Detecting High-engaging Breaking News Rumors in Social Media

SARAH A. ALKHODAIR, College of Computer and Information Sciences, King Saud University, Saudi Arabia
BENJAMIN C. M. FUNG, School of Information Studies, McGill University, Canada
STEVEN H. H. DING, School of Computing, Queen's University, Canada
WILLIAM K. CHEUNG, Department of Computer Science, Hong Kong Baptist University, Hong Kong
SHIH-CHIA HUANG, Department of Electronic Engineering, National Taipei University of Technology, Taiwan

Users from all over the world increasingly adopt social media for newsgathering, especially during breaking news. Breaking news is an unexpected event that is currently developing. Early stages of breaking news are usually associated with lots of unverified information, i.e., rumors. Efficiently detecting and acting upon rumors in a timely fashion is of high importance to minimize their harmful effects. Yet, not all rumors have the potential to spread in social media. High-engaging rumors are those written in a manner that ensures achievement of the highest prevalence among the recipients. They are difficult to detect, spread very fast, and can cause serious damages to society. In this paper, we propose a new multi-task CNN-attention-based neural network architecture to jointly learn the two tasks of breaking news rumors detection and breaking news rumors popularity prediction in social media. The proposed model learns the salient semantic similarities among important features for detecting high-engaging breaking news rumors and separates them from the rest of the input text. Extensive experiments on five real-life datasets of breaking news suggest that our proposed model outperforms all baselines and is capable of detecting breaking news rumors and predicting their future popularity with high accuracy.

CCS Concepts: • **Computing methodologies → Natural language processing**; **Machine learning**; **Supervised learning by classification**; **Neural networks**; • **Information systems** → *Social networking sites.*

Additional Key Words and Phrases: social media, deep learning, breaking news, rumor detection

ACM Reference Format:
Sarah A. Alkhodair, Benjamin C. M. Fung, Steven H. H. Ding, William K. Cheung, and Shih-Chia Huang. 20??. Detecting High-engaging Breaking News Rumors in Social Media. *ACM Trans. Manag. Inform. Syst.* ?, ?, Article ? (? 20??), 10 pages. https://doi.org/10.1145/3416703

1 INTRODUCTION

Social media highly impacts people's knowledge and perception of the world. The convenience, accessibility of real-time information, and the diversity of the available sources have attracted more people to gather their news in social media every day [Matsa and Shearer 2018]. In fact, the global median for getting news from social media every day has become 42% in 2018 [Mitchell et al. 2018]. The lack of fact-checking and source-verification in social media facilitates the spread of rumors. When these rumors become viral, they may result in extremely damaging consequences in just a few minutes.

A rumor is "a story or a statement whose truth value is unverified" [Allport and Postman 1965]. Accordingly, a rumor does not have to be false; it can be deemed later to be true or false. There are two categories of rumors based on their temporal characteristics [Zubiaga et al. 2018]: *long-standing rumors* that are well-discussed for long periods of time, and *breaking news rumors* that are generally unseen before and emerge extremely fast during the breaking news evolution. Breaking news refers to "newly received information about an event that is currently occurring or developing"[1]. In contrast to regular news, breaking news dramatically evolves over time with no sufficient details about what happened and what will happen, and is associated with unexpected sequence of sub-topics that mostly do not occur in existing data.

Early stages of breaking news diffusion are usually associated with many rumors. In fact, the volume of rumors is directly proportional to the importance of and interest in the topic to individuals [Allport and Postman 1965]. Detecting and acting upon such rumors in a timely fashion to minimize their harmful effect is an extremely difficult and crucial task. Yet, not all rumors have the potential to spread in social media. High-engaging breaking news rumors are those written in a manner that ensures they achieve the highest prevalence among the recipients. These rumors are difficult to detect and have the potential to become extremely viral in social media for several reasons. First, the mental state of recipients during breaking news is one that is ready to accept any information without thinking or analysis its contents [Muhammad 2017]. Furthermore, during breaking news emergencies, people closely follow up with any information update regarding its current development. Moreover, emotions often govern the sharing act in social media [Berger and Milkman 2009]. Rumors are intended to touch and satisfy the primal emotional needs of recipients, such as fear, anger, or anxiety [Muhammad 2017]. More importantly, rumors are intentionally written to mimic how verified information is reported. Thus, in addition to the compelling writing, they are believable, expressive, informative, and answer questions people want to know [Crescimbene et al. 2012; Difonzo and Bordia 2007].

Detecting high-engaging breaking news rumors in social media helps prioritize the rumors verification process during breaking

Authors' addresses: Sarah A. Alkhodair, salkhudair@ksu.edu.sa, College of Computer and Information Sciences, King Saud University, Riyadh, Saudi Arabia; Benjamin C. M. Fung, ben.fung@mcgill.ca, School of Information Studies, McGill University, Montreal, Canada; Steven H. H. Ding, ding@cs.queensu.ca, School of Computing, Queen's University, Kingston, Canada; William K. Cheung, william@comp.hkbu.edu.hk, Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong; Shih-Chia Huang, schuang@ntut.edu.tw, Department of Electronic Engineering, National Taipei University of Technology, Taipei 106, Taiwan.

2158-656X/20??/?-ART? $15.00
https://doi.org/10.1145/3416703

[1]Source: https://en.oxforddictionaries.com/definition/breaking_news, retrieved on Feb 3, 2019

news emergencies. Flagging rumors that are expected to become viral in social media helps authorities react fast toward such rumors and curb/slow their spread, which in turn reduces their damaging consequences." However, this is a challenging task. First, breaking news covers topics that may not exist in the training dataset. The existing data may also lack similar, or related topics. In this case, detecting high-engaging breaking news rumors requires zero-shot learning for real-time detection. Second, a breaking news rumor does not require the existence of explicit links between recipients to propagate throughout social media networks. In this case, the task of predicting the popularity of a rumor micro-post based on the structure of the social network is not applicable. Also, in contrast to long-standing rumors, breaking news rumors have a distinctive life-cycle where a single verification statement from the authorities can quickly curb the diffusion of a breaking news rumor [Ozturk et al. 2015; Takayasu et al. 2015]. Thus, using early observations of a rumor's dynamics might not give a reliable estimation of its future popularity. Also, handling breaking news rumors is a time-critical task. In such a case, even waiting for a few minutes for enough data to be available might render the results of rumor detection and popularity prediction models useless in terms of minimizing the harmful effect of rumors. This is because more than 50% of the sharing act happens within the first ten minutes after posting the micro-post in social media [Zaman et al. 2014]. This percentage becomes much higher during breaking news and emergencies. Thus, after a few minutes, the damaging consequences of a high-engaging breaking news rumor is more likely to have already happened. Finally, the characteristics of high-engaging rumors are very much in line with high-engaging posts in social media in general. This makes the process of distinguishing high-engaging rumors from high-engaging non-rumours much more challenging than detecting rumors in general.

To address these challenges, we propose a multi-task deep learning model that can incrementally learn the shared latent features among the two tasks of *breaking news rumors detection* and *breaking news rumors popularity prediction* and use these features to train the model with the objective of detecting high-engaging breaking news rumors in social media. In our design, we jointly train a *word embedding learning* model, *word2vec* [Mikolov et al. 2013], with an unsupervised objective to learn the word embedding and train a multi-task deep learning model with a supervised objective of high-engaging breaking news rumor detection. We propose to train a word2vec model on the fly with the input of a high-engaging breaking news rumor detection model. Keeping the word2vec model parallel to the detection model and using it to update the embedding space help the proposed model to incrementally learn the distributed vector representations of words in the input text, capture the deep latent features and their correlations from it, and use them to build the detection model. This helps our proposed model better handle the issues of emerging topics of breaking news such as new *Out-Of-Vocabulary (OOV)* and topic-shifts that do not exist in the training data. Furthermore, we use a *Convolutional Neural Network (CNN)* model and a *Self-attention mechanism* as shared feature extractors in the model. The CNN model helps capture several classes of semantic features while the attention mechanism guides the model to weight differently to the input sequence and locates important features for predicting the final class. This helps the proposed model learn the salient semantic similarities among important words and phrases

In this paper, we tackle the problem of detecting breaking news rumors that are most likely to achieve high engagement rates in social media. The main contributions of this work are summarized as follow. **(1)** This is the first work that tackles the problem of detecting high-engaging breaking news rumors in social media. **(2)** We propose a new multi-task CNN-attention-based neural network architecture to *jointly* learn the two tasks of breaking news rumors detection and breaking news rumors popularity prediction in social media. The proposed model learns the salient semantic similarities among important features for detecting high-engaging breaking news rumors and separates them from the rest of the input text. **(3)** Extensive experiments on five real-life breaking news datasets suggest that our proposed model is capable of detecting high-engaging breaking news rumors in social media, and it outperforms all baselines in terms of precision, recall, and F1.

Most existing work on rumor detection suffers from two issues. First, it focuses on long-standing rumors rather than breaking news rumors and aims at tracking the diffusion of rumors, classifying opinions expressed toward them, or predicting their veracity. Second, it assumes that rumors are always false and proposes models to detect these false rumors [Zubiaga et al. 2018]. In contrast, this work aims at detecting breaking news rumors, rather than long-standing rumors, regardless of their truth value. The goal is to flag high-engaging micro-post rumors during the rapid diffusion of breaking news to help reduce their damaging consequences. Most existing work on popularity prediction estimates the future popularity of a micro-post based on the early observations of its dynamics, the network structure, or both. In contrast, our proposed model does not need a collection of early observations nor is it based on the network structure.

2 RELATED WORK

2.1 Rumor Detection

Rumor detection and analysis in social media has been an attractive field of study in the last few years. Existing work in this field falls in one of four categories: rumor detection, rumor tracking, rumor stance classification, and rumor veracity classification [Zubiaga et al. 2018]. Rumor detection is the first and most important task. The goal is to detect which unverified information is spreading across social media. Yet, there has been very little work in this field. The first rumor detection method was proposed by Zhao et al. [Zhao et al. 2015b]. The proposed method starts by detecting "signal tweets". These tweets are then grouped into different clusters, each representing a rumor. Next, each cluster is summarized and the summary is used to retrieve more related tweets. Finally, the clusters are ranked by their likelihood of being rumors. Their method is based entirely on using a list of user-defined regular expressions to detect the "signal tweets". Thus, for their method to better handle new unseen stories, this list needs to be revised periodically. Zubiaga et al. [Zubiaga et al. 2016b] proposed a rumor detection model based on a sequential classifier. The proposed model classifies a tweet as a rumor or non-rumor based on previously encountered data. This method achieves higher performance than the previous work [Zhao et al. 2015b]. However, it suffers from the cold start problem where the performance of the proposed sequential classifier model on detecting a breaking news rumor depends on the sequence of related micro-posts encountered so far [Zubiaga et al. 2016b]. Alkhodair et al. [Alkhodair et al. 2019] proposed a semi-supervised model that employs representation learning and deep learning models to learn and exploit the lexical and temporal features of rumor micro-posts. The proposed model outperforms the state-of-the-art model [Zubiaga et al. 2016b] in detecting breaking news rumors in terms of accuracy. However, their work aims at detecting breaking news rumors in social media regardless of their popularity. Overwhelming authorities with huge volumes of new rumors every day may reduce their ability to handle them in a timely-fashion, especially during breaking news emergencies. In contrast, our work aims at both detecting breaking news rumors and predicting which of them are most likely to become viral in social media, therefore, need immediate attention.

2.2 Popularity Prediction

Most existing work in this field of study investigates the problem of predicting the future popularity of a micro-post by observing its dynamics for some time after it is posted. For example, Zaman et al. [Zaman et al. 2014] proposed to use network information and the time path of previous retweets to predict the future popularity of a tweet. Also, Zhao et al. [Zhao et al. 2015a] proposed a model to predict the future retweet intensity of a tweet as the product of its "infectivity" and the excitation effect of all of its previous retweets. Similarly, Mishra et al. [Mishra et al. 2016] proposed to combine features-based approaches with point process models to predict the future popularity of a tweet based on its previous retweets. Other proposed models predict the future popularity of a tweet by observing its retweeting dynamics of the first day [Yan et al. 2016] and the first two hours [Chen and Tan 2018] after it is posted. All these methods predict the future popularity of a micro-post based on the early observations of its dynamics in social media. This requires monitoring its popularity for some time after it is posted to gather sufficient observations for a reliable prediction. Such methods are not applicable when dealing with breaking news and emergencies. In contrast, our proposed model does not need a collection of early observations to predict the future popularity of a rumor micro-post.

There are other lines of research that investigate the problem of predicting the popularity of a news article prior to publishing. For example, Bandari et al. [Bandari et al. 2012] used Twitter data to predict the future popularity of a news article before it is published. Similarly, Abbar et al. [Abbar et al. 2018] proposed to predict the popularity of a new news article based on the recent popularity of its topic and similar articles. In addition to dealing with news articles rather than the short text of social media micro-posts, this line of work requires a collection of related posts, similar articles, or similar topics which does not exist in the case of breaking news. In contrast, our proposed model does not need a collection of related posts or topics to predict the future popularity of a rumor micro-post.

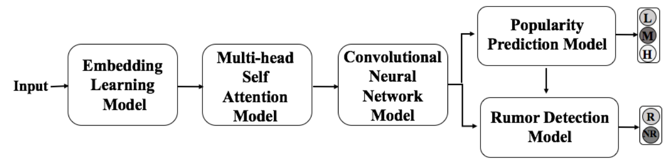

Fig. 1. The proposed joint learning model for detecting high-engaging breaking news rumors in social media

3 JOINT LEARNING MODEL FOR DETECTING HIGH-ENGAGING BREAKING NEWS RUMORS

3.1 Problem Statement

High-engaging breaking news rumors are intentionally written in a way that satisfies the primal emotional needs of recipients and mimics the writing style of verified information. Thus, we also used different supportive feature sets, such as stylometric and emotional triggers features, as inputs to our proposed model to help improve its detection performance. The research problem of detecting high-engaging breaking news rumors can be defined as follows. Let $mp = \langle w_1, ..., w_T \rangle$ be a sequence of words in a micro-post mp of length T, and let F be the set of its associated supportive features. Given mp and F as inputs, the goal is to simultaneously classify mp as either a rumor or a non-rumor by assigning a label from $RL = \{R, NR\}$ and predicting the engagement rate it will achieve by assigning a label from $PL = \{High, Moderate, Low\}$.

3.2 Proposed Model

Figure 1 illustrates an overview of our proposed joint learning model. The proposed model consists of five components, namely the embedding learning model, the multi-head self-attention model, the convolutional neural network model, the popularity prediction model, and the rumor detection model.

3.2.1 Embedding Learning Model. This model takes a text corpus as an input and produces real-valued low-dimensional vector representations for words that appear in that corpus. These vector representations are called *word embeddings*. In this work, we train a *skip-gram* word2vec model [Mikolov et al. 2013] to convert the sequence of words in each micro-post mp into a sequence of word embedding vectors $X = \langle x_1, \ldots, x_T \rangle$ that is passed to the subsequent multi-head self-attention model. In our work, we keep the word2vec model parallel to the detection model and use it to update the embedding space on the fly to help the proposed model incrementally learn the distributed vector representations of words in the input text. Alkhodair et al. [Alkhodair et al. 2019] have shown that such training strategy significantly outperforms other embedding training strategies when dealing with breaking news rumors and can help the training process overcome the challenges of the cross-topic and the OOV in breaking news rumors detection."

3.2.2 Multi-head Self-attention Model. Attention mechanisms have been used mainly to guide a deep learning model to attend differently to the input sequence and locate important features to predict the final class. Recently, a *Self-attention* mechanism has been proposed as a part of a machine translator architecture called the *Transformer* [Vaswani et al. 2017]. *Self-attention* mechanism works as follows. First, it calculates the dot-product of a weight matrix

$W_{attn} \in \mathbb{R}^{\mathcal{D} \times 3\mathcal{D}}$ by each word embedding x_i in X. Next, it splits the result through dimension to generate three matrices of size $\mathcal{D}$ known as query $\mathcal{Q}$, key $\mathcal{K}$, and value $\mathcal{V}$ matrices. Finally, the attention is calculated as follows:

$$Self\text{-}attention(\mathcal{Q}, \mathcal{K}, \mathcal{V}) = softmax(\frac{\mathcal{Q}\mathcal{K}^\top}{\sqrt{\mathcal{D}}})\mathcal{V} \tag{1}$$

In our work, we use a *Multi-head Self-attention* mechanism to allow the learning model to "jointly attend to information from different representation subspaces at different positions" [Vaswani et al. 2017]. *Multi-head Self-attention* mechanism works as follows. First, the queries, keys, and values are linearly projected $\mathcal{H}$ times with different learned linear projections to $\mathcal{D}_\mathcal{K}$, $\mathcal{D}_\mathcal{K}$, and $\mathcal{D}_\mathcal{V}$ dimensions, respectively [Vaswani et al. 2017]. Next, the self-attention value for each projected version, i.e., head, is calculated as follows:

$$head_i = Self\text{-}attention(\mathcal{Q}\mathcal{P}_i^\mathcal{Q}, \mathcal{K}\mathcal{P}_i^\mathcal{K}, \mathcal{V}\mathcal{P}_i^\mathcal{V}) \tag{2}$$

where $\mathcal{P}_i^\mathcal{Q} \in \mathbb{R}^{\mathcal{D} \times \mathcal{D}_\mathcal{K}}$, $\mathcal{P}_i^\mathcal{K} \in \mathbb{R}^{\mathcal{D} \times \mathcal{D}_\mathcal{K}}$ and $\mathcal{P}_i^\mathcal{V} \in \mathbb{R}^{\mathcal{D} \times \mathcal{D}_\mathcal{V}}$ are the projection parameters. Finally, these attentions are concatenated and projected again to compute the final multi-head attention as follows:

$$Multi\text{-}headAttn(\mathcal{Q}, \mathcal{K}, \mathcal{V}) = concat(head_1, \ldots, head_\mathcal{H})\mathcal{P}^O \tag{3}$$

where $\mathcal{P}^O \in \mathbb{R}^{\mathcal{H}\mathcal{D}_\mathcal{V} \times \mathcal{D}}$ is the projection parameter. By the end of this model, the sequence of word embedding vectors $X = \langle x_1, \ldots, x_T \rangle$ of each micro-post is converted into a sequence of attended-vectors $X^{att} = \left\langle x_1^{att}, \ldots, x_T^{att} \right\rangle$. X^{att} is then transformed into a matrix representation $\mathcal{A} \in \mathbb{R}^{T * \mathcal{D}}$ to be passed to the convolutional neural network model.

3.2.3 Convolutional Neural Network Model. *Convolutional Neural Networks (CNNs)* are feed forward neural networks that typically consist of convolutional layers followed by pooling layers. We chose to use CNNs rather than *Recurrent Neural Networks (RNNs)* for our model due to the nature of textual data in social media websites which imposed more challenges on the high-engaging breaking news rumor detection task. A micro-post in social media websites is usually very short, noisy, incomplete, and rarely follows the correct grammatical structure of a sentence. RNNs can model long sequential context dependencies in a text. They are trained to capture textual patterns across time. However, one limitation of RNNs is that later words in a sentence (micro-post) will have more influence on the final sentence representation than formers words [Zhou et al. 2017]. This is a problem when dealing with textual data in social media since important words can appear anywhere in a micro-post. CNNs overcome this limitation by assigning the same level of importance to each word in a sentence regardless of its position [Minaee et al. 2020; Zhou et al. 2017]. Furthermore, several recent studies have shown that CNN's performance is better than RNN's on classification tasks that depend on the recognition of key-phrases such as sentiment analysis and question-answer matching problems [Hu et al. 2018; Minaee et al. 2020; Yin et al. 2017]. For the task of rumor detection of breaking news events that have not been seen before, the classification task mainly depends on the detection of the important key-phrases in the text such as confusion, wondering, angry, and sadness terms rather than modeling of the context dependencies used in RNNs. For example, given a micro-post about a breaking news event, key-phrases such as "Just heard that", "Is this true?", "unbelievable", or "Check this out" can be sufficient to classify a micro-post about a previously unseen topic as a rumor or not. Furthermore, CNNs are faster than RNNs which is a major advantage in the case of high-engaging breaking news rumor detection.

In our work, a CNN is used as a feature extractor to extract the latent features from the input textual data as follows. First, the convolutional layers apply convolutional filters over the input matrix to produce different feature maps. Then, these feature maps are fed through pooling layers to induce a fixed length features vector of the micro-post. By varying the size of the convolutional filters in CNNs, the model can detect patterns of different sizes (n-grams) regardless of their position in a micro-post. This helps the model overcome the problem of unstructured textual data in social media and capture the important word patterns in a micro-post.

Formally, given an input matrix $\mathcal{A} \in \mathbb{R}^{T * \mathcal{D}}$ that represents a micro-post consisting of T words, each represented by a $\mathcal{D}$-dimensional vector of real values, the convolutional layer will repeatedly apply the linear filters on sub-matrices of $\mathcal{A}$ as follows [Zhang and Wallace 2017]:

$$O_{cnv}^i = \mathbf{W} \cdot \mathcal{A}\,[i : i + r - 1] \tag{4}$$

where O_{cnv}^i is the output of the convolutional operator after applying the i^{th} filter, r is the region size or the height of the filter, $\mathbf{W}$ is the weight matrix of the filter, and $\mathcal{A}\,[i : i + r - 1]$ represents the sub-matrix of $\mathcal{A}$ from row i to row $i + r - 1$. A feature map c_i of the i^{th} filter is then calculated as follows:

$$c_i = \mathbf{a}\left(O_{cnv}^i + \mathbf{b}\right) \tag{5}$$

where $\mathbf{a}$ is the activation function and $\mathbf{b} \in \mathbb{R}$ is a bias term. The feature maps C are then fed into a 1-max pooling layer to generate a univariate feature vector from each feature map c_i. The univariate feature vectors are then concatenated to form a fixed-size feature vector FV to be passed to the subsequent models for the final prediction task.

3.2.4 Popularity Prediction Model. This model learns the future popularity of a micro-post in social media as a function of its *Engagement Rate*, which is a widely-adopted measure to evaluate the quality of a micro-post on different social media platforms [Comcowich 2015; Socialbakers 2013]. Let the *EngVolume* be the total impact of a micro-post *mp*, calculated as the total count of *likes*, *shares*, and *comments* received by *mp*; let the *BaseVolume* be the number of users with direct exposure to *mp*. The engagement rate is calculated as follows:

$$EngRate(mp) = \frac{EngVolume(mp)}{BaseVolume(mp)}, \tag{6}$$

Since it is difficult to perform classification or regression for the whole engagement rate range, we reformulate the popularity prediction task as a multi-classification task that assigns each micro-post *mp* to one of three levels of popularity as follows: **low** if $EngRate(mp) < 0.02\%$, **moderate** if $0.02\% < EngRate(mp) < 0.33\%$, and **high** if $< EngRate(mp) > 0.33\%$. The popularity levels were defined based on analyzing millions of influencers' accounts in Twitter by *Scrunch* [2].

This model learns the popularity prediction task as follows. First, it takes the feature vector FV generated by the convolutional neural network model and passes it through a fully connected layer to generate an internal feature vector of the popularity prediction model, namely FV_p. Then, FV_p is passed through a softmax layer with *sigmoid* function to predict the popularity level of a micro-post. Although sigmoid is not an ideal choice for multi-class classification task, we used it to relax the classification problem and allow the model to learn the features of each popularity level independently. The intuition behind this choice is the lack of enough data and the nature of breaking news events. Since we reformulated the task of popularity prediction into a multi-class classification task instead of performing the regression on the whole engagement rate range, a micro-post might exist on the border that separates two classes/levels of popularity. Micro-posts that sit on both sides of the border have very similar features in regard to the popularity-level it will achieve. Allowing the model to learn the features of each class independently helps it overcome the issue of penalizing one class to predict the other. This is crucial to the problem of detecting high-engaging breaking news rumors in the sense that a rumor micro-post can be very dangerous if it is ignored because it was classified to belong to the moderate-popularity level rather than the high-popularity level. The sigmoid function generates a probability output of each popularity-level independently. This gives the model more flexibility on giving rumors around the border high probabilities for both classes. The model will, for each micro-post, return the class with the maximum probability output as the predicted class and, at the same time, can flag a rumor to have a high-probability to belong to another class. The objective function is optimized by minimizing the following cross-entropy popularity prediction loss function L_P:

$$L_P = -\frac{1}{\mathsf{N}} \sum_{i=1}^{\mathsf{N}} \sum_{j=1}^{\mathsf{M}_p} \left[\left(\mathsf{y}_p^i = j \right) F_p \left(\mathsf{x}^i, \Delta \right) \right] \tag{7}$$

where N, M_p, x^i, and y_p^i represent the number of training samples, number of classes, the i^{th} training example, and its actual class, respectively. Δ denotes the model parameters, and $F_p\left(\mathsf{x}^i, \Delta\right)$ represents the predicted popularity class.

3.2.5 Rumor Detection Model. This model takes two inputs: the feature vector FV generated by the convolutional neural network model and the feature vector FV_p generated by the popularity prediction model. The intuition behind using the output of the popularity prediction model as an input for the breaking news rumor detection model is to leverage from the popularity prediction model, at the example level, in improving the rumor detection model. The *context-sensitive MTL* model proposed in [Silver et al. 2008] for lifelong learning systems has showed that, by experimenting on five real-life domains, using additional contextual inputs improves the predictive performance of the primary task of the model. Inspired by their work, to further improve the breaking news rumor detection model, we use the popularity prediction-specific context as an additional input to the rumor detection model to achieve knowledge sharing at the domain as well as the example level.

[2]Source: https://blog.scrunch.com/what-is-a-good-engagement-rate-on-twitter, retrieved on April 20, 2019

Table 1. Percentages of rumors and non-rumors tweets in the PHEME datasets

Breaking News	Rumors	Non-rumors
Charlie Hebdo	458 (22.0%)	1,621 (78.0%)
Ferguson	284 (24.8%)	859 (75.2%)
Germanwings Crash	238 (50.7%)	231 (49.3%)
Ottawa Shooting	470 (52.8%)	420 (47.2%)
Sydney Siege	522 (42.8%)	699 (57.2%)

This model learns the rumor detection task as follows. First, the feature vector FV is fed through fully connected layers and the output is concatenated with the FV_p feature vector. The merged vector is then fed through a softmax layer with *tanh* function to predict whether or not the input micro-post is a rumor. The objective function here is to minimize the cross-entropy rumor detection loss function L_R as follows:

$$L_R = -\frac{1}{\mathsf{N}} \sum_{i=1}^{\mathsf{N}} \sum_{j=1}^{\mathsf{M}_R} \left[\left(\mathsf{y}_R^i = j \right) F_R \left(\mathsf{x}^i, \Omega \right) \right] \tag{8}$$

where N, M_R, x^i, and y_R^i represent the number of training samples, number of classes, the i^{th} training example, and its actual class, respectively. Ω denotes the model parameters and $F_R\left(\mathsf{x}^i, \Omega\right)$ represents the predicted rumor class. Finally, the joint loss for the entire joint learning model is calculated as the following unified loss:

$$L_{uni} = \lambda L_P + L_R \tag{9}$$

where λ is a weighting factor to be learned.

By designing our model this way, we leverage from the shared characteristics between popular micro-posts in social media and rumors by using the feature vector FV_p generated by the popularity prediction model to help improve the task of breaking news rumors detection in social media.

4 EXPERIMENTS

The objectives of the experiments are to evaluate the classification performance of the proposed joint learning model and to evaluate the effect of joint learning on the single task of breaking news rumors detection and the task of breaking news rumors popularity prediction.

4.1 Dataset

In our experiments, we used five sets of real-life, publicly accessible tweets from *PHEME* [Zubiaga et al. 2016a], where each set is related to a different breaking news story and contains both rumors and non-rumors tweets. Table 1 shows the statistics of the PHEME dataset.

To evaluate the effect of using different features on the classification performance of our proposed model as well as the baseline

Table 2. The list of primal emotions and the associated emotional triggers emoticons used in social media

Emotion	Emotional triggers emoticons
Anger	>:S , >:{ , >: , >:[, >:| , x-@ , :@ , :-@ , :-/ , :-\ , :/ , :\
Disgust	:&, , :-&
Fear	:-o , :-O , :o , :O , :-$, :$, o_O , O_o
Joy	:-) , :) , ;-) , ;) , :-)) , :-D , :D , ;-D , ;D , :-p , :p , ;-p :^) , ;^) , :o) , ;o) , :') , :-] , :] , ;-] , ;] , :-> , :> , :~) ;p , =-D , =D
Sadness	:-(, :(, :-((, =(, :-[, :[, :-< , :< , :~(, :^(, :o(, :'(
Surprise	:-o , :-O , :o , :O , :-$, :$, o_O , O_o

Table 3. Mean ± variance of precision (P), recall (R), and F1 scores of detecting high-engaging breaking news rumors using different features sets and variations of our proposed joint learning model

Model	Features	P	R	F1
Multi-task CNN-based model	Txt	0.711 ±0.002	0.775 ±0.007	0.742 ±0.005
	TS	0.712 ±0.002	0.784 ±0.001	0.746 ±0.001
	TE	0.725 ±0.000	0.772 ±0.002	0.748 ±0.001
	TSE	0.710 ±0.002	0.790 ±0.010	0.748 ±0.006
Multi-task CNN-Attn-based model	Txt	0.712 ±0.001	0.794 ±0.006	0.753 ±0.002
	TS	0.716 ±0.001	**0.797** ±0.001	0.756 ±0.001
	TE	**0.731** ±0.000	0.791 ±0.001	**0.761** ±0.000
	TSE	0.721 ±0.001	0.792 ±0.001	0.755 ±0.001

models, we experimented with the following feature sets as our input:

- **Text of micro-posts (Txt).**
- **Text and stylometric features (TS).** The stylometric features used in this work are *capital letters ratio* and the *counts* of each (**?** , **!**, **.**, ", ', **words**, **URL**, #, **@**, **emoicons**, **,**, **&**, **;**, **:**) in the text of the micro-post.
- **Text and emotional triggers features (TE).** For the emotional trigger words, we employed the *NRC Word-Emotion Association Lexicon (EmoLex)* [Mohammad and Turney 2010]. This lexicon covers words that are associated with the eight primal emotions: *anger, fear, disgust, sadness, anticipation, surprise, joy, and trust.* Furthermore, we leverage from the *emoticons* associated with these primal emotions in social media. Table 2 lists the emotional triggers emoticons we adopt in this work.
- **Text, stylometric, and emotional triggers features (TSE).**

4.2 Experimental Settings

To simulate a real-life breaking news scenario, we performed a 5-fold cross-validation as follows. In each run, we used the datasets of four breaking news stories as our training data. Then, we used the fifth dataset to evaluate the classification performance of the proposed model and the baseline models in terms of precision, recall, and F1. By designing our experiment this way, we insure that the dataset used for the evaluation in each of the five runs represents breaking news rumors of unseen topics. Furthermore, to insure the stability of the reported results of deep learning models and get a more robust estimation of their classification performances, we did five repetitions of each run of the 5-fold cross-validation for each model. Then, we reported their classification performance as the *mean* ± *variance* of the precision, recall, and F1 scores across the five repetitions of the 5-fold cross-validation instead of a single 5-fold cross-validation run. To capture the imbalance in the data sets, the precision, recall, and F1 are calculated as the micro-averaged-value across all classes. Random search was used for tuning the hyper-parameters of the deep learning models.

4.3 Experimental Results

4.3.1 Classification Performance of the Proposed Joint Learning Model. This is the first work that tackles the problem of detecting high-engaging breaking news rumors in social media. To evaluate the classification performance of our proposed joint learning model, we compared the following variations of the proposed model:

- **Multi-task CNN-based model.** A multi-task joint learning model with a CNN as the feature extractor.
- **Multi-task CNN-Attn-based model.** A multi-task joint learning model with both self-attention and CNN models as feature extractors.

In this experiment, we trained the model to simultaneously learn the two tasks of breaking news rumor detection and breaking news popularity prediction. We used different feature sets as inputs and reported the classification performance results of each model as the *mean* ± *variance* of precision, recall, and F1. Table 3 shows the obtained results. Bold values indicate the best classification performance among all models. As shown in the table, the *Multi-task CNN-Attn-based* model along with the *Text and Emotional* features outperformed all other models in detecting high-engaging breaking news rumors in terms of precision and F1, while it outperformed all other models in terms of recall when the *Text and Stylometric* features are used as its input. The results also show that, for each features set, our proposed *Multi-task CNN-Attn-based* model yielded a better classification performance than the *Multi-task CNN-based* model. This suggests that using both the convolutional neural network model and the self-attention model as shared feature extractors helps the proposed model to better capture the salient features and the semantic similarities among important words and phrases in the input text for the task of detecting high-engaging breaking news rumors in social media. In our design, the convolutional neural network model is used to learn several classes of semantic features in the input text of a micro-post, while the self-attention layer is used to emphasize important words that contribute more to the meaning of the micro-post.

Table 4. Mean ± variance of precision (P), recall (R), and F1 scores of the two tasks of breaking news rumors detection and popularity prediction across all five runs for the single-task baseline classifiers and our proposed joint learning models using different features sets

Model	**Features**	**Breaking News Rumors Detection**			**Breaking News Rumors Popularity Prediction**		
		P	**R**	**F1**	**P**	**R**	**F1**
NB	Txt	0.512	0.523	0.517	0.222	0.335	0.267
	TS	0.536	0.527	0.531	0.315	0.347	0.330
	TE	0.530	0.529	0.529	0.299	0.300	0.300
	TSE	0.546	0.536	0.541	0.319	0.360	0.338
SVM	Txt	0.307	0.500	0.380	0.224	0.333	0.268
	TS	0.494	0.533	0.513	0.347	0.346	0.347
	TE	0.437	0.534	0.481	0.257	0.334	0.291
	TSE	0.502	0.534	0.518	0.361	0.357	0.359
RF	Txt	0.320	0.500	0.390	0.224	0.333	0.268
	TS	0.321	0.499	0.391	0.234	0.338	0.277
	TE	0.340	0.398	0.367	0.257	0.321	0.285
	TSE	0.332	0.504	0.400	0.307	0.379	0.339
CNN-based model	Txt	0.514 ±0.000	0.554 ±0.016	0.533 ±0.009	0.811 ±0.000002	0.959 ±0.0003	0.879 ±0.00004
	TS	0.519 ±0.000	0.559 ±0.004	0.538 ±0.001	0.810 ±0.00001	0.988 ±0.00007	0.890 ±0.000003
	TE	0.532 ±0.000	0.564 ±0.004	0.548 ±0.0009	0.834 ±0.00004	0.886 ±0.003	0.859 ±0.0006
	TSE	0.551 ±0.000	0.599 ±0.001	0.574 ±0.0002	**0.841** ±0.00003	0.943 ±0.0002	0.889 ±0.00002
CNN-Attn-based model	Txt	0.514 ±0.000	0.618 ±0.052	0.561 ±0.045	0.805 ±0.00001	0.963 ±0.007	0.878 ±0.001
	TS	0.523 ±0.00001	0.635 ±0.001	0.574 ±0.0002	0.805 ±0.000	0.999 ±0.000	0.890 ±0.000
	TE	0.519 ±0.00002	0.666 ±0.011	0.583 ±0.0009	0.794 ±0.0007	0.999 ±0.000003	0.885 ±0.0003
	TSE	0.518 ±0.00005	**0.728** ±0.029	0.605 ±0.003	0.806 ±0.000	**1.000** ±0.000	**0.892** ±0.000
Multi-task CNN-based model	Txt	0.616 ±0.003	0.576 ±0.014	0.595 ±0.003	0.811 ±0.00003	0.974 ±0.001	0.885 ±0.0001
	TS	0.613 ±0.003	0.589 ±0.002	0.601 ±0.002	0.811 ±0.00003	0.978 ±0.0005	0.887 ±0.00003
	TE	0.638 ±0.001	0.565 ±0.003	0.599 ±0.0004	0.812 ±0.00002	0.979 ±0.0005	0.888 ±0.0001
	TSE	0.605 ±0.004	0.618 ±0.020	0.612 ±0.002	0.815 ±0.00003	0.973 ±0.0003	0.887 ±0.00003
Multi-task CNN-Attn-based model	Txt	0.613 ±0.003	0.609 ±0.011	0.611 ±0.0003	0.811 ±0.0002	0.979 ±0.0002	0.887 ±0.00004
	TS	0.621 ±0.002	0.613 ±0.002	0.617 ±0.0004	0.811 ±0.00003	0.981 ±0.001	0.888 ±0.00004
	TE	**0.647** ±0.0002	0.600 ±0.001	**0.623** ±0.001	0.814 ±0.00001	0.982 ±0.001	0.890 ±0.0001
	TSE	0.630 ±0.002	0.600 ±0.001	0.615 ±0.0001	0.813 ±0.0005	0.978 ±0.0005	0.888 ±0.0003

We further evaluated the classification performance of the two variations of the proposed joint learning model: the *Multi-task CNN-based* model and the *Multi-task CNN-Attn-based* model using the *Area Under Curve - Receiver Operating Characteristic (AUC-ROC)* curves. AUC-ROC curves are used to evaluate the classification performance of classification models at various threshold settings. The *Receiver Operating Characteristic (ROC)* curve is plotted in a 2-dimensional space where the x-axis represents the *False Positive Rate*

(FPR) and the y-axis represents the *True Positive Rate (TPR)*. The *Area Under Curve (AUC)* score measures the ability of the classification model to separate or distinguish between the different classes. The higher the AUC score, the better the classification performance of the model is.

In this experiment, we performed a 5-fold cross-validation and plotted the ROC curve for each run as well as the mean ROC curve across all five runs for the two models. We also calculated the AUC score for each run as well as the Mean ± variance of the AUC scores across all five runs for each model. Figure 2 shows the obtained results. As shown in Figure 2.a, the mean AUC score is 0.63 ± 0.14 for the *Multi-task CNN-based* model while the mean AUC score is 0.69 ± 0.08 for the *Multi-task CNN-Attn-based* model as shown in Figure 2.b. These results show that our proposed *Multi-task CNN-Attn-based* model significantly outperformed the baseline *Multi-task CNN-based* model in detecting high-engaging breaking news rumors in terms of the AUC-ROC score. This suggests that the design of our proposed joint learning model helps it better capture the features for the task of detecting high-engaging breaking news rumors in social media.

4.3.2 Joint Learning Effect on the Single Tasks of Breaking News Rumors Detection and Breaking News Rumors Popularity Prediction. The basic assumption is that, when jointly learning multiple tasks, the shared knowledge among them will lead to better predictive performance compared to independently learning each task [Thung and Wee 2018]. To evaluate the effect of the joint learning on the classification performance of the single tasks of breaking news rumors detection as well as breaking news rumors popularity prediction, we compared our joint learning model with the following single-task baseline classifiers:

- **Naive Bayes (NB), Support Vector Machine (SVM), and Random Forest (RF) classifiers.** We implemented a NB classifier, a SVM classifier, and a RF classifier using *R* [R Core Team 2013] and *RStudio* [3]. A pretrained word2vec model was used to replace the text of each micro-post with its *word embeddings* representation.
- **CNN-based model.** A single-task model with a CNN model as the feature extractor.
- **CNN-Attn-based model.** A single-task model with self-attention and CNN models as feature extractors.

In this experiment, we used different features sets as our inputs. For each feature set, we trained two models of each baseline single-task classifier. One to learn the task of breaking news rumors detection and another to learn the task of breaking news rumors popularity prediction. Table 4 shows the obtained results reported as the precision, recall, and F1 scores of the NB, SVM, and RF classifiers, and as the *mean ± variance* of precision, recall, and F1 of each deep learning model. Bold values indicate the best classification performance among all models. The following subsections discuss the effect of the joint learning on the classification performance of each task.

Joint learning effect on breaking news rumors detection task: Table 4 shows that *Multi-task CNN-Attn-based* model along with *Text and Emotional* features had the best classification performance in terms of precision and F1, while *CNN-Attn-based* model along with *Text, Stylometric, and Emotional* features yielded the best classification performance in terms of recall. These results suggest that jointly learning the two tasks of breaking news rumors detection and breaking news rumors popularity prediction by our proposed model had significantly improved the classification performance of breaking news rumors detection in terms of precision and F1

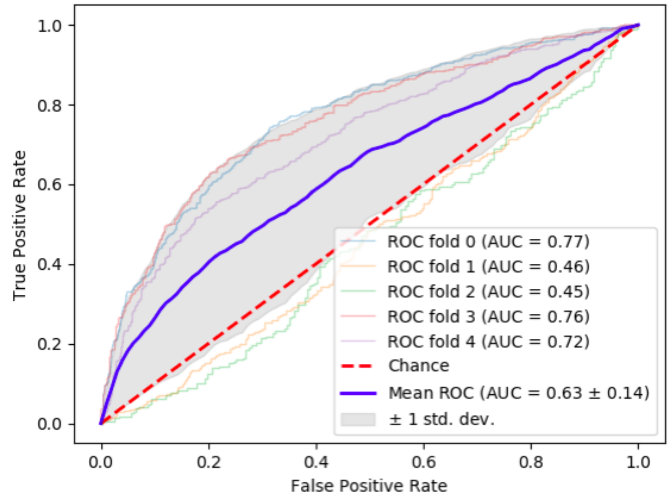

Fig. 2.a: AUC-ROC for the Multi-task CNN-based model

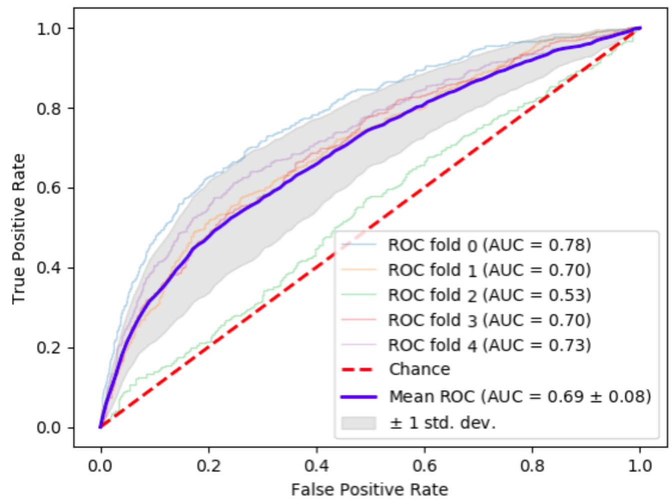

Fig. 2.b: AUC-ROC for the Multi-task CNN-Attn-based model

Fig. 2. Receiver Operating Characteristic (ROC) curves for the two variations of the proposed joint learning model showing the ROC curves and the Area Under Curve (AUC) scores for each of the five runs and the Mean ± variance of the AUC scores across all five runs. Figure 2.a shows the obtained results for the Multi-task CNN-based model and Figure 2.b shows the obtained results for the Multi-task CNN-Attn-based model

[3]Source: https://www.rstudio.com/, retrieved on January 12, 2019

over all single-task classifiers in our experiment. This shows how the design of our proposed model helps leverage from the shared characteristics between the two tasks in improving the breaking news rumors detection task.

Joint learning effect on breaking news rumors popularity prediction task: Table 4 shows that *CNN-based* model along with *Text, Stylometric, and Emotional* features had the best performance in terms of precision, while *CNN-Attn-based* model along with *Text, Stylometric, and Emotional* features yielded the best classification performance in terms of recall and F1. These results suggest that the joint learning did not improve the task of breaking news rumors popularity prediction. In fact, single-task deep learning models yielded the best overall classification performance. Nevertheless, our proposed joint learning model still achieved a high classification performance in terms of precision and recall and a comparable classification performance in terms of F1 to all single-task deep learning classifiers in our experiment. This shows that our joint learning model is capable of learning important latent features for predicting the popularity of breaking news rumors in social media with high accuracy without the need for gathering the early observations of its dynamics or the need for a collection of related or similar topics and posts.

4.3.3 Discussion on Feature Sets.

We discuss the effect of including the *emotional triggers* and the *stylometric* features, in addition to the text, on the classification performance of breaking news rumors detection and breaking news popularity prediction. We start by inspecting the results in Table 4 to determine, for each model, which feature sets yielded the best classification performance in terms of precision, recall, and F1. We observed that using the *Text, Stylometric, and Emotional* features sets yielded the best classification performance in terms of precision of the rumors detection task and in terms of recall and F1 of the rumors popularity prediction task in all models except the *Multi-task CNN-Attn-based model*. It also yielded the best classification performance in terms of precision of the rumors detection task for three out of the five single-task models and the best classification performance in terms of recall and F1 of the rumors popularity prediction task for four out of the five single-task models. We also observed that, for the proposed joint learning model, using the *Text and Emotional* features sets yielded the best classification performance in terms of precision and F1 for rumors detection and in terms of precision, recall, and F1 for popularity prediction.

These observations suggest the following. First, using *emotional triggers* as well as *stylometric* features can effectively help a classification model better learn the tasks of breaking news rumors detection and the breaking news rumors popularity prediction in social media. Second, in most of the cases, the best classification performances of the two tasks were achieved using the same sets of features. Hence, the high accuracy of jointly learning the two tasks. Finally, although incrementally learning the word embeddings of the input text helps mitigate the topic-shift and OOV issues of breaking news rumors detection, including emotional triggers and stylometric features helps improve the task of detecting high-engaging breaking news rumors in social media.

5 CONCLUSION

With the increased adaptation of social media as a major source of gathering worldwide breaking news, detecting rumors and acting upon them in a timely fashion becomes an extremely challenging task. Breaking news rumors, if not identified as early as possible, may have extremely damaging consequences. Yet, not all rumors will spread in social media. detecting high-engaging rumors helps authorities prioritize the rumor verification process during breaking news. In this work, we tackle the problem of detecting high-engaging breaking news rumors in social media by proposing a model that jointly learns the tasks of breaking news rumor detection and breaking news rumors popularity prediction. Our experiments on real-life datasets show that our joint learning model outperforms the baseline classifiers and is capable of detecting high-engaging breaking news rumors with high accuracy.

ACKNOWLEDGMENTS

The first author is supported by King Saud University in Riyadh, Saudi Arabia, and the Saudi Arabian Cultural Bureau in Canada. The research was conducted during her doctoral study at Concordia University in Canada. The second author is supported by the Discovery Grants (RGPIN-2018-03872) from the Natural Sciences and Engineering Research Council of Canada (NSERC) and Canada Research Chairs Program (950-232791).

REFERENCES

S. Abbar, C. Castillo, and A. Sanfilippo. 2018. To Post or Not to Post: Using Online Trends to Predict Popularity of Offline Content. In *Proc. of the 29th on Hypertext and Social Media (HT '18)*. ACM.

S. A. Alkhodair, S. H. H. Ding, B. C. M. Fung, and J. Liu. 2019. Detecting Breaking News Rumors of Emerging Topics in Social Media. *Information Processing and Management Journal - A Special Issue on Mining Social Influence and Actionable Insights from Social Networks* (2019).

G. W. Allport and L. J. Postman. 1965. *The psychology of rumor.* Russell & Russell.

R. Bandari, S. Asur, and B. A. Huberman. 2012. The Pulse of News in Social Media: Forecasting Popularity. *CoRR* (2012).

J. A. Berger and K. Milkman. 2009. What Makes Online Content Viral? *Journal of Marketing Research* 49 (2009).

F. Chen and W. H. Tan. 2018. Marked Self-exciting Point Process Modelling of Information Diffusion on Twitter.

W. Comcowich. 2015. *Social Media Measurement: How to Determine Your Engagement Rate.* https://glean.info/social-media-measurement-how-to-determine-your-engagement-rate/

M. Crescimbene, F. La Longa, and T. Lanza. 2012. The science of rumors. *Annals of geophysics* 55 (2012).

N. Difonzo and P. Bordia. 2007. Rumors influence: Toward a dynamic social impact theory of rumor. *Science of social influence: Advances and future progress* (2007).

Y. Hu, Y. Li, T. Yang, and Q. Pan. 2018. Short Text Classification With A Convolutional Neural Networks Based Method.

K. E. Matsa and E. Shearer. 2018. *News Use Across Social Media Platforms 2018.* Report. Pew Research Center.

T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proc. of the 26th International Conference on Neural Information Processing Systems*. Curran Associates Inc., USA.

S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao. 2020. Deep Learning Based Text Classification: A Comprehensive Review.

S. Mishra, M.-A. Rizoiu, and L. Xie. 2016. Feature Driven and Point Process Approaches for Popularity Prediction. In *Proc. of the 25th ACM International on Conference on Information and Knowledge Management (CIKM '16)*. ACM.

A. Mitchell, K. Simmons, K. E. Matsa, and L. Silver. 2018. *People in poorer countries just as likely to use social media for news as those in wealthier countries.* Report. Pew Research Center.

S. M. Mohammad and P. D. Turney. 2010. Emotions Evoked by Common Words and Phrases: Using Mechanical Turk to Create an Emotion Lexicon. In *Proc. of the*

NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text (CAAGET '10). Association for Computational Linguistics.

E. S. Muhammad. 2017. The Psychology of Rumors: Why They Spread & How To Avoid Being Misled.

P. Ozturk, H. Li, and Y. Sakamoto. 2015. Combating Rumor Spread on Social Media: The Effectiveness of Refutation and Warning. In *Proc. of the 48th Hawaii International Conference on System Sciences (HICSS '15)*.

R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. http://www.R-project.org/

D. Silver, R. Poirier, and D. Currie. 2008. Inductive transfer with context-sensitive neural networks. *Machine Learning* 73 (12 2008), 313–336. https://doi.org/10.1007/s10994-008-5088-0

Socialbakers. 2013. *Engagement Rate: A Metric You Can Count On*. https://www.socialbakers.com/blog/1427-engagement-rate-a-metric-you-can-count-on

M. Takayasu, K. Sato, Y. Sano, K. Yamada, W. Miura, and H. Takayasu. 2015. Rumor Diffusion and Convergence during the 3.11 Earthquake: A Twitter Case Study. *PLoS ONE* 10 (2015).

K. Thung and C.-Y. Wee. 2018. A brief review on multi-task learning. *Multimedia Tools and Applications* 77 (08 2018). https://doi.org/10.1007/s11042-018-6463-x

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30*.

Y. Yan, Z. Tan, X. Gao, S. Tang, and G. Chen. 2016. STH-Bass: A Spatial-Temporal Heterogeneous Bass Model to Predict Single-Tweet Popularity. In *Database Systems for Advanced Applications*, S. B. Navathe, W. Wu, S. Shekhar, X. Du, S. X. Wang, and H. Xiong (Eds.). Springer International Publishing.

W. Yin, K. Kann, M. Yu, and H. Schütze. 2017. Comparative Study of CNN and RNN for Natural Language Processing. (02 2017).

T. Zaman, E. B. Fox, and E. T. Bradlow. 2014. A Bayesian approach for predicting the popularity of tweets. *The Annals of Applied Statistics* (2014).

Y. Zhang and B. Wallace. 2017. A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification. In *Proc. of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Asian Federation of Natural Language Processing.

Q. Zhao, M. A. Erdogdu, H. Y. He, A. Rajaraman, and J. Leskovec. 2015a. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In *Proc. of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. ACM.

Z. Zhao, P. Resnick, and Q. Mei. 2015b. Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts. In *Proc. of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland.

P. Zhou, S. Zheng, J. Xu, Z. Qi, H. Bao, and B. Xu. 2017. Joint Extraction of Multiple Relations and Entities by Using a Hybrid Neural Network. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, M. Sun, X. Wang, B. Chang, and D. Xiong (Eds.). Springer International Publishing, 135–146.

A. Zubiaga, A. Aker, K. Bontcheva, M. Liakata, and R. Procter. 2018. Detection and Resolution of Rumours in Social Media: A Survey. *Comput. Surveys* 51 (2018).

A. Zubiaga, G. W. S. Hoi, M. Liakata, and R. Procter. 2016a. PHEME dataset of rumours and non rumours.

A. Zubiaga, M. Liakata, and R. Procter. 2016b. Learning Reporting Dynamics during Breaking News for Rumour Detection in Social Media. *CoRR* (2016).