A clustering-based topic modelling algorithm for twitter using word networks. In *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 2009–2018.
- [22] Junhua Liu, Yung Chuen Ng, Kristin L. Wood, and Kwan Hui Lim. 2020. IPOD: A Large-scale Industrial and Professional Occupation Dataset. In *Proceedings of the 2020 ACM Conference on Computer Supported Cooperative Work and Social Computing Companion (CSCW’20)*. 323–328.
- [23] Junhua Liu, Trisha Singhal, Lucienne T.M. Blessing, Kristin L. Wood, and Kwan Hui Lim. 2020. EPIC30M: An Epidemics Corpus of Over 30 Million Relevant Tweets. In *Proceedings of the 2020 IEEE International Conference on Big Data (BigData’20)*. 1206–1215.
- [24] Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *Proceedings of the 7th International Conference on Learning Representations (ICLR’17)*.
- [25] Sukanya Manna and Haruto Nakai. 2019. Effectiveness of Word Embeddings on Classifiers: A Case study with tweets. In *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*. IEEE, 158–161.
- [26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [27] Nic Newman. 2011. Mainstream media and the distribution of news in the age of social media. (2011).
- [28] Dat Tien Nguyen, Kamela Ali Al Mannai, Shafiq Joty, Hassan Sajjad, Muhammad Imran, and Prasenjit Mitra. 2017. Robust classification of crisis-related data on social networks using convolutional neural networks. In *Eleventh International AAAI Conference on Web and Social Media*.
- [29] Dat Tien Nguyen, Shafiq Joty, Muhammad Imran, Hassan Sajjad, and Prasenjit Mitra. 2016. Applications of online deep learning for crisis response using social media information. (2016).
- [30] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. Crisislex: A lexicon for collecting and filtering microblogged communications in crises. In *Eighth international AAAI conference on weblogs and social media*.
- [31] Alexandra Olteanu, Sarah Vieweg, and Carlos Castillo. 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.
- [32] Beverly Estephany Parilla-Ferrer, PL Fernandez, and JT Ballena. 2014. Automatic classification of disaster-related tweets. In *Proc. International conference on Innovative Engineering Technologies (ICIET)*, Vol. 62.
- [33] Bahman Pedrood and Hemant Purohit. 2018. Mining help intent on twitter during disasters via transfer learning with sparse coding. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*. Springer, 141–153.
- [34] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [35] Sasa Petrovic, Miles Osborne, Richard McCreadie, Craig Macdonald, Iadh Ounis, and Luke Shrimpton. 2013. Can twitter replace newswire for breaking news?. In *Proceedings of the international AAAI conference on web and social media*, Vol. 7.
- [36] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019).
- [37] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3980–3990.
- [38] Zafar Saeed, Rabeeh Ayaz Abbasi, Onaiza Maqbool, Abida Sadaf, Imran Razzak, Ali Daud, Naif Radi Aljohani, and Guandong Xu. 2019. What’s happening around the world? A survey and framework on event detection techniques on Twitter. *Journal of Grid Computing* 17, 2 (2019), 279–312.
- [39] Naina Said, Kashif Ahmad, Michael Riegler, Konstantin Pogorelov, Laiq Hassan, Nasir Ahmad, and Nicola Conci. 2019. Natural disasters detection in social media and satellite imagery: a survey. *Multimedia Tools and Applications* 78, 22 (2019), 31267–31302.
- [40] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*. 851–860.
- [41] Luke S Snyder, Morteza Karimzadeh, Christina Stober, and David S Ebert. 2019. Situational Awareness Enhanced through Social Media Analytics: A Survey of First Responders. In *Proceedings of the 2019 IEEE International Symposium on Technologies for Homeland Security (HST)*.
- [42] Prashant Solanki, Kwan Hui Lim, and Aaron Harwood. 2021. User Identification across Social Networking Sites using User Profiles and Posting Patterns. In *Proceedings of the 2021 International Joint Conference on Neural Networks (IJCNN’21)*.

- [43] Internet Live Statistics. 2020. Twitter usage statistics. *Internet* (2020). <https://www.internetlivestats.com/twitter-statistics/>
- [44] Kevin Stowe, Michael Paul, Martha Palmer, Leysia Palen, and Kenneth M Anderson. 2016. Identifying and categorizing disaster-related tweets. In *Proceedings of The fourth international workshop on natural language processing for social media*. 1–6.
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [46] Yuan Wang and Yiyi Yang. 2020. Dialogic communication on social media: How organizations use Twitter to build dialogic relationships with their publics. *Computers in Human Behavior* 104 (2020), 106183.
- [47] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*. 5754–5764.
- [48] Chao Zhang, Dongming Lei, Quan Yuan, Honglei Zhuang, Lance Kaplan, Shaowen Wang, and Jiawei Han. 2018. Geoburst+: effective and real-time local event detection in geo-tagged tweet streams. *ACM Transactions on Intelligent Systems and Technology* 9, 3 (2018), 34.
- [49] Shanshan Zhang and Slobodan Vucetic. 2016. Semi-supervised discovery of informative tweets during the emerging disasters. *arXiv preprint arXiv:1610.03750* (2016).
- [50] Arkaitz Zubiaga, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Peter Tolmie. 2016. Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PloS one* 11, 3 (2016).