

**Please cite the Published Version**

Jain, Deepak Kumar, Kumar, Akshi  and Shrivastava, Akshat (2022) CanarDeep: a hybrid deep neural model with mixed fusion for rumour detection in social data streams. *Neural Computing and Applications*, 34 (18). pp. 15129-15140. ISSN 0941-0643

**DOI:** <https://doi.org/10.1007/s00521-021-06743-8>

**Publisher:** Springer (part of Springer Nature)

**Version:** Accepted Version

**Downloaded from:** <https://e-space.mmu.ac.uk/629487/>

**Additional Information:** This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <https://doi.org/10.1007/s00521-021-06743-8>

**Enquiries:**

If you have questions about this document, contact [openresearch@mmu.ac.uk](mailto:openresearch@mmu.ac.uk). Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

# ***CanarDeep*: A Hybrid Deep Neural Model with Mixed Fusion for Rumour Detection in Social Data Streams**

Deepak Kumar Jain<sup>1</sup>, Akshi Kumar<sup>2\*</sup>, Akshat Shrivastava<sup>3</sup>

<sup>1</sup>Key Laboratory of Intelligent Air-Ground Cooperative Control for Universities in Chongqing, College of Automation, Chongqing University of Posts and Telecommunications, Chongqing, China

<sup>2</sup>Department of Information Technology, Netaji Subhas University of Technology, Delhi, India

<sup>3</sup>Department of Computer Science & Engineering, Delhi Technological University, Delhi, India

\**akshi.kumar@nsut.ac.in*

**Abstract.** The unrelenting trend of doctored narratives, content spamming, fake news and rumour dissemination on social media can lead to grave consequences that range from online intimidating and trolling to lynching and riots in real-life. It has therefore become vital to use computational techniques that can detect rumours, do fact-checking and inhibit its amplification. In this paper, we put forward a model for rumour detection in streaming data on social platforms. The proposed *CanarDeep* model is a hybrid deep neural model that combines the predictions of a hierarchical attention network (HAN) and a multi-layer perceptron (MLP) learned using context-based (text + meta-features) and user-based features respectively. The concatenated context feature vector is generated using feature-level fusion strategy to train HAN. Eventually, a decision-level late fusion strategy using logical OR combines the individual classifier prediction and outputs the final label as rumour or non-rumour. The results demonstrate improved performance to the existing state-of-the-art approach on the benchmark PHEME dataset with a 4.45% gain in F1-score. The model can facilitate well-time intervention and curtail the risk of widespread rumours in streaming social media by raising an alert to the moderators.

**Keywords:** Rumour; Social data streams; Deep learning; Data fusion

## **Declarations**

**Funding:** No Funding has been received.

**Conflicts of interest/Competing interests:** The authors certify that there is no conflict of interest in the subject matter discussed in this manuscript.

**Availability of data and material:** Publicly accessible data has been used by the authors.

**Code availability:** Can be made available on request.

**Authors' contributions:** All the authors have equally contributed in the manuscript preparation.

**Ethics approval:** The work conducted is not plagiarized. No one has been harmed in this work.

**Consent to participate:** All the authors have given consent to submit the manuscript.

**Consent for publication:** Authors provide their consent for the publication.

## **1. Introduction**

The blatant social media technology has an obvious, potentially enormous dark side. The *tsunami* of polluted information threatens the underpinning trust between users and their experience online. The amplification of the most abominable views in social data streams, the virality of incorrect information at a lightning speed and the anonymization of our public discourse divulge the vulnerabilities associated. Moreover, the information disorders (misinformation, disinformation and mal-information) [1] vary in impact and mechanization which is not limited to a single social media platform. The primary function of the misleading and fabricated content is to make sense in a wrapped version of reality with intentions ranging from hurtful or harmless to the harmful one. The economics of social media favors these 'information statements' that are instrumentally relevant and often consumed within our online echo chambers and filter bubbles. Undoubtedly, doctored narratives, fake news, rumours, polarized content, deceptive propaganda, content spamming and fabricated hoaxes define the taxonomy of online content within the periphery of social media abuse [2-4].

A rumour is any statement that is not yet confirmed at the time of posting, irrespective of whether it's true or false [5]. It gets viral at great speed and is facilely believed particularly during public crisis. As soon as users start engaging with the post by "liking", "replying", "commenting", "forwarding", and "sharing",

technology proliferates diffusion at an unparalleled rate. The COVID-19 public health crisis made this vulnerability highly visible, with numerous rumours circulating through social media. To examine post legitimacy, many social media platforms have put in systematic efforts to moderate posts and improve online accountability. These include implementing the mandatory code of practice for content filtering using a combination of artificial intelligence and user feedback. But the guidelines are not transparent to the users and the content moderators have to deal with information overload and exposure to harmful and undesirable content. Indeed the information credibility analysis on social media platforms is complex as it diversifies along three dimensions as post-level credibility, user-level credibility and topic-level credibility [6, 7]. Moreover, the problem with debunking rumours is that the exchange and distribution is at such an unprecedented rate that makes it difficult to scale up the combat operations to curb the viral spread. At the same time, the diffusion process of rumors often has an associated chronology. Often spread as breaking news with high virality and cascading effect, the rumours subside with time and eventually die out. It is reasonably practical to capture the changing characteristics of rumour spread as with time veracity and stance are comprehensible. But the damage done within the early hours of release is huge as it evokes high-arousal emotions. Thus, early detection of rumours (identifying if an online post is a rumor or non-rumor) and combating its viral spread is a pressing priority. A rumour detection system that identifies an online posting as a candidate rumour in its early stages can be effectively used for well-timed intervention by alerting the moderators to take corrective measures that impede the spreading of inappropriate content.

Pertinent studies report an enormous use of machine learning techniques to combat online content fabrication but have limited success due to lack of context-awareness. Recent research trends report the use of deep learning models such as Convolutional neural networks (CNNs) [8] and recurrent neural networks (RNNs) [9] for various natural language understanding tasks on streaming data of online media [10]. A superior hierarchical attention based document classification model, called hierarchical attention network (HAN) has been put forward by Yang et al. Commonly built from bidirectional RNNs [11], a HAN consists of gated recurrent units (GRUs) or Long-short-term-memory model (LSTMs) with attention mechanisms. It is conceptual based on “hierarchies” where the output from lower levels in hierarchical structure become the input to upper level. The HAN architecture characterizes the hierarchically derived knowledge in a document with sentence and words as building blocks at the corresponding level of hierarchy. Falsehood on Twitter spreads like a wildfire and may be compiled as a single tweet, a tweet thread (multi-part tweet) or a conversational thread. Using HAN lets discrete contribution of compound fragments of a tweet (tweet-sentence-word) to its quintessence by implementing two attention mechanisms. The contextual connotation of these fragments is taken into account to construct the document representation. This research aims to cope with the language complexities by using HAN with deep contextualized language representations from pre-trained ELMo (Embeddings from Language Models) word embedding [12].

As an effectual technique to tweet-level analysis of rumours, an attention-based hybrid deep learning model for rumour detection from social data streams is put forward. The proposed *CanarDeep* model detects rumourous posts in real-time by using the learned features from both HAN and multi-layer perceptron (MLP). The model derives its nomenclature from the French word “*Canard*” which means ‘*unfounded, groundless or false report or story*’ and “*Deep*” from the ‘*deep neural network*’ built to detect rumours and combat its viral spread in online streams. A primary approach to decipher the truth value of a post is to look for some user-based and text-based evidence. Some meta-features such as re-post count, and morpho-syntactic (exclamations) & typographic (capitalization, quotes) markers can also serve as non-trivial cues. Therefore, in the mixed-fusion hybrid deep neural model, rumour detection is done by the individual classification model, namely HAN and MLP using context-based features (textual + meta-features) and user-based features respectively. The context-based features are initially fused to form a concatenated feature vector. At both word and sentence levels, HAN uses a bi-directional GRU with attention [13]. The attention mechanism enables to find the most important words and sentences in a document while taking the context into consideration. On the other hand, the MLP generates its output based on the user-based features using the back-propagation algorithm and adjusting the weights of the neurons accordingly. MLP is a simple yet highly sought-after model when it comes to binary classification using discrete numerical features [14]. Finally, a late fusion method (Logical OR operation) is used to capture the results of the two classifier sub-networks and classify the output as a decision. The primary contributions of the work are:

- The hybrid of HAN and MLP is used to classify online streaming posts as rumours & non-rumours.
- HAN is trained using context-based features (the textual content & post meta-data) and MLP is trained using the user-based features (discrete numeric features related to the user profile).
- Two types of data fusion techniques are used, that is, feature-level (early) fusion is done to concatenate context-based features whereas the output generated from individual classifiers is

combined to produce a final common decision using a logical OR for decision-level (late) fusion.

- Comparison with the state-of-the-art model [15] to validate the improvements and consistency in improvements across the events in the benchmark Twitter PHEME dataset [16].

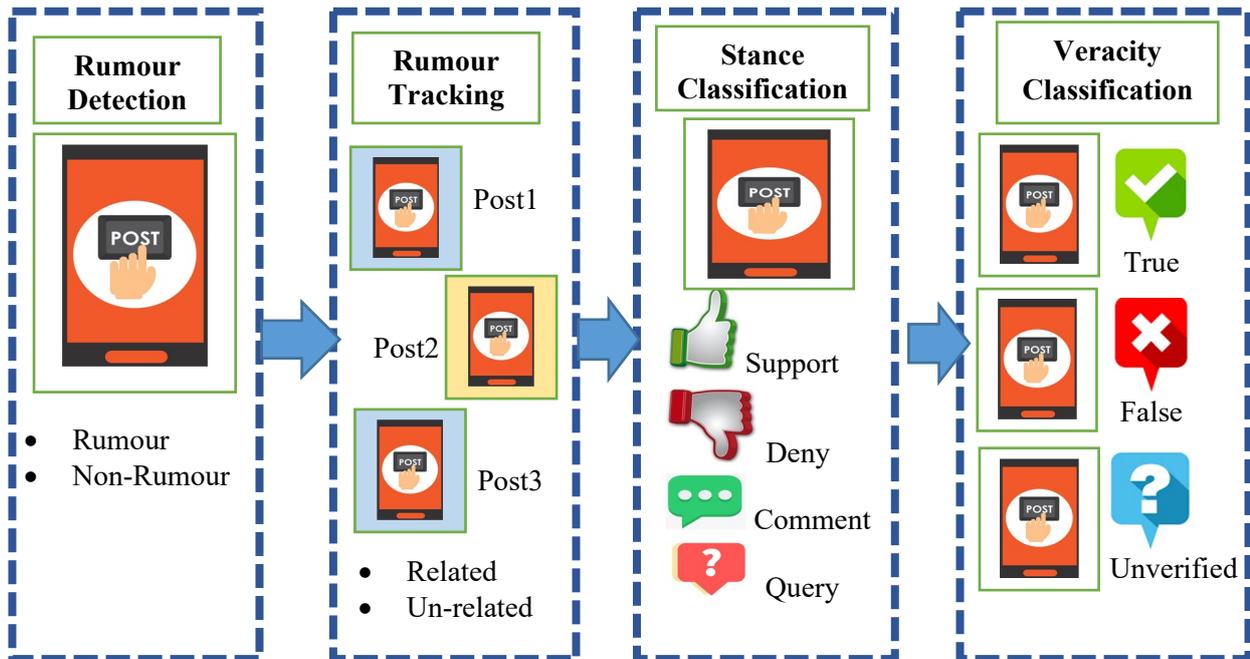
The organization of the paper is as follows: related work within the area of rumour detection in social media is given in section 2. Section 3 gives description of the mixed-fusion hybrid of HAN-MLP for rumour detection followed by the results in section 4 and conclusion in section 5.

## 2. Related Work

Various primary [5, 15] and secondary [4, 17] literature studies have been reported on rumour detection in microblogs. The variety of techniques used to classify tweets as rumour are dominated by the use of machine learning and more recently deep learning. Diverse combinations of features which include content-based, post-based, user profile based and rumour diffusion attribute have been used to train the learning models. A range of studies focus on the use of text-only based features from the tweet. Takahashi and Igata [18] used keywords and retweet ratio for rumour detection in Twitter data. Kumar et al. [19] compared machine learning and deep learning models trained using Tf-idf features and one-hot encoding respectively on the PHEME and reported Bi-LSTM as the best model. Bhattacharjee et al. [20] used the Glove word embedding to train a LSTM for rumor detection on the PHEME dataset. The use of RNN has also been reported with [21] and without [22] attention mechanism. Several studies have reported the combination of features. A hybrid of CNN and RNN was proposed by Nguyen et al. in 2017 [23] using text-based and temporal features. Guo et al. [24] suggested a word, post and sub-event level hierarchical network using attention-based Bi-LSTM encoder combined with social features. The results were evaluated on Sina Weibo and Twitter datasets. In 2020, Zubiaga et al. [16] used a conditional random field [CRF] model to detect rumours in the PHEME dataset. The authors used both content and context based features to train the model. Vijeev et al. [25] used both user features and content features to train three supervised learning methods on PHEME. Our recently published research [26] describes an attention-based Residual network (ARN) to detect rumour in the same Twitter benchmark dataset and reports an improvement in recall and F1 score in comparison to Zubaiga’s CRF model.

## 3. The Proposed *CanarDeep* Model

A comprehensive rumour classification system [3] consists of four core components, namely rumour detection, rumour tracking, stance classification and veracity classification as shown in fig.1. Automatic detection of rumours is typically a natural language processing (NLP) task to label online streaming posts as a rumour or a non-rumour. It intends to recognize potential rumours, especially the emerging ones.



**Fig. 1.** Rumour Classification system

The proposed *CanarDeep* model corresponds to the first component, i.e., the recognition of potential rumours. The publicly available PHEME dataset [15] for rumour detection which is compiled using event-

based tweets for classifying tweets into rumourous or non-rumourous categories is used. It is a collection of 5802 tweets, out of which 1972 are labelled as rumours and the remaining 3830 as non-rumours. The data is initially pre-processed to transform it into a well-formed data for analytics. Preprocessing was done by converting the tokens to lowercase, removing the hyperlinks, unwanted characters, symbols, whitespace and stop-words from the text, performing spell check, lemmatization and stemming. *StandardScaler* from the *sklearn* library for Python to scale all the user profile based features. Evaluation of a rumour strongly relies on its context which is formulated using the textual content and the post meta-data. Characteristics of user profile also help in building trust in the content posted and therefore to train our model we use the context-based, the user-based features and the rumour/non-rumour label to train our model. A total of 11 context-based features and 5 user-based features are used as shown in fig.2.

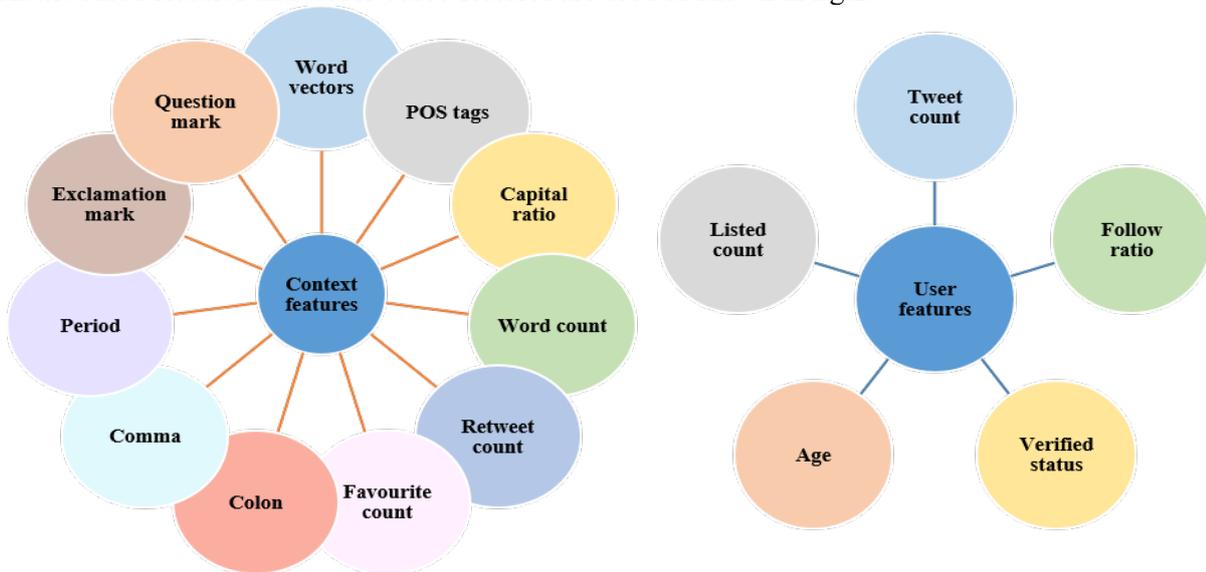


Fig. 2. Context-based and user-based features

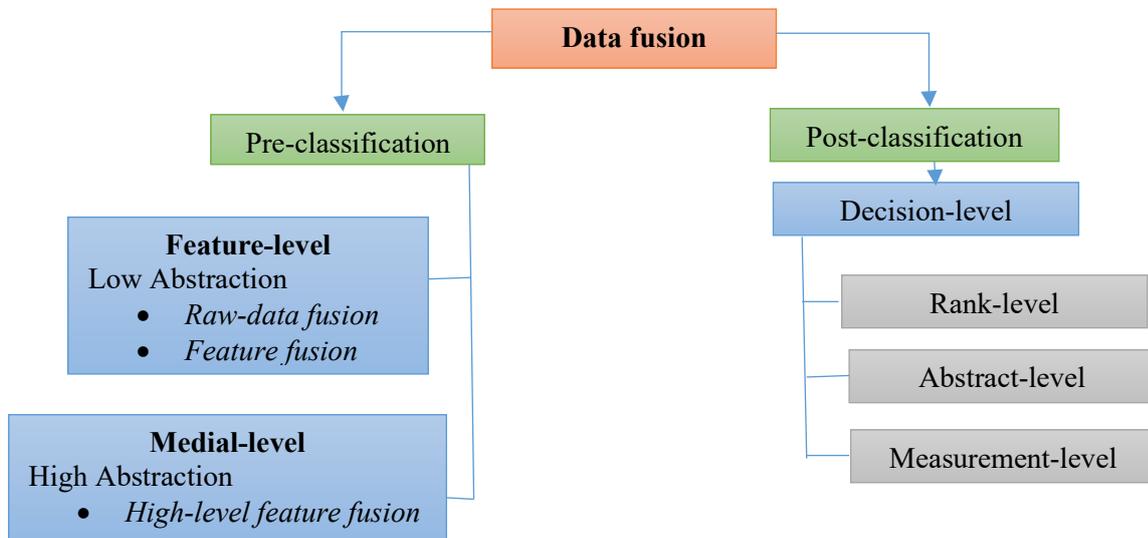
The description of these features is as follows:

- **Context-Based Features:** These entail textual content of the post as well as the content-based cues called meta-features associated with the post.
  - *Word Vectors:* Word vectors map words or phrases from vocabulary to a corresponding vector of real numbers which used to find word predictions, word similarities/semantics [27]. In *CanarDeep*, the word vector is built using the ELMo 5.5B word embedding.
  - *POS Tags:* These tags do the part-of-speech annotation. They tag each word in the tweet with the respective grammar tag such as nouns, adverbs, adjectives, etc.
  - *Capital Ratio:* Most of the time, the word which has been spelled in capital letters tends to have more impact than the word written in lower case.
  - *Word Count:* It depicts the count of words in a tweet.
  - *Question Mark:* Question marks sometimes represent the uncertainty of a saying, disrespect, impatience, or lack of tactfulness. Its presence or absence in a tweet is represented as a binary feature.
  - *Exclamation Mark:* The exclamation marks in the tweets express surprise, astonishment, or any strong emotion resulting in additional emphasis.
  - *Period:* Punctuation might represent good writing and hence quality reporting. Its presence or absence in a tweet is represented as a binary feature.
  - *Colon:* The use of a colon in tweets helps the user to add two independent clauses, thus allowing them to add two complete thoughts that stand alone as complete sentences. The presence of a colon may suggest careful reporting. Its presence or absence in a tweet is represented as a binary feature.
  - *Comma:* The use of a comma in a tweet suggests quality reporting. Its presence or absence in a tweet is represented as a binary feature.
  - *Favorite Count:* This depicts the number of users who have signaled a particular tweet as their favorite. The higher the count, the more are the chances of it not being a rumour because a higher favorite count shows that people believe in that tweet.
  - *Retweet Count:* This depicts the number of users who have retweeted a particular tweet.

Retweeting is defined as the sharing of a tweet by a user so that the user's followers can also read the tweet. If the retweet count of a tweet is high, there is a good chance that the tweet is not a rumour because the users trusted it enough to share it.

- **User-based Features:** These comprise of user profile-based discrete numeric features as follows:
  - *Tweet Count:* The total number of tweets a user had posted on Twitter.
  - *Listed Count:* The count of lists a user is a part of, i.e., the number of times they were added to a list by other users.
  - *Follow Ratio:* The reputation of a user is assessed on the basis of the count of their followers. But, sometimes, the count of followers does not reflect the true prominence of a user. For example, some users follow many others in order to be followed back. Keeping this scenario in mind, we take the follow ratio as a feature, which is the number of followers someone has divided by the number of people following them. It is basically the follower to following ratio.
  - *Age:* The age of a twitter user pertains to the years they have been using Twitter. It is the time from the setting up of the account to current tweet time.
  - *Verified:* The account verification status of the user. A verified user account has traceability and accountability making such accounts least probable to spread rumours in comparison to an unverified account.

There are two primary fusion strategies for multiple input types, namely the model-free fusion and model-level fusion (medial). Early fusion (feature-level) and late fusion (decision-level) are types of model-free fusion. Early fusion combines the input types to form a single input vector to train a classifier. It basically refers to the fusing (concatenation) of features from multiple data sources to create a new feature set. The new feature vector has higher discriminative power in comparison to the individual input feature vectors. Whereas, in late fusion individual classifiers are modelled using discrete input-types and the predictions of the individual classifiers are combined to decide the final output [28, 29]. Model-level fusion pools the benefits of both of these strategies by concatenating high-level feature representations from diverse classifiers. Fig.3 depicts these types of data fusion.



**Fig. 3.** Data fusion strategies

In this research, two types of data fusion strategies are used. The early fusion strategy is used to combine the textual content with the meta-features of the post to generate a context-based feature vector. That is, the meta-data features which are categorical in nature are concatenated to the word vector generated by the embedding layer as a feature-level fusion strategy. Late fusion is applied to combine the decisions of multiple classifiers, namely HAN and MLP, trained using context-based and user-based features respectively, to produce a final prediction. There are various abstract methods to accomplish decision-level fusion, such as using logical operators, votes or weighted majority. The proposed hybrid of HAN and MLP is trained using the user-based and context-based features with a mix of fusion strategies. Fig.4 represents the architecture of the proposed *CanarDeep* model.

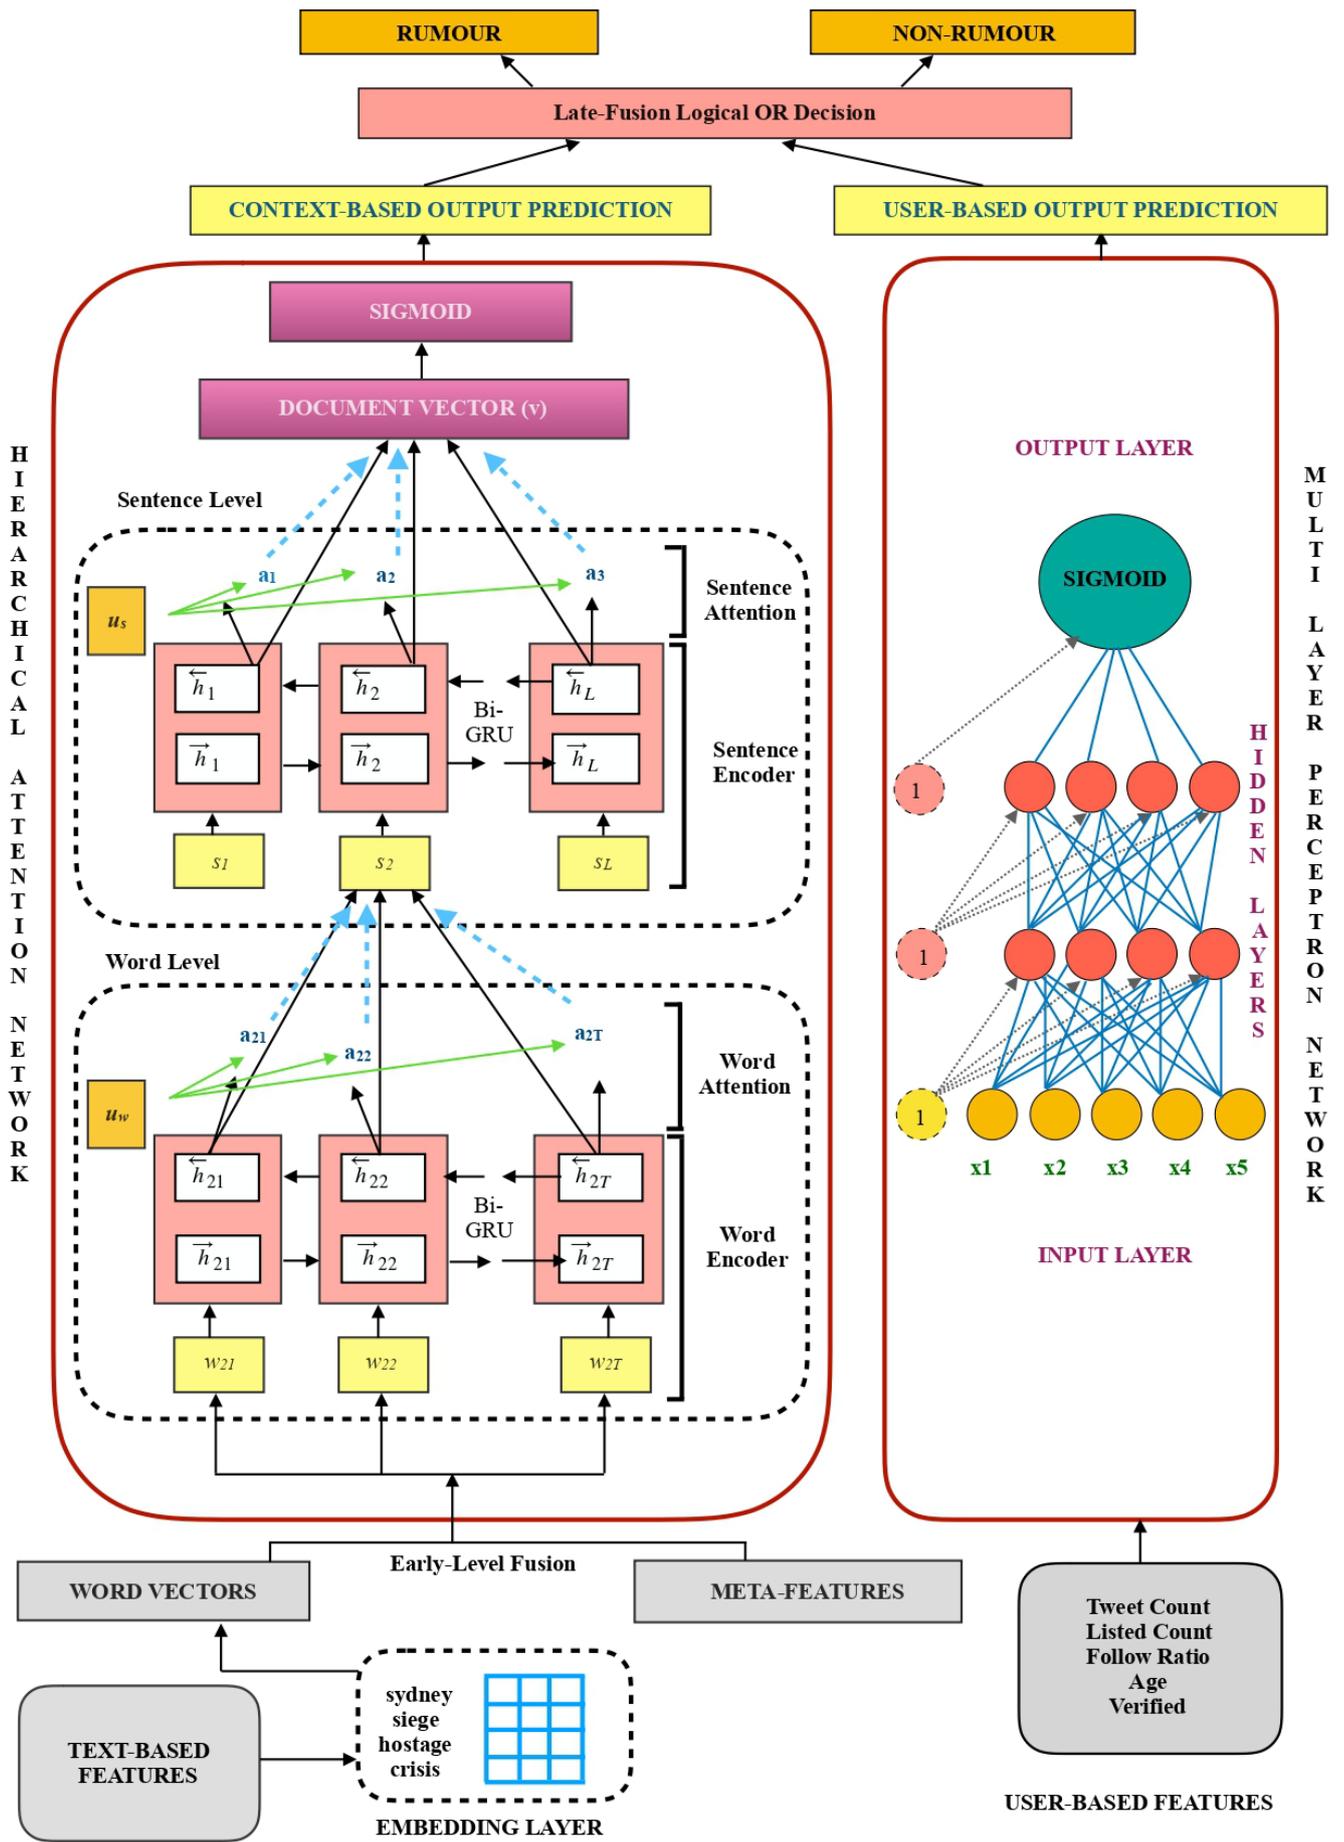


Fig.4. The Proposed CanarDeep Model

### 3.1. Hierarchical Attention Network (HAN)

To learn the context-based features, the architecture of the HAN classifier has an embedding layer, encoders and attention layers. In this research, the context-based ELMo 5.5B model [12] is used to generate the word embedding to seed the classifier. The bi-directional GRU with attention layer is then used to capture the essential meaning of the document. The model includes two levels of attention - word-level and sentence-level. The differential contribution of various parts of the document to its essential meaning and changing context is considered when constructing the representation of the document. The encoders extract relevant context and the attention layers compute the degree of relevance of the sequence of tokens with respect to the document. The bi-directional GRU with attention layer is firstly used at word-level and repeated at the sentence-level as follows:

- **Word-level encoder and attention**

An embedding matrix maps the discrete categorical variables into a vector of continuous numbers. For a given sentence with words  $w_{it}$ ,  $t \in [0, T]$ , a word to the vector embedding matrix,  $W_e, x_{it} = W_e w_{it}$  is built using ELMo. The vectorized tokens are the inputs for the next layer, the word encoding layer. In this research, a bidirectional GRU is applied to get annotations of the words by summarizing information. The forward GRU  $\vec{f}$  reads the sentence  $s_i$  from  $w_{i1}$  to  $w_{iT}$  as given in (1):

$$\vec{h}_{it} = \overrightarrow{GRU}(x_{it}), t \in [1, T] \quad (1)$$

The backward GRU  $\overleftarrow{f}$  reads sentences from  $w_{iT}$  to  $w_{i1}$ , as given in (2):

$$\overleftarrow{h}_{it} = \overleftarrow{GRU}(x_{it}), t \in [T, 1] \quad (2)$$

The annotation of the word  $w_{it}$  is calculated by combining the forward and backward hidden states i.e.,  $h_{it} = [\vec{h}_{it}, \overleftarrow{h}_{it}]$ .

The word-level attention layer extracts the words with major contributions to the essential meaning. These words form the sentence vector. That is, using a one-layer multi-layer perceptron (MLP) we calculate  $u_{it}$  to get the hidden representation of the  $h_{it}$  as given in (3):

$$u_{it} = \tanh(W_w h_{it} + b_w) \quad (3)$$

The layer uses the tanh function to ensure that the network does not falter. The tanh function achieves this by correcting input values to be between  $-1$  and  $1$  and also maps zeros to near-zero. Next, the normalized importance,  $a_{it}$  is generated using a sigmoid or a softmax function that calculates the impact of the word as a similarity of  $u_{it}$  with a word level context vector  $u_w$  as given in (4):

$$a_{it} = \frac{\exp(u_{it}^T u_w)}{\sum_t \exp(u_{it}^T u_w)} \quad (4)$$

After that, the weighted total of the word annotations is calculated using (5) thus generating the sentence vector  $s_i$ :

$$s_i = \sum_t a_{it} h_{it} \quad (5)$$

- **Sentence-level encoder and attention**

Similar to the word level, now the entire network is re-run on sentence that is focusing on the sentence  $i$ . There is no embedding layer as we already get sentence vectors  $s_i$  from word level as input. The bidirectional GRU sentence encoder is given in (6) & (7):

$$\vec{h}_i = \overrightarrow{GRU}(s_i), i \in [1, L] \quad (6)$$

$$\overleftarrow{h}_i = \overleftarrow{GRU}(s_i), t \in [L, 1] \quad (7)$$

The annotation of sentence  $i$  is computed by concatenating the forward and backward hidden states i.e.,  $h_i = [\vec{h}_i, \overleftarrow{h}_i]$ .

Similar to the word-level attention layer, the sentence-level attention layer extracts sentences that significantly convey the meaning of the document. Therefore, a one-layer MLP is used to calculate  $u_i$  and get the hidden representation of  $h_i$  as given in (8):

$$u_i = \tanh(W_s h_i + b_s) \quad (8)$$

Likewise, the normalized importance,  $a_i$  is also generated using a sigmoid or a softmax function. The similitude of  $u_i$  with a sentence-level context vector  $u_s$  is calculated to characterize the significance of the sentence as a as given in (9):

$$a_i = \frac{\exp(u_i^T u_s)}{\sum_i \exp(u_i^T u_s)} \quad (9)$$

The weighted total of the sentence annotations is then computed to generate the document vector. The document vector  $v_i$  summarizes all the data of the sentences present in the document as given in (10):

$$v_i = \sum_i a_i h_i \quad (10)$$

Trainable weights and biases are again randomly initialized and jointly learned during the training process as shown in (11) and (12).

$$u_i = \tanh(W_s h_i + b_s) \quad (11)$$

$$a_i = \frac{\exp(u_i^T u_s)}{\sum_i \exp(u_i^T u_s)} \quad (12)$$

The final output is a document vector  $v_i$  which can be used as features for document classification as depicted by (13)

$$v_i = \sum_i a_i h_i. \quad (13)$$

Finally, the document vector  $v$  is used to generate the final class. The final layer which uses a sigmoid activation takes as input  $v$  to output the classification result of the HAN sub-network as given in (14)

$$p = \text{sigmoid}(W_c v + b_c) \quad (14)$$

The training loss is taken as the negative of the log-likelihood of the correct labels as shown in (15)

$$L = -\sum_d \log p d_j \quad (15)$$

where,  $j$  stands for the label of document  $d$ .

### 3.2. Multi-Layer Perceptron (MLP)

The MLP forms the second classifier of our hybrid model, and it takes the user-based features as input. It does so by using the back-propagation algorithm and adjusting the weights of the neurons after every backward pass, thus minimizing the error as much as possible and finally attaining convergence [30]. The MLP architecture used in this research is as follows: the input layer consists of five neurons, the two hidden layers consist of four neurons each, and the output layer has a single neuron with a sigmoid activation function which generates the final output class, i.e. rumour (positive class) or non-rumour (negative class). Next, we fuse the outputs from the two classifiers, i.e. HAN and MLP using the Logical OR operation to generate the final class for a particular post.

### 3.3. Decision-level Fusion for Classification

In the *CanarDeep* model, the output decisions from both the classifiers, i.e., HAN and MLP, are fused using the Logical OR operation. If the decision from either classifier is that the given input is a rumour, then the input is classified as a rumour. The given input is classified as a non-rumour if and only if both the classifiers decide that the input is a non-rumour. Table 1 illustrates the logical OR operation.

**Table 1.** Decision-level fusion using logical OR

HAN	MLP	Decision
+ (Rumour)	+ (Rumour)	Rumour
- (Non- Rumour )	- (Non- Rumour )	Non-Rumour
+ ( Rumour )	- (Non-Rumour)	Rumour
- (Non-Rumour )	+ ( Rumour )	Rumour

The OR operation helps to debunk a rumour with maximum possibility. If both the classifiers detect the post as rumour then it is irrefutably a rumour. Textual content and its meta-features provide valuable markers to indicate a rumour and therefore even if only the context-based classifier is indicative of rumour, the output is marked as rumour. This is because rumours are driven based on psychology and behavior of users which may alter with change in beliefs, confusion and anxiety or due to uncertainty. Hence, even a non-suspicious account can spread rumours. Likewise, if a user profile is identified suspicious using the user-based classifier, it is also marked as rumour. The primary notion is that intelligent bots and masqueraded profiles tend to use professional services for believable content writing tactics which can often be missed by the context-based classifier. Thus, the ORing decides to discard a post only if both the classifiers classify it as a non-rumour.

#### 4. Results and Discussion

Two events in the PHEME dataset suffer from the class imbalance problem. These are Charlie Hebdo (CH) and Ferguson (FS), where the data is skewed towards the non-rumour category in comparison to the rumour category, while the other three i.e, Germanwings (GW), Ottawa Shooting (OS), and Sydney Siege (SS) do not. To resolve this, the performance of the *CanarDeep* model is examined with respect to each of the individual events to analyze how well the model performs across the dataset. Data skewness or class imbalance sabotages a classification task and using accuracy as a performance metric may lead to incorrect interpretation & evaluation [31, 32]. A relative-to-each-class measure (like ROC-AUC) is thus preferred over the absolute measure (accuracy). F1-Score, Precision, and Recall have been used to evaluate the classifier performance.

There are various parameters that have been used for both the sub-networks of our proposed model during the experiment. The values of these parameters are shown in table 2.

**Table 2.** Parameters used for HAN and MLP classifier

HAN	
Parameter	Value
Embedding Dimension	300
Bi-GRU Units	150
Hidden Units	300
Return Sequences	True
Trainable	True
Non-Linearity Function	ReLU
Loss Function	Binary Crossentropy
Optimizer	Adam
Dropout	0.5
Word Embedding	ELMo 5.5B
Batch Size	128
Epochs	7
Maximum Vocabulary Size	20000
Maximum Sentence Length	50
Maximum Sentence Number	5
MLP	
Max Iterations	10
Solver	Adam
Learning Rate Initializer	0.01
Batch Size	200
No. of Hidden Layers	2
Units in each Hidden Layer	4
Tolerance	1e-4
Activation Function	ReLU

A good understanding of how the proposed model performed on individual events can be obtained by taking a look at each event’s confusion matrix. The confusion matrices [29, 30] are calculated by counting four values for each event, namely, True Positives (TP), False Positives (FP), False Negatives (FN) and True Negatives (TN). The confusion matrices for each PHEME dataset event are shown in fig. 5(a) through 5(e), where with horizontal axis represents the actual class and the vertical axis represents the predicted class.

	Rumour	Non-Rumour
Rumour	168	70
Non-Rumour	46	185

Fig. 5(a). Germanwings

	Rumour	Non-Rumour
Rumour	410	112
Non-Rumour	240	459

Fig. 5(b). Sydney Siege

	Rumour	Non-Rumour
Rumour	135	149
Non-Rumour	215	644

Fig. 5(c). Ferguson

	Rumour	Non-Rumour
Rumour	344	126
Non-Rumour	144	276

Fig. 5(d). Ottawa Shooting

	Rumour	Non-Rumour
Rumour	293	165
Non-Rumour	285	1336

Fig. 5(e). Charlie Hebdo

The AUC scores of the five events present in the PHEME dataset range from 0.64 to 0.74. Fig. 6(a) through 6(e) displays the ROC curve for individual events. The proposed *CanarDeep* classifier performs better for two events, namely, the Germanwings and Sydney Siege in comparison to the other three.

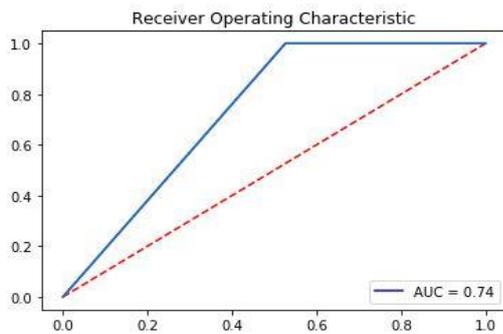


Fig. 6(a). ROC-Germanwings

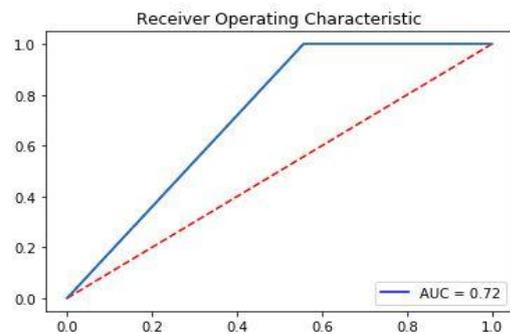


Fig. 6(b). ROC-Sydney Siege

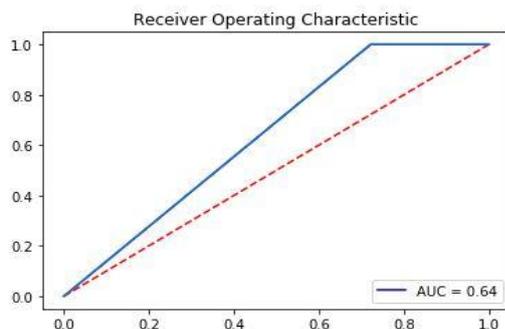


Fig. 6(c). ROC-Ferguson

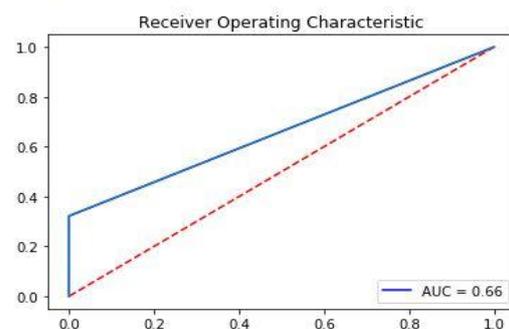
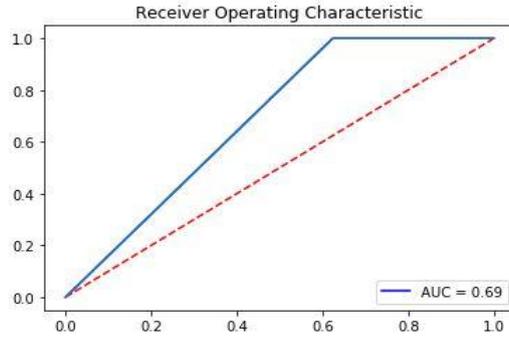


Fig. 6(d). ROC-Ottawa Shooting



**Fig. 6(e).** ROC-Charlie Hebdo

Our model performs fairly uniformly for all the five individual events in the dataset, as shown in table 3. Precision values generated by the *CanarDeep* model outperform the CRF classifier [15] for three events. Recall values achieved by the *CanarDeep* model are also higher across four out of the five events with CH being an exception. The proposed model manages to achieve a precision-recall equilibrium, transcending the state-of-the-art.

**Table 3.** Classifier performance for individual events

Event	Metric	CRF [15]	<i>CanarDeep</i>
GW	P	0.743	0.753
	R	0.668	0.755
	F1	0.704	0.754
CH	P	0.545	0.732
	R	0.762	0.698
	F1	0.636	0.715
OS	P	0.841	0.695
	R	0.585	0.696
	F1	0.690	0.695
SS	P	0.764	0.721
	R	0.385	0.717
	F1	0.512	0.719
FS	P	0.566	0.613
	R	0.394	0.599
	F1	0.465	0.606

GW: Germanwings, CH: Charlie Hebdo, OS: Ottawa Shooting, SS: Sydney Siege, FS: Ferguson, P: Precision, R: Recall, F1: F1 score

The SS event has the highest number of tweets among the events and does not suffer from the class imbalance problem too. The highest improvement in the F1 score is observed for this event with an improvement of 40.43%, as noted in table 3. A significant improvement in the recall score is observed for four events accounting for *CanarDeep*'s overall superior performance over the CRF classifier. Moreover, the proposed model does not let the event datasets suffering from class imbalance hamper its performance, another qualitative improvement over the CRF classifier. As compared to CRF which has F1 score of 0.636 for CH and 0.465 for FS, *CanarDeep* achieves F1 of 0.715 for CH and 0.606 for FS respectively. The proposed model follows a similar trend with datasets that are free from the class imbalance problem. Superior F1 scores are obtained for the GW, SS and OS events in comparison to the CRF classifier. Our previous work using attention-based residual network (ARN) [26] which reported an improvement over the state-of-the-art CRF in terms of recall and F1 scores was also compared to *CanarDeep*. The effectiveness of the *CanarDeep* model is assessed using the whole dataset, i.e., all the five events combined. Table 4 depicts the performance results on the complete data.

**Table 4.** Performance on complete dataset

Model	Metrics		
	P	R	F1
CRF [15]	0.667	0.556	0.607
ARN [26]	0.662	0.570	0.612
<i>CanarDeep</i>	0.685	0.592	0.634

Thus, training the hybrid HAN and MLP using user-based and context-based features with a mix of fusion strategies outperforms the existing state-of-the-art model for all the three metrics.

## 5. Conclusion

Rumours thrive in adversarial situations due to their high manipulative power and the inability of naïve users to differentiate between legit and false content. It is thus imperative to determine the potential rumour candidate at its emergence. As a key to classify online rumours, a deep learning model was proposed which coalesced information from two classifiers to label tweets as rumour in benchmark PHEME dataset. Two different input types, namely context-based and user-based were learned separately using the classifiers (HAN for context-based and MLP for user-based). The output predictions of these were then combined using a logical OR decision-level operation to finally classify the tweets. The benefit of using early-level fusion to concatenate textual and meta-features as context-based features is that it does not isolate interactions between correlated features whereas the advantage of using decision-level fusion for final output is that the model need not synchronize between different types of features. The robustness of the technique is validated for both individual events and the whole dataset. The experimental evaluation reveals superior performance in comparison with the existing state-of-the-art with a 4.45% gain in F1-score.

As rumours can harm and do irreparable damage, it is equally important to achieve high performance with interpretability and verifiability of decisions. We intend to use explainable artificial intelligence (XAI) to realize the action traceability and build robust, trustworthy and unbiased learning models to detect rumours in social data streams. Further future work also includes using characteristics of breaking-news and long-standing rumours for training learning models. Also, the current research only uses the rumour and non-rumour categories to label the post but future studies need to develop models that identify fine-grain rumour categories such as dread rumours, wish rumours, wedge-driving rumours and reputation rumours.

## References

1. Bounegru, L., Gray, J., Venturini, T., & Mauri, M. (2018). A Field Guide to 'Fake News' and Other Information Disorders. *A Field Guide to "Fake News" and Other Information Disorders: A Collection of Recipes for Those Who Love to Cook with Digital Methods*, Public Data Lab, Amsterdam (2018).
2. Li, G., Dong, M., Yang, F., Zeng, J., Yuan, J., Jin, C., ... & Zheng, B. (2020). Misinformation-oriented expert finding in social networks. *World Wide Web*, 23(2), 693-714.
3. Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., & Procter, R. (2018). Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)*, 51(2), 1-36.
4. Chen, C., Wen, S., Zhang, J., Xiang, Y., Oliver, J., Alelaiwi, A., & Hassan, M. M. (2017). Investigating the deceptive information in Twitter spam. *Future Generation Computer Systems*, 72, 319-326.
5. Kumar, A., & Sangwan, S. R. (2019). Rumor detection using machine learning techniques on social media. In *International conference on innovative computing and communications* (pp. 213-221). Springer, Singapore.
6. Alrubaiyan, M., Al-Qurishi, M., Alamri, A., Al-Rakhami, M., Hassan, M. M., & Fortino, G. (2018). Credibility in online social networks: A survey. *IEEE Access*, 7, 2828-2855.
7. Kumar, A., & Sharma, H. (2020). PROD: A potential rumour origin detection model using supervised machine learning. In *International Conference on Intelligent Computing and Smart Communication 2019* (pp. 1269-1276). Springer, Singapore.
8. Gu J, Wang Z, Kuen J, Ma L, Shahroudy A, Shuai B, Liu T, Wang X, Wang G, Cai J, Chen T. (2018). Recent advances in convolutional neural networks. *Pattern Recognition*, 77, 354-77.
9. Otter, D. W., Medina, J. R., & Kalita, J. K. (2020). A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2), 604-624.
10. Young T, Hazarika D, Poria S, Cambria E. (2018). Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence magazine*, 13(3), 55-75.
11. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., & Hovy, E. (2016, June). Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies* (pp. 1480-1489).
12. Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
13. Kumar, A. (2021). Contextual semantics using hierarchical attention network for sentiment classification in social internet-of-things. *Multimedia Tools & Application* <https://doi.org/10.1007/s11042-021-11262-8>
14. Ting, F. F., & Sim, K. S. (2017, November). Self-regulated multilayer perceptron neural network for breast cancer classification. In *2017 International Conference on Robotics, Automation and Sciences (ICORAS)* (pp. 1-5). IEEE.
15. Zubiaga, A., Liakata, M., & Procter, R. (2017, September). Exploiting context for rumour detection in social media. In *International Conference on Social Informatics* (pp. 109-123). Springer, Cham.

16. Zubiaga A, Wong Sak Hoi G, Liakata M, Procter R. PHEME dataset of rumours and non-rumours [Internet]. figshare; 2016 [cited 2020 May 24]. Available from: [https://figshare.com/articles/PHEME\\_dataset\\_of\\_rumours\\_and\\_non-rumours/4010619/1](https://figshare.com/articles/PHEME_dataset_of_rumours_and_non-rumours/4010619/1)
17. Cao, J., Guo, J., Li, X., Jin, Z., Guo, H., & Li, J. (2018). Automatic rumor detection on microblogs: A survey. *arXiv preprint arXiv:1807.03505*.
18. Takahashi, T., & Igata, N. (2012, November). Rumor detection on twitter. In *The 6th International Conference on Soft Computing and Intelligent Systems, and The 13th International Symposium on Advanced Intelligence Systems* (pp. 452-457). IEEE.
19. Kumar, A., Singh, V., Ali, T., Pal, S., & Singh, J. (2020). Empirical evaluation of shallow and deep classifiers for rumor detection. In *Advances in Computing and Intelligent Systems* (pp. 239-252). Springer, Singapore.
20. Bhattacharjee, U., Srijiith, P. K., & Desarkar, M. S. (2019, January). Term specific tf-idf boosting for detection of rumours in social networks. In *2019 11th International Conference on Communication Systems & Networks (COMSNETS)* (pp. 726-731). IEEE.
21. Chen, T., Li, X., Yin, H., & Zhang, J. (2018, June). Call attention to rumors: Deep attention based recurrent neural networks for early rumor detection. In *Pacific-Asia conference on knowledge discovery and data mining* (pp. 40-52). Springer, Cham.
22. Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B. J., Wong, K. F., & Cha, M. (2016). Detecting rumors from microblogs with recurrent neural networks, In *Proceedings of IJCAI*. 1, 3, 5, 6
23. Nguyen, T. N., Li, C., & Niederée, C. (2017, September). On early-stage debunking rumors on twitter: Leveraging the wisdom of weak learners. In *International Conference on Social Informatics* (pp. 141-158). Springer, Cham.
24. Guo, H., Cao, J., Zhang, Y., Guo, J., & Li, J. (2018, October). Rumor detection with hierarchical social attention network. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (pp. 943-951).
25. Vijeev, A., Mahapatra, A., Shyamkrishna, A., & Murthy, S. (2018, September). A hybrid approach to rumour detection in microblogging platforms. In *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)* (pp. 337-342). IEEE.
26. Kumar, A., Shrivastava, A. (2020). Rumour Detection in Benchmark Dataset using Attention-Based Residual Networks. *International Journal of Advanced Science and Technology*, 29(3), 14682 -. Retrieved from <http://sersc.org/journals/index.php/IJAST/article/view/31956>
27. Kumar, A. (2021). Rumour Stance Classification using A Hybrid of Capsule Network and Multi-Layer Perceptron. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 12(13), 4110-4120.
28. Kumar, A., Sachdeva, N. Multimodal cyberbullying detection using capsule network with dynamic routing and deep convolutional neural network. *Multimedia Systems* (2021). <https://doi.org/10.1007/s00530-020-00747-5>
29. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3), e0118432.
30. Hand, D. J. (2012). Assessing the performance of classification methods. *International Statistical Review*, 80(3), 400-414.
31. Sangwan, S. R., & Bhatia, M. P. S. (2020). D-BullyRumbler: a safety rumble strip to resolve online denigration bullying using a hybrid filter-wrapper approach. *Multimedia Systems*, 1-17.
32. Kumar, A., Dikshit, S., & Albuquerque, V. H. C. (2021). Explainable Artificial Intelligence for Sarcasm Detection in Dialogues. *Wireless Communications and Mobile Computing*, 2021.