

Article

HAT4RD: Hierarchical Adversarial Training for Rumor Detection in Social Media

Shiwen Ni ¹, Jiawen Li ^{1,2,3} and Hung-Yu Kao ^{1,*}¹ Department of Computer Science and Information Engineering, National Cheng Kung University, Tainan 70101, Taiwan² Maritime College, Guangdong Ocean University, Zhanjiang 524000, China³ Technical Research Center for Ship Intelligence and Safety Engineering of Guangdong Province, Zhanjiang 524000, China* Correspondence: hykao@mail.ncku.edu.tw

Abstract: With the development of social media, social communication has changed. While this facilitates people's communication and access to information, it also provides an ideal platform for spreading rumors. In normal or critical situations, rumors can affect people's judgment and even endanger social security. However, natural language is high-dimensional and sparse, and the same rumor may be expressed in hundreds of ways on social media. As such, the robustness and generalization of the current rumor detection model are in question. We proposed a novel hierarchical adversarial training method for rumor detection (HAT4RD) on social media. Specifically, HAT4RD is based on gradient ascent by adding adversarial perturbations to the embedding layers of post-level and event-level modules to deceive the detector. At the same time, the detector uses stochastic gradient descent to minimize the adversarial risk to learn a more robust model. In this way, the post-level and event-level sample spaces are enhanced, and we verified the robustness of our model under a variety of adversarial attacks. Moreover, visual experiments indicate that the proposed model drifts into an area with a flat loss landscape, thereby, leading to better generalization. We evaluate our proposed method on three public rumor datasets from two commonly used social platforms (Twitter and Weibo). Our experimental results demonstrate that our model achieved better results compared with the state-of-the-art methods.

Keywords: rumor detection; adversarial training; deep learning; social media



Citation: Ni, S.; Li, J.; Kao, H.-Y. HAT4RD: Hierarchical Adversarial Training for Rumor Detection in Social Media. *Sensors* **2022**, *22*, 6652. <https://doi.org/10.3390/s22176652>

Academic Editor: Rebeca P. Díaz Redondo

Received: 1 August 2022

Accepted: 29 August 2022

Published: 2 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Today, social media is a popular news source for many people. However, without automatic rumor-detection systems, social media can be a breeding ground for rumors. Rumors can seriously affect people's lives [1]. For instance, during the early outbreak of the current COVID-19 pandemic, rumors about a national lockdown in the United States fueled panic buying in groceries and toilet papers, disrupting the supply chain, exacerbating the demand-supply gap and worsening the issue of food insecurity among the socioeconomically disadvantaged and other vulnerable populations [2]. Setting up automatic rumor detection is therefore essential.

Automatic rumor detection is extremely challenging, and the greatest difficulty lies in spotting camouflaged rumors. As the saying goes, “a rumor has a hundred mouths”; these words indicate that the ways rumors are expressed constantly change as they spread. Some malicious rumormongers may deliberately modify rumor text information to escape manual detection [3]. Variability and disguise are the main characteristics of rumors, which means that a robust automatic rumor detection model is necessary. Unfortunately, most current rumor detection models are not robust enough to spot the various changes and disguises during the rumor propagation process.

As shown in Figure 1, we simulated the constantly changing process of rumors during their propagation and found that the general deep-learning model was too sensitive to sentence changes and disguise. A BERT-base [4] model trained on the rumor dataset PHEME [5] had a prediction confidence of 0.85 for the rumor “Police say shots fired at 3 #ottawa sites National War Memorial, Parliament Hill, and now Rideau shopping centre”; however, when the input is changed to “According to the government authority report: The shootings took place at three #ottawa locations the National War Memorial parliament Hill and now the Rideau shopping centre”, the model’s prediction confidence decreased from 0.85 to 0.47. However, the main meaning and label of the input rumor text did not change but the model prediction was incorrect. This result shows that the robustness and generalization of a traditional rumor detection model are poor, and the changes of a few words while the meaning of the sentence remains the same may cause significant changes in the prediction results.




Figure 1. A well-trained BERT-base (The BERT-base code is completed by Hugging Face (<https://github.com/huggingface/transformers> (accessed on 8 May 2022))) on model for rumor detection. The rumor detection model labeled a rumor as a non-rumor when the words underwent small changes but the meaning remained the same.

To alleviate that problem, we designed a novel rumor detection model called HAT4RD to enhance the generalization ability and robustness of an automatic rumor detection model. Our model detects rumors based on an event, which includes a source post and a certain number of replies. To make full use of the tweet object information and obtain a high-level representation, we took a hierarchical architecture as the skeleton of our model.

To enhance the robustness of our model, adversarial training is included in our model. Using more adversarial data to train the model can enhance the robustness and generalization of the model. However, natural language text space is sparse, and it is impossible to exhaust all possible changes manually to train a robust model. Thus, we perturb the sample space of post-level and event-level, respectively, to comprehensively improve the robustness of the model against changes in the text. The main contributions of this paper can be summarized as follows:

- We first propose a hierarchical adversarial training method that encourages the model to provide robust predictions under the perturbed post-level and event-level embedding spaces.
- We evaluate the proposed model HAT4RD on three real-world datasets. The experimental results demonstrate that our model outperforms state-of-the-art models.
- We prove through experiments that the proposed hierarchical adversarial training method can enhance the robustness and generalization of the model and prevent the model from being deceived by disguised rumors.

2. Related Work

2.1. Rumor Detection

With the development of artificial intelligence, existing automated rumor detection methods are mainly based on deep neural networks. MA et al. [6] were the first to use a deep learning network, an RNN (Recurrent Neural Network)-based model, for automatic misinformation detection. Chen et al. [7], Yu et al. [8] proposed an attention mechanism into an RNN or CNN (Convolutional Neural Network) model to process a certain number of sequential posts for debunking rumors. Ajao et al. [9] proposed a framework combining CNN and LSTM (Long Short-Term Memory) to classify rumors.

Shu et al. [10] delved into an explainable rumor detection model by using both news content and user comments. Guo et al. [11], Sujana et al. [12] detected rumors by creating a hierarchical neural network to obtain higher-level textual information representations. Yang et al. [13] proposed a rumor detection model that can handle both text and images. Ruchansky et al. [14] analyzed articles and extracted user characteristics to debunk rumors. Ma et al. [15] constructed a recursive neural network to handle conversational structure. Their model was presented as a bottom-up and top-down propagation tree-structured neural network.

Li et al. [16,17] used a variable-structure graph neural network to simulate rumor propagation and obtain more precise information representations in the rumor detection task. Ni et al. [1] used multi-view attention networks to simultaneously capture clue words in the rumor text and suspicious users in the propagation structure. Gumaiei et al. [18] proposed an extreme gradient boosting (XGBoost) classifier for rumor detection of Arabic tweets. Li et al. [19] combined objective facts and subjective views for an evidence-based rumor detection. No rumor detection model currently takes adversarial robustness into account.

2.2. Adversarial Training

Adversarial training is an important method to enhance the robustness of neural networks. Szegedy et al. [20] first proposed the theory of adversarial training by adding small generated perturbations on input images. The perturbed image pixels were later named as adversarial examples. Goodfellow et al. [21] proposed a fast adversarial example generation approach to attempt to obtain the perturbation value that maximizes adversarial loss. Jia and Liang [22] were the first to adopt adversarial example generation for natural language processing tasks.

Zhao et al. [23] found that when adopting the gradient-based adversarial training method on natural language processing tasks, the generated adversarial examples were invalid characters or word sequences. Gong et al. [24] utilized word vectors as the input for deep-learning models; however, this also generated words that could not be matched with any words in the word embedding space. Ni et al. [25] proposed a random masked weight adversarial training method to improve generalization of neural networks. However, thus far there is no adversarial training method designed for rumor-specific hierarchical structures.

3. Problem Definition

We define false information that is socially inconsistent with facts to be rumors. Furthermore, we define the task of rumor detection as determining whether it is a rumor based on the relevant information (such as the text content, comments and propagation patterns) of microblog posted on social media platforms. We treat the original post and its reply posts together as an event (see Figure 2 for a real-world example of an event) for rumor detection. A whole event as the final decision-making unit contains a wealth of internal logic and user stance information.

Multiple events in the dataset are defined as $D = \{E_1, E_2, \dots, E_{|e|}\}$. An event consists of a source post and several reply posts, $E_j = \{P_s, P_1, P_2, \dots, P_{|p|}\}$. It should be noted that different events are composed of different numbers of posts, and a post is composed of different words, meaning our model needs to be able to process variable-length sequence

information with a hierarchical structure. The event-level classifier can perform learning via labeled event data, that is, $E_j = \{P_s, P_1, P_2, \dots, P_{|p|}\} \rightarrow y_j$. In addition, because an event contains multiple posts, we make the posts within the same event share labels. The post-level classifier $P_n = \{x_1, x_2, \dots, x_{|x|}\} \rightarrow y_n$ can, therefore, be established.




Figure 2. Posts and events on social media (hierarchical structure of the post-level and event-level).

4. The Proposed Model HAT4RD

4.1. Preliminaries

Rumors in social media have a hierarchical structure of post-level and event-level. In response to this special data structure, we built the HAT4RD model based on the hierarchical BiLSTM (Bi-directional Long Short-Term Memory), which can be divided into the post-level module and event-level module, as shown in Figure 3. Hierarchical Adversarial Training (HAT) is a novel adversarial training method based on the hierarchical structure model. The overall hierarchical adversarial training procedure is shown in Algorithm 1. Taking the text of all posts under the event as input, we calculate the embedding of each word by Glove [26] word vectors to obtain the input of post-level BiLSTM. The formula is as follows:

$$I_p = \{x_1, x_2, \dots, x_n\} \quad (1)$$

where x_i is the pre-trained word vector, I_p is the input of post-level BiLSTM, and all the vectors with the posts as the unit pass through the post-level BiLSTM layer in proper order. For each time point t , the formula is as follows:

$$h_t^p = \text{BiLSTM}_p(x_i, h_{t-1}^p) \quad (2)$$

The cell state h_t^p of the uppermost LSTM_p at the last time point is used as the result of the post encoding. Due to the use of the bidirectional structure, the final state of both directions is joint, and an event can be represented by a matrix in which each column is a vector representing a post. The formula is as follows:

$$O_p = [h_s^p, h_1^p, h_2^p, \dots, h_{|p|}^p] \quad (3)$$

where h_s^p is the embedding of the source post. h_i^p is the embedding of a reply post, O_p is the output of post-level BiLSTM, and I_e is the input of event-level BiLSTM. The formula is as follows:

$$I_e = O_p = [h_s^p, h_1^p, h_2^p, \dots, h_{|p|}^p] \quad (4)$$

For the next module, the event-level BiLSTM encoding process is similar to post-level BiLSTM. The difference can be seen in the input data unit; post-level BiLSTM uses a post vector composed of word vectors, while event-level BiLSTM uses an event vector composed of post vectors. The formula is as follows:

$$h_t^e = \text{BiLSTM}_e(h_t^p, h_{t-1}^e) \quad (5)$$

In the rumor detection task, the state h_t^e of the event-level BiLSTM, the last layer at the last time point can be understood as a comprehensive representation of all posts.




Figure 3. The architecture of the proposed model HAT4RD.

Based on the principle of multi-task learning, rumor post classification and rumor event classification are highly related, and the parameters of the post-level module are shared in the two tasks. A post-level auxiliary classifier and an event-level primary classifier were therefore included in the hierarchical model. The post-level auxiliary classifier is mainly for accelerating training and preventing “vanishing gradient”. Two classifiers were used to obtain post-level prediction results and event-level prediction results. The formula is as follows:

$$\hat{y}_p = \text{softmax}(W_p \cdot h_t^p + b_p) \quad (6)$$

$$\hat{y}_e = \text{softmax}(W_e \cdot h_t^e + b_e) \quad (7)$$

where $\hat{y}_p$ and $\hat{y}_e$ are the post and event classification results, respectively; W_p and W_e are the weights of the fully connected layers; and b_p and b_e are the biases. The goal of each training process is to minimize the standard deviation between the predicted and output values using the following loss function:

$$L_p = -y \log(\hat{y}_{pr} - (1 - y_p) \log(1 - \hat{y}_{pn})) \quad (8)$$

$$L_e = -y \log(\hat{y}_{er} - (1 - y_e) \log(1 - \hat{y}_{en})) \quad (9)$$

$$L_t = \alpha L_p + (1 - \alpha) L_e \quad (10)$$

where L_p and L_e are the post-level loss and event-level loss, respectively. α is the loss coefficient weight to control L_p and L_e . L_t is the total loss of the entire rumor detection model used to update the parameters. y is the real label; $\hat{y}_r$ and $\hat{y}_n$ are the two labels predicted by the model: rumor and non-rumor. The gradient of the model was calculated according to Loss L_t . The formula is as follows:

$$g = \nabla_{\theta} L_t(\theta, x, y) \quad (11)$$

Algorithm 1 Hierarchical adversarial training algorithm

Input: Training samples $\mathcal{X}$, perturbation coefficient ϵ_p and ϵ_e , Loss coefficient weight α , Learning rate τ , Parameter: θ

```

1: for epoch = 1 ...  $N_{ep}$  do
2:   for  $(x, y) \in \mathcal{X}$  do
3:     Forward-propagation calculation Loss:
4:      $L_p \leftarrow -y \log(\hat{y}_{pr}) - (1 - y_p) \log(1 - \hat{y}_{pn})$ 
5:      $L_e \leftarrow -y \log(\hat{y}_{er}) - (1 - y_e) \log(1 - \hat{y}_{en})$ 
6:      $L_t \leftarrow \alpha L_p + (1 - \alpha) L_e$ 
7:     Backward-propagation calculation gradient:
8:      $g_p \leftarrow \nabla_{x_p} L_t(\theta, x_p, (y_p, y_e)); g_e \leftarrow \nabla_{x_e} L_e(\theta, x_e, y_e)$ 
9:     Compute hierarchical adversarial perturbation:
10:     $r_p \leftarrow \epsilon_p \cdot g_p / \|g_p\|_2; r_e \leftarrow \epsilon_e \cdot g_e / \|g_e\|_2$ 
11:    Forward-Backward-propagation calculation adversarial gradient:
12:     $g_{adv}^p \leftarrow \nabla_{\theta} L_{t_{adv}}^p(\theta, x_p + r_p, (y_p, y_e))$ 
13:     $g_{adv}^e \leftarrow \nabla_{\theta} L_{e_{adv}}^e(\theta, x_e + r_e, y_e)$ 
14:    Update parameter:
15:     $\theta \leftarrow \theta - \tau(g + g_{adv}^p + g_{adv}^e)$ 
16:   end for
17: end for
18: Output:  $\theta$ 

```

4.2. Hierarchical Adversarial Training

The above is a forward propagation under standard training of the model. To enhance the robustness of our model, a hierarchical adversarial training method is adopted. This adversarial optimization process was expressed with the following Min-Max formula:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim D} \left\{ \max_{\delta_p, \delta_e \in S} [L_t(\theta, x_p + \delta_p, (y_p, y_e)) + L_e(\theta, x_e + \delta_e, y_e)] \right\} \quad (12)$$

where δ_p and δ_e are the perturbations of the post-level input x_p and event-level input x_e under maximization of the internal risk. We, respectively, estimated these values by linearizing $\nabla_{x_p} L_t(\theta, x_p, (y_p, y_e))$ and $\nabla_{x_e} L_e(\theta, x_e, y_e)$ around x_p and x_e . Using the $\nabla_{x_p} L_t(\theta, x_p, (y_p, y_e))$ and $\nabla_{x_e} L_e(\theta, x_e, y_e)$ linear approximation in Equations (13) and (14) and the L2 norm constraint, the resulting adversarial perturbations are:

$$\delta_p = \epsilon_p \cdot \frac{\nabla_{x_p} L_t(\theta, x_p, (y_p, y_e))}{\|\nabla_{x_p} L_t(\theta, x_p, (y_p, y_e))\|_2} \quad (13)$$

$$\delta_e = \epsilon_e \cdot \frac{\nabla_{x_e} L_e(\theta, x_e, y_e)}{\|\nabla_{x_e} L_e(\theta, x_e, y_e)\|_2} \quad (14)$$

where ϵ_p and ϵ_e are the perturbation coefficients. Note that the value of the perturbation δ_p is calculated based on the back-propagation of the total Loss instead L_t of L_p , because the addition of the perturbation δ_p makes L_p and L_e increase at the same time.

4.2.1. Post-Level Adversarial Training

After a normal forward and backward propagation, δ_p and δ_e were calculated according to the gradient. Using post-level adversarial training, we added word-level perturbation to the word vector to obtain the input of post-level BiLSTM, and the formula is as follows:

$$I_{p_{adv}} = \{x_1 + \delta_1^p, x_2 + \delta_2^p, \dots, x_n + \delta_n^p\} \quad (15)$$

where $I_{p_{adv}}$ is the adversarial input of post-level BiLSTM, and δ_n^p is the post-level perturbation added to the word vector x_n . All the vectors with the posts as the unit then pass through the post-level BiLSTM layer in proper order. For each time point t , the formula is as follows:

$$h_t^{p_{adv}} = \text{BiLSTM}_p(x_i + \delta_i^p, h_{t-1}^{p_{adv}}). \quad (16)$$

The adversarial cell state $h_t^{p_{adv}}$ of the uppermost LSTM_p at the last time point is used as the result of the post encoding. Due to the use of the bidirectional structure, the final state of both directions is joint, and an event can be represented by a matrix in which each column is a vector representing a post. The formula is as follows:

$$O_{p_{adv}} = [h_s^{p_{adv}}, h_1^{p_{adv}}, h_2^{p_{adv}}, \dots, h_{|p|}^{p_{adv}}] \quad (17)$$

where $h_s^{p_{adv}}$ is the adversarial result of the post-level BiLSTM, that is, the embedding of the source post. $h_i^{p_{adv}}$ is the adversarial embedding of the reply post, and $O_{p_{adv}}$ is the adversarial output of post-level BiLSTM and input of event-level BiLSTM. The formula is as follows:

$$h_t^{e_{adv}} = \text{BiLSTM}_p(h_t^{p_{adv}} + \delta_t^e, h_{t-1}^{e_{adv}}) \quad (18)$$

Finally, h_t^e was replaced with $h_t^{e_{adv}}$ and the adversarial loss $L_{p_{adv}}^p$, $L_{e_{adv}}^p$ and $L_{t_{adv}}^p$ of post-level perturbation can be calculated using Equations (6)–(9). The post-level adversarial gradient g_{adv}^p is calculated based on the result of backpropagation. The formula is as follows:

$$g_{adv}^p = \nabla_{\theta} L_{t_{adv}}^p(\theta, x_p + \delta_p, (y_p, y_e)) \quad (19)$$

4.2.2. Event-Level Adversarial Training

We next performed event-level adversarial training and repeated the process of Equations (1)–(3) to obtain the post vector. Event-level perturbation was then added to the post vector to obtain the adversarial input of event-level BiLSTM, and the formula is as follows:

$$I_{e_{adv}} = \{h_s^p + \delta_s^p, h_1^p + \delta_1^p, h_2^p + \delta_2^p, \dots, h_{|p|}^p + \delta_{|p|}^p\} \quad (20)$$

In the same way, input $I_{e_{adv}}$ into the event-level BiLSTM to obtain the final event representation vector $h_t^{e_{adv}}$, replace h_t^e with $h_t^{e_{adv}}$ and calculate the adversarial loss $L_{e_{adv}}^e$ of event-level perturbation through Equations (6)–(9). Finally, the post-level adversarial gradient g_{adv}^e is calculated based on backpropagation. The formula is as follows:

$$g_{adv}^e = \nabla_{\theta} L_{e_{adv}}^e(\theta, x_e + \delta_e, y_e) \quad (21)$$

Finally, the gradient is calculated by the standard training; the gradient calculated by the post-level adversarial training and the gradient calculated by the event-level adversarial training were used to update the model parameters. The parameter update process is expressed as:

$$\theta \leftarrow \theta - \tau(g + g_{adv}^p + g_{adv}^e) \quad (22)$$

where τ is the learning rate.

5. Experiments

5.1. Datasets

Three well-known public rumor datasets, PHEME 2017, PHEME 2018 [5] and WEIBO [6], were used to evaluate our method HAT4RD. Among them, the original data of PHEME 2017 and PHEME 2018 are from the Twitter social platform, and the language is English; the original data of WEIBO is from the Sina Weibo social platform, and the language is Simplified Chinese. In these three datasets, each event is composed of a source post and several reply posts. The statistical details of these three datasets are shown in Table 1.

“Users” represents the number of users in the datasets; “Posts” represents the number of posts in the datasets; “Event” represents the number of events in the datasets (that is, the number of source posts); “Avg words/post” represents the average number of words contained in a post; “Avg posts/event” represents the average number of posts contained in an event; “Rumor” represents the number of rumors in the datasets; “Non-rumor” represents the number of non-rumors in the datasets; and “Balance degree” represents the percentage of rumors in the datasets.

Table 1. The statistics of datasets used in this paper.

Statistic	PHEME 2017	PHEME 2018	WEIBO
Users	49,345	50,593	2,746,818
Posts	103,212	105,354	3,805,656
Events	5802	6425	4664
Avg words/post	13.6	13.6	23.2
Avg posts/event	17.8	16.3	816.0
Rumor	1972	2402	2313
Non-rumor	3830	4023	2351
Balance degree	34.00%	37.40%	49.59%

5.2. Evaluation Metrics

For a fair comparison, we adopted the same evaluation metrics used in previous work [19]. Therefore, the Accuracy, Precision Recall and F1-measure (F1) were adopted for evaluation, which is described in the following equations:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (23)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (24)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (25)$$

$$\text{F1_Measure} = \frac{2 * \text{Recall} * \text{Precision}}{\text{Recall} + \text{Precision}} \quad (26)$$

where TP are the true positive, TN are the true negative, FP are the false positive and FN are the false negative predictions.

5.3. Experimental Settings

Following the work of [19], the datasets were split for our experiment: 80% for training, 10% for validation and 10% for testing. We trained all the models by employing the derivative of the loss function through backpropagation and used the Adam optimizer [27] to update the parameters. From post text to its embedding, we used Glove’s [26] pre-trained 300-dim word vector.

For the hyperparameters, the maximum value of vocabulary was 80,000; the batch size was 64, the dropout rate was 5, the BiLSTM hidden size unit was 512, the loss coefficient weight α was 0.1, the learning rate was 0.0001, and the perturbation coefficients ϵ_p and ϵ_e were 1.0 and 0.3. Our proposed model was finally trained for 100 epochs with early

stopping. In addition, all experiments were run under the following hardware environment: CPU: Intel(R) Core(TM) i7-8700 CPU@3.20GHz, GPU: GeForce RTX 2080, 10G.

5.4. Performance Comparison

Our HAT4RD model was compared with other well-known rumor detection models to evaluate our model's rumor debunking performance.

- SVM-BOW: a rumor detection naive baseline, which is an SVM that uses bag-of-words for word representation [15].
- TextCNN: a rumor detection naive baseline based on deep convolutional neural networks [28].
- BiLSTM: a RNN-based bidirectional model that detects rumor by considering the bidirectional information [29].
- BERT: a well-known pre-trained language model. We fine-tuned a BERT-base to detect rumors [4].
- CSI: a state-of-the-art model detecting rumor by scoring users based on their behavior [14].
- CRNN: a hybrid model that combines recurrent neural network and convolutional neural network to detect rumors [30].
- RDM: a rumor detection model that integrates reinforcement learning and deep learning for early rumor detection [31].
- CSRD: a rumor detection model that classifies rumors by simulating comments' conversation structure using GraphSAGE [16].
- EHCS-Con: a model exploited the user's homogeneity by using the node2vec mechanism encoding user's follow-followers relationship for rumor detection [17].
- LOSIRD: a state-of-the-art rumor detection model that leverages objective facts and subjective views for interpretable rumor detection [19].

5.5. Main Experiment Results

The results of different rumor detection models are compared in Table 2; the HAT4RD clearly performed the best in terms of rumor detection compared to the other methods based on the three datasets with 92.5% accuracy on PHEME 2017, 93.7% on PHEME 2018 and 94.8% on WEIBO. In addition, the precision, recall and F1 were all higher than 91% in the HAT4RD model. Our HAT4RD improved on the F1 value of the SOTA model by about 1.5% on the dataset WEIBO. These results demonstrate the effectiveness of the hierarchical structure model and hierarchical adversarial training in rumor detection. However, the SVM-BOW result is poor because the traditional statistical machine-learning method could not handle this complicated task.

The results of the CNN, BiLSTM, BERT and RDM models were poorer than ours due to their insufficient information extraction capabilities. The models are based on post-processing information and cannot obtain a high-level representation from the hierarchy. Compared to other models, our HAT4RD model has a hierarchical structure and performs different levels of adversarial training. This enhances both the post-level and event-level sample space and improves the robustness and generalization of the rumor detection model.

5.6. Ablation Analysis


To evaluate the effectiveness of every component of the proposed HAT4RD, we removed each one of them from the entire model for comparison. "ALL" denotes the entire model HAT4RD with all components, including post-level adversarial training (PA), event-level adversarial training (EA), the post-level auxiliary classifier (PC) and event-level primary classifier (EC). After the removing, we obtained the sub-models "-PA", "-EA", "-PC" and "-EC", respectively. "-PA-PC" means that both the post-level adversarial training and auxiliary classifier were removed. "-PA-EA" denotes the reduced HAT4RD without both post-level adversarial training and event-level adversarial training. The results are shown in Figure 4.

Table 2. The results of different methods on three datasets. We report their average of five runs.

Method	PHEME 2017				PHEME 2018				WEIBO			
	Acc	Pre	Rec	F1	Acc	Pre	Rec	F1	Acc	Pre	Rec	F1
SVM-BOW	66.9	53.5	52.4	51.9	68.8	51.8	51.2	50.4	72.3	63.5	67.4	65.6
TextCNN	78.7	73.7	70.2	71.0	79.4	73.2	67.3	68.6	84.2	73.3	77.9	75.5
BiLSTM	79.5	76.3	69.1	70.6	79.6	72.7	67.7	68.9	85.7	83.1	89.6	86.4
BERT	86.5	85.9	85.1	85.5	84.4	83.4	83.5	83.5	90.7	89.4	89.7	89.5
CSI	85.7	84.3	85.9	85.1	85.1	83.6	85.5	84.5	91.4	90.4	90.7	90.5
CRNN	85.5	84.6	85.4	85.0	86.2	85.7	85.6	85.6	91.1	90.2	91.8	91.0
RDM	87.3	81.7	82.3	82.0	85.8	84.7	85.9	85.2	92.7	91.6	93.7	92.6
CSRD	90.0	89.3	86.9	88.1	91.9	89.2	92.3	90.7	92.4	91.5	91.7	91.6
EHCS-Con	91.2	90.5	90.5	90.5	92.3	92.3	92.3	92.4	93.0	92.2	92.6	92.4
LOSIRD [†]	91.4	91.5	90.0	90.6	92.5	92.2	92.4	92.3	93.2	92.3	92.7	92.5
HAT4RD[‡]	92.5	91.7	91.1	91.7	93.8	93.1	93.6	93.4	94.8	93.8	94.2	94.0

[†] The state-of-the-art model; [‡] Our model.

It can be observed that every component plays a significant role in improving the performance of HAT4RD. HAT4RD outperforms ALL-PA and ALL-EA, which shows that the post-level adversarial training and event-level adversarial training are indeed helpful in rumor detection. Both ALL-PA and ALL-EA were better than ALL-PA-EA, which shows that hierarchical adversarial training was more efficient than single-level adversarial training. The performance of ALL-PC was lower than that of HAT4RD, proving that the post-level auxiliary classifier contributes to the learning and convergence of the model.

**Figure 4.** HAT4RD ablation analysis in accuracy.

5.7. Early Rumor Detection

Our model's performance in early rumor detection was evaluated. To simulate the early stage rumor detection scenarios in the real world, nine different size test sets from PHEME 2017, PHEME 2018 and WEIBO were created. Each test set contained a certain number of posts, ranging from 5 to 45. We found that HAT4RD could detect rumors with an approximate 91% accuracy rate with only five posts as illustrated in Figure 5. Compared to the other models, our model uses hierarchical adversarial training and continuously generates optimal adversarial samples to join the training. It, therefore, has good generalization despite limited information.




Figure 5. Early rumor detection accuracy.

5.8. Robustness Analysis

We used *OpenAttack* (<https://github.com/thunlp/OpenAttack> (accessed on 8 May 2022)) [32] to conduct a variety of adversarial attacks on the models and compared the robustness of various recent models. FSGM draws from [21], which is a gradient-based adversarial attack method. HotFlip [33] uses gradient-based word or character substitution to attack. PWWS [34] uses a greedy word substitution order determined by the word saliency and weighted by the classification probability. As shown in Tables 3 and 4, our model can maintain the minimum performance degradation under the three adversarial attacks compared to other baseline models. In particular for gradient-based attacks, the robustness of our model is clear. Under the attack of FSGM, the performance of our model only dropped by about 10%. Under the attacks of HotFlip and PWWS, our model HAT4RD was also significantly more robust than other models.

Table 3. The classification accuracy of models on the PHEME 2017, PHEME 2018 datasets and the perturbed datasets using different attacking methods.

Method	PHEME 2017				PHEME 2018			
	Original	FSGM	HotFlip	PWWS	Original	FSGM	HotFlip	PWWS
Bert	0.865	0.432	0.426	0.425	0.834	0.412	0.431	0.442
EHCS-Con	0.912	0.547	0.576	0.531	0.923	0.568	0.513	0.412
LOSIRD	0.914	0.576	0.415	0.287	0.922	0.534	0.501	0.491
HAT4RD	0.925	0.846	0.786	0.534	0.932	0.835	0.744	0.615

Table 4. The classification accuracy of models on the WEIBO chinese dataset and the perturbed datasets using different attacking methods.

Method	WEIBO			
	Original	FSGM	HotFlip	PWWS
Bert	0.907	0.513	0.442	0.474
EHCS-Con	0.930	0.673	0.543	0.484
LOSIRD	0.932	0.624	0.654	0.446
HAT4RD	0.948	0.853	0.726	0.696

5.9. The Impact of Hierarchical Adversarial Training on the Loss Landscape

To further visually analyze the effectiveness of the hierarchical adversarial training method, we drew the high-dimensional non-convex loss function with a visualization method (<https://github.com/tomgoldstein/loss-landscape> (accessed on 8 May 2022)) proposed by [35]. We visualize the loss landscapes around the minima of the empirical risk selected by standard and hierarchical adversarial training with the same model structure.

The 2D and 3D views are plotted in Figure 6. We defined two direction vectors, d_x and d_y with the same dimensions as θ , drawn from a Gaussian distribution with zero mean and a scale of the same order of magnitude as the variance of layer weights. We then chose a center point θ^* and added a linear combination of α and β to obtain a loss that is a function of the contribution of the two random direction vectors.

$$f(d_x, d_y) = L(\theta^* + \alpha d_x + \beta d_y). \quad (27)$$

The results show that the hierarchical adversarial training method indeed selects flatter loss landscapes by dynamically generating post-level perturbation and event-level perturbation. Having a flatter loss function indicates that the model is more robust in input features and can prevent the model from overfitting. Empirically, many studies have shown that a flatter loss landscape usually means better generalization [36–38].




Figure 6. 2D and 3D visualization of the minima of the loss function selected by standard training (a–f) and hierarchical adversarial training (a*–f*) on the PHEME 17 (a,a*,d,d*), PHEME 18 (b,b*,e,e*) and WEIBO (c,c*,f,f*) datasets.

6. Conclusions

Herein, we proposed a new hierarchical adversarial training for rumor detection that considers the camouflages and variability of rumors from an adversarial perspective. Dynamically generating perturbations on the post-level and event-level embedding vectors enhanced the model’s robustness and generalization.

The evaluations of three real-world rumor detection datasets on social media showed that our HAT4RD model outperformed the state-of-the-art methods. Numerically, our proposed HAT4RD was 1.1%, 1.1% and 1.5% higher compared with the F1 of the state-of-the-art model LOSIRD on the three public rumor detection datasets, respectively. The early rumor detection performance of our model also outperformed the other models.

We examined the contribution of each part to the model performance through ablation experiments. Moreover, visual experiments proved that the hierarchical adversarial training method we proposed can optimize the model for a flatter loss landscape. Our HAT4RD model is general and can be applied to data on any topic, as long as the data is posted on social media (e.g., Twitter and Weibo). The ability of our model depends on the training dataset. We only need to add the corresponding data to the model training to detect rumors of different topics.

7. Future Work

Robustness and generalization are the focus of rumor detection. In the future, we can integrate features, such as text and images for multi-modal adversarial training, to

further enhance the model. In addition, for the unique structure of posts and events, we propose that graph neural networks will also be a good research direction, and graph neural networks can be combined with adversarial training to obtain graph adversarial training. We think this will be an interesting research direction. As rumor data collection and labeling are complicated and time-consuming, the recently popular prompt learning based on pre-trained language models for few-shot rumor detection is also worth studying.

8. Limitations

Finally, our model currently has certain limitations. Since our model includes hierarchical adversarial training, the training time is longer than the general model. Moreover, although our hierarchical adversarial training improves the robustness of the model, our model still has room for improvement due to the diversity of rumors and the sparsity of natural language.

Author Contributions: Conceptualization, S.N. and J.L.; methodology, S.N. and J.L.; investigation, S.N. and H.-Y.K.; writing—original draft preparation, S.N.; writing—review and editing, S.N. and J.L.; supervision, H.-Y.K.; funding acquisition, H.-Y.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Ministry of Science and Technology, Taiwan, under Grant MOST 111-2221-E-006-001.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ni, S.; Li, J.; Kao, H.Y. MVAN: Multi-View Attention Networks for Fake News Detection on Social Media. *IEEE Access* **2021**, *9*, 106907–106917. [\[CrossRef\]](#)
2. Tasnim, S.; Hossain, M.M.; Mazumder, H. Impact of rumors and misinformation on COVID-19 in social media. *J. Prev. Med. Public Health* **2020**, *53*, 171–174. [\[CrossRef\]](#) [\[PubMed\]](#)
3. Ni, S.; Li, J.; Kao, H.Y. True or False: Does the Deep Learning Model Learn to Detect Rumors? In Proceedings of the 2021 International Conference on Technologies and Applications of Artificial Intelligence (TAAI), Taichung, Taiwan, 18–20 November 2021; pp. 119–124.
4. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Minneapolis, MN, USA, 2–7 June 2019; Volume 1, pp. 4171–4186.
5. Kochkina, E.; Liakata, M.; Zubiaga, A. *PHEME Dataset for Rumour Detection and Veracity Classification*; Figshare: London, UK, 2018.
6. Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B.J.; Wong, K.F.; Cha, M. Detecting rumors from microblogs with recurrent neural networks. In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016), New York, NY, USA, 9–15 July 2016; pp. 3818–3824.
7. Chen, T.; Li, X.; Yin, H.; Zhang, J. Call attention to rumors: Deep attention based recurrent neural networks for early rumor detection. In *Trends and Applications in Knowledge Discovery and Data Mining, Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining, Melbourne, VIC, Australia, 3 June 2018*; Springer: Cham, Switzerland, 2018; pp. 40–52.
8. Yu, F.; Liu, Q.; Wu, S.; Wang, L.; Tan, T. Attention-based convolutional approach for misinformation identification from massive and noisy microblog posts. *Comput. Secur.* **2019**, *83*, 106–121. [\[CrossRef\]](#)
9. Ajao, O.; Bhowmik, D.; Zargari, S. Fake news identification on twitter with hybrid cnn and rnn models. In Proceedings of the 9th International Conference on Social Media and Society, Copenhagen, Denmark, 18–20 July 2018; pp. 226–230.
10. Shu, K.; Cui, L.; Wang, S.; Lee, D.; Liu, H. defend: Explainable fake news detection. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 395–405.
11. Guo, H.; Cao, J.; Zhang, Y.; Guo, J.; Li, J. Rumor detection with hierarchical social attention network. In Proceedings of the 27th ACM International Conference on Information and Knowledge Management, Torino, Italy, 22–26 October 2018; pp. 943–951.
12. Sujana, Y.; Li, J.; Kao, H.Y. Rumor Detection on Twitter Using Multiloss Hierarchical BiLSTM with an Attenuation Factor. In Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing, Suzhou, China, 4–7 December 2020; pp. 18–26.

13. Yang, Y.; Zheng, L.; Zhang, J.; Cui, Q.; Li, Z.; Yu, P.S. TI-CNN: Convolutional neural networks for fake news detection. *arXiv* **2018**, arXiv:1806.00749.
14. Ruchansky, N.; Seo, S.; Liu, Y. CSI: A hybrid deep model for fake news detection. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, Singapore, 6–10 November 2017; pp. 797–806.
15. Ma, J.; Gao, W.; Wong, K.F. Rumor detection on twitter with tree-structured recursive neural networks. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, Melbourne, VIC, Australia, 15–20 July 2018.
16. Li, J.; Sujana, Y.; Kao, H.Y. Exploiting microblog conversation structures to detect rumors. In Proceedings of the 28th International Conference on Computational Linguistics, Barcelona, Spain, 8–13 December 2020; pp. 5420–5429.
17. Li, J.; Ni, S.; Kao, H.Y. Birds of a Feather Rumor Together? Exploring Homogeneity and Conversation Structure in Social Media for Rumor Detection. *IEEE Access* **2020**, *8*, 212865–212875. [[CrossRef](#)]
18. Gumaei, A.; Al-Rakhami, M.S.; Hassan, M.M.; De Albuquerque, V.H.C.; Camacho, D. An effective approach for rumor detection of Arabic tweets using extreme gradient boosting method. *Trans. Asian-Low-Resour. Lang. Inf. Process.* **2022**, *21*, 1–16. [[CrossRef](#)]
19. Li, J.; Ni, S.; Kao, H.Y. Meet The Truth: Leverage Objective Facts and Subjective Views for Interpretable Rumor Detection. In Proceedings of the Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, Online, 1–6 August 2021.
20. Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; Fergus, R. Intriguing properties of neural networks. In Proceedings of the 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, 14–16 April 2014.
21. Goodfellow, I.J.; Shlens, J.; Szegedy, C. Explaining and harnessing adversarial examples. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015.
22. Jia, R.; Liang, P. Adversarial Examples for Evaluating Reading Comprehension Systems. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark, 7–11 September 2017; pp. 2021–2031.
23. Zhao, Z.; Dua, D.; Singh, S. Generating Natural Adversarial Examples. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
24. Gong, Z.; Wang, W.; Li, B.; Song, D.; Ku, W.S. Adversarial texts with gradient methods. *arXiv* **2018**, arXiv:1801.07175.
25. Ni, S.; Li, J.; Kao, H.Y. DropAttack: A Masked Weight Adversarial Training Method to Improve Generalization of Neural Networks. *arXiv* **2021**, arXiv:2108.12805.
26. Pennington, J.; Socher, R.; Manning, C.D. Glove: Global vectors for word representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543.
27. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
28. Chen, Y.C.; Liu, Z.Y.; Kao, H.Y. Ikm at semeval-2017 task 8: Convolutional neural networks for stance detection and rumor verification. In Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), Vancouver, BC, Canada, 3–4 August 2017; pp. 465–469.
29. Augenstein, I.; Rocktäschel, T.; Vlachos, A.; Bontcheva, K. Stance detection with bidirectional conditional encoding. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, ACL, Austin, TX, USA, 1–5 November 2016; pp. 876–885.
30. Liu, Y.; Wu, Y.F.B. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
31. Zhou, K.; Shu, C.; Li, B.; Lau, J.H. Early rumour detection. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Minneapolis, MN, USA, 2–7 June 2019; Long and Short Papers; Volume 1, pp. 1614–1623.
32. Zeng, G.; Qi, F.; Zhou, Q.; Zhang, T.; Hou, B.; Zang, Y.; Liu, Z.; Sun, M. Openattack: An open-source textual adversarial attack toolkit. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations, Online, 1–6 August 2021; pp. 363–371.
33. Ebrahimi, J.; Rao, A.; Lowd, D.; Dou, D. HotFlip: White-Box Adversarial Examples for Text Classification. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, Melbourne, VIC, Australia, 15–20 July 2018; Short Papers; Volume 2, pp. 31–36.
34. Ren, S.; Deng, Y.; He, K.; Che, W. Generating natural language adversarial examples through probability weighted word saliency. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 1085–1097.
35. Li, H.; Xu, Z.; Taylor, G.; Studer, C.; Goldstein, T. Visualizing the loss landscape of neural nets. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018; pp. 6391–6401.
36. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)] [[PubMed](#)]
37. Keskar, N.S.; Nocedal, J.; Tang, P.T.P.; Mudigere, D.; Smelyanskiy, M. On large-batch training for deep learning: Generalization gap and sharp minima. In Proceedings of the 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017.
38. Ishida, T.; Yamane, I.; Sakai, T.; Niu, G.; Sugiyama, M. Do We Need Zero Training Loss After Achieving Zero Training Error? In Proceedings of the International Conference on Machine Learning, PMLR, Online, 13–18 July 2020; pp. 4604–4614.