

Probing Spurious Correlations in Popular Event-Based Rumor Detection Benchmarks

Jiaying Wu^(✉) and Bryan Hooi

School of Computing, National University of Singapore, Singapore, Singapore
jiayingwu@u.nus.edu, bhooi@comp.nus.edu.sg

Abstract. As social media becomes a hotbed for the spread of misinformation, the crucial task of rumor detection has witnessed promising advances fostered by open-source benchmark datasets. Despite being widely used, we find that these datasets suffer from spurious correlations, which are ignored by existing studies and lead to severe overestimation of existing rumor detection performance. The spurious correlations stem from three causes: (1) *event-based* data collection and labeling schemes assign the same veracity label to multiple highly similar posts from the same underlying event; (2) merging multiple data sources spuriously relates source identities to veracity labels; and (3) labeling bias. In this paper, we closely investigate three of the most popular rumor detection benchmark datasets (i.e., Twitter15, Twitter16 and PHEME), and propose *event-separated rumor detection* to eliminate spurious cues. Under the event-separated setting, we observe that the accuracy of existing state-of-the-art models drops significantly by over 40%, becoming only comparable to a simple neural classifier. To better address this task, we propose Publisher Style Aggregation (PSA), a generalizable approach that aggregates publisher posting records to learn writing style and veracity stance. Extensive experiments demonstrate that our method outperforms existing baselines in terms of effectiveness, efficiency and generalizability.

Keywords: Rumor detection · Spurious correlations · Benchmarks · Text mining · Social network

1 Introduction

In the battle against escalating online misinformation, recent years have witnessed growing interest in automatic rumor detection on social media. Numerous real-world rumor detection datasets including TWITTER15 [13], TWITTER16 [13] and PHEME [15] have emerged as valuable resources fueling continuous development in this field. From early feature engineering models [1] to recent content-based [5,20] and propagation-based [6,29] methods, the ever-evolving approaches have achieved promising advances.

However, most existing methods ignore the *spurious attribute-label correlations* induced by dataset construction pitfalls, which arise from dataset-related artifacts instead of relationships generalizable to practical real-world settings. The

commonly adopted *event-based* data collection framework first fact-checks news-worthy events, and then automatically scrapes a large number of highly similar microblogs (e.g., tweets) containing the same event keywords. Some benchmark datasets also merge data samples from multiple existing sources to balance their class distribution. These factors consequently induce event-label and source-label correlations, which may not hold in practice.

Negligence of spurious cues can lead to unfair over-predictions that limit model generalization and adaptability. Similar issues have been identified in several natural language processing tasks including sentiment classification [17], argument reasoning comprehension [18] and fact verification [35], but the task of social media rumor detection remains underexplored.

Hence, in this paper, we make an effort to thoroughly investigate the causes of spurious correlations in existing rumor detection benchmark datasets, and take solid steps to counteract their impact. Specifically, we identify three causes of spurious correlations: (1) event-based data collection and labeling strategies associate event keywords with veracity labels; (2) merging data sources for label balancing creates correlations between source-specific propagation patterns and microblog veracity; and (3) event-level annotation strategies give rise to labeling bias. Under the post-level data splitting scheme commonly adopted by existing approaches, the prevalence of spurious cues can bring about numerous shared spurious correlations between the training and test data. For instance, the training and test data might even contain identical microblog texts, leading to data leakage (see examples in Fig. 1). If left unchecked, these correlations can lead to severe overestimation of model performance.

To offset the impact of spurious cues, we study a more practical task, namely *event-separated rumor detection*, where the test data contains microblogs from a set of events unseen during training. Without prior knowledge of event-specific cues in the test set, we empirically demonstrate stark performance deterioration of existing approaches, e.g. state-of-the-art rumor detection accuracy plummets from 90.2% to 44.3% on TWITTER16 [13], one of the most widely adopted datasets.

Striving for reliable rumor detection, we propose Publisher Style Aggregation (PSA), a novel method inspired by human fact-checking logic (i.e. reading through a user’s homepage to determine user stance and credibility). Specifically, our approach (1) encodes the textual features of source posts and user comments; (2) learns publisher-specific features based on multiple microblog instances produced by each source post publisher; and (3) augments each local microblog representation with its corresponding global publisher representation.

We evaluate PSA on the event-separated rumor detection task using 3 real-world benchmark datasets and compare it against 8 state-of-the-art baselines. Extensive experiments show that PSA outperforms its best competitors by a significant margin, respectively boosting test accuracy and F1 score by 14.18% and 15.26% on average across all 3 datasets. Furthermore, we empirically demonstrate the efficiency and generalizability of PSA via two experimental objectives, namely early rumor detection and cross-dataset rumor detection. Our code is publicly available at: <https://github.com/jiayingwu19/PSA>.

Dataset	Tweet Content	Freq.
TWITTER15	we are running out of chocolate, warns world's largest chocolate manufacturer URL URL	5
	lego letter from the 1970s still offers a powerful message to parents 40 years later URL URL	4
TWITTER16	steve jobs was adopted, his biological father was abdufhattah jandali, a syrian muslim.	12
	this is how people in china are riding escalators after a horrific accident URL URL	5

Steve Jobs was adopted. His biological father was Abdufhattah Jandali, a Syrian Muslim.
9:38 AM · Sep 28, 2015

Steve Jobs was adopted. His biological father was Abdufhattah Jandali, a Syrian Muslim.
9:37 AM · Oct 22, 2015

...

Steve Jobs was adopted. His biological father was Abdufhattah Jandali, a Syrian Muslim.
1:12 AM · Nov 25, 2015

Fig. 1. Automated event-based scraping results in numerous duplicate microblog texts in benchmark datasets TWITTER15 and TWITTER16, causing data leakage under random splitting. The highlighted words are event keywords obtained from the Snopes fact-checking website, in line with the datasets’ data collection scheme (see Section 3.1).

2 Related Work

Social Media Rumor Detection Real-world rumor detection datasets, with microblog posts and propagation patterns retrieved from social media platforms such as Twitter [12,13,15] and Weibo [4], form the bedrock of rapidly evolving approaches.

Recent advances in automated rumor detection typically adopt neural network based frameworks. *Content-based* approaches utilize microblog and comment features. For instance, [4,9] respectively employ Recurrent Neural Networks and Convolutional Neural Networks to model the variations of text and user representations over time. Hierarchical attention networks [10] and pre-trained language models [11] have also proven effective. Another line of work leverages *propagation-based* information diffusion patterns to encode information flow along user interaction edges. Some inject structural awareness into recursive neural networks [20] and multi-head attention [21,22,23], while others achieve success with Graph Neural Networks [6,40].

Closely related to our topic is fake news detection. [38] trains an event discriminator to overlook domain-specific knowledge under the multi-modal setting, [39] formulates domain-agnostic fake news detection as a continual learning problem, [37] studies the case with limited labeling budget, and [7,36] take advantage of auxiliary user descriptions and large-scale user corpus.

Existing methods either overlook the publisher-microblog relationship or require external knowledge (e.g. images and additional user description). In contrast, we seek to achieve generalizable rumor detection by capturing rumor-indicative publisher characteristics based on aggregation of multiple microblog data samples.

Investigation of Spurious Correlations Despite the promising performance of deep learning models, reliance on dataset-related cues has been observed in a wide range of tasks including text classification [17], natural language inference [32] and visual question answering [30]. In fact-checking scenarios, language models can capture underlying identities of news sites [19], and rumor instances can

possess time-sensitive characteristics [11]. Spurious artifacts lead to model failure on out-of-domain test instances, as empirically observed by [23,29,41].

However, systematic investigation into social media rumor detection remains unexplored. We bridge this gap by discussing three types of spurious correlations specific to this topic, and provide a solution to offset the impact of spurious correlations (i.e. event-separated rumor detection).

3 Spurious Correlations in Event-Based Datasets

3.1 Event-Based Data Collection

In this subsection, we outline the event-based data collection scheme adopted by benchmark datasets.

Newsworthy Event Selection: Newsworthy events serve as vital information sources, from which rumors and non-rumors arise and diffuse on social media. Existing studies either collect events from leading fact-checking websites (e.g., Snopes, Emergent, and PolitiFact) [4,12,14], or obtain candidate events identified by professionals [15].

Keyword-Based Microblog Retrieval: To facilitate mass collection, existing datasets typically adopt automated *event-based* data collection strategies, i.e. for each event, (1) extract keywords from its claim; (2) scrape microblogs via keyword-based search; and (3) select influential microblogs. Event keywords are mostly *neutral* (e.g. places, people or objects), carrying little or no stance.

Microblog Labeling Scheme: Existing rumor detection datasets conduct fact-checking at either event-level [4,12,14] or post-level [15]. While *event-level* labeling assigns all source posts under an event with the same event-level fact-checking label, *post-level* labeling annotates every source post independently. Although both event- and post-level annotations are performed by trained professionals, the former is more vulnerable to data selection pitfalls, on which we elaborate in Section 3.2.

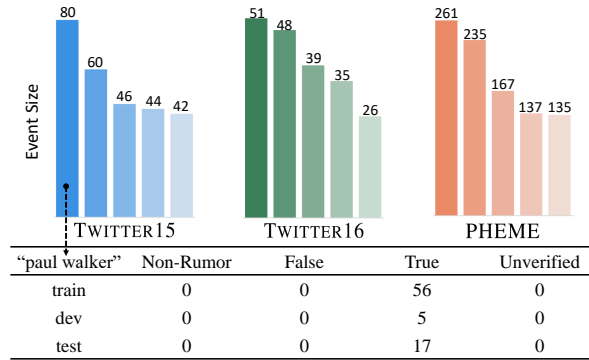
3.2 Possible Causes of Spurious Correlations

We investigate three of the most popular event-based rumor detection benchmark datasets containing source posts, propagation structures and conversation threads, namely TWITTER15 [13], TWITTER16 [13] and PHEME [15], and summarize the dataset statistics in Table 1. As TWITTER15 and TWITTER16 both consist of class-balanced tweets with abundant interactions, we also sample class-balanced tweets involving at least 10 users in PHEME for a fair comparison.

Intra-Event Textual Similarity Under each event, the automated keyword-based microblog retrieval framework collects a large number of highly similar keyword-sharing samples with the same label, even obtaining identical microblog texts (Fig. 1). Consequently, the correlations between event keywords and class labels result in strong textual cues that generalize poorly beyond the current event. Under the post-level data splitting scheme adopted by existing works,

Table 1. Dataset statistics.

Dataset	TWITTER15	TWITTER16	PHEME
Labeling Scheme	Event-Level	Event-Level	Post-Level
# of Events	298	182	9
# of Source Posts	1490	818	973
# of Non-Rumors	374	205	245
# of False Rumors	370	205	241
# of True Rumors	372	207	244
# of Unverified Rumors	374	201	243
# of Distinct Users	480,984	289,675	12,905
# of User Interactions	622,927	362,713	21,169

**Fig. 2.** The size of largest events in three datasets; “paul walker” directly correlates with “True” label in the post-level random splits adopted by the SOTA method [22].

these cues would scatter into different splits, creating shared correlations between the training and test data that do not hold in the real world. We illustrate such correlations via TWITTER15’s largest event about the death of Paul Walker (Fig. 2). All 80 microblogs reporting this event are assigned the “True” label, among which 78 contain the keywords “paul walker”. These event-specific keywords produce a strong correlation between “paul walker” and the “True” label. Under post-level random splitting, these samples spread across different data splits, creating undesirable textual similarity between the training and test data. As shown in Fig. 2, the datasets are dominated by such large-size events. Specifically, the top-5 largest events cover 96.09% of data samples in PHEME, while large-size events (containing more than 5 keyword-sharing tweets) cover more than 70% of samples in TWITTER15 and TWITTER16. Large event sizes lead to the prevalence of event-specific keyword-label correlations, further exacerbating the problem.

Merge of Data Sources For label balancing purposes, TWITTER15 and TWITTER16 merge tweets from multiple sources including [4,12,16], and scrape additional news events from verified media accounts. While the events covered by different sources do not overlap, direct 1-1 correlations between data sources and




Fig. 3. Source-label correlations in TWITTER16 suggest underlying spurious properties. IJCAI and PLOS_ONE refer to [4] and [16], respectively.



Fig. 4. One mislabeled instance from TWITTER15, where the source tweet, inconsistent with the Snopes claim, is wrongly assigned the same event-level label.

labels can possibly induce spurious correlations between data source features and the labels. As demonstrated in Fig. 3, the user interaction count (comments and reposts) and interaction time range of tweets from each source form distinctive patterns. For instance, all tweets from PLOS_ONE [16] are “True”, spread very quickly and tend to arouse less interactions. These source-specific propagation patterns could possibly be exploited by graph- or temporal-based models, which we empirically demonstrate in Section 6.5 (Table 4).

Labeling Bias While automated event-based data retrieval and event-level annotations allow for easier construction of large-scale benchmark datasets, the lack of post-level scrutiny induces vulnerability to labeling bias. For instance, as shown in Fig. 4, Snopes marked a MH17-related event claim as “False” due to image fabrication. However, in view of highly similar keywords “Malaysian Airlines” and “photo”, the data collection framework retrieved an MH17-relevant tweet linking to an authentic photo by Reuters and mistakenly labeled it as “False” in TWITTER15. Such mislabelings exacerbate our previously mentioned problem of intra-event textual similarity, making the resulting keyword-label correlations stronger but more deceptive. To make the best use of valuable data resources, we suggest that future approaches incorporate techniques that are robust to label noise (e.g. noise-tolerant training [34]).

4 Event-Separated Rumor Detection

4.1 Problem Formulation

Social media rumor detection aims to learn a classification model that is able to detect and fact-check rumors. Let $\mathcal{T} = \{T_1, T_2, \dots, T_N\}$ be a rumor detection dataset of size N , and $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$ be the corresponding ground-truth labels, with $y_i \in \mathcal{C} = \{1, \dots, C\}$. Each microblog instance T_i consists of a source post publisher u_i , a source post p_i , and related comments c_i^1, \dots, c_i^k . p_i (c_i^j) has a corresponding textual feature vector \mathbf{r}_i (\mathbf{r}_i^j). Denote the event behind microblog instance T_i as e_i . Consequently, training data \mathcal{T}_{tr} and test data \mathcal{T}_{te} respectively contain events $\mathcal{E}_{tr} = \{e_i | T_i \in \mathcal{T}_{tr}\}$ and $\mathcal{E}_{te} = \{e_i | T_i \in \mathcal{T}_{te}\}$.

Most existing approaches ignore the underlying microblog-event relationship and adopt event-mixed post-level data splits, resulting in significant overlap between \mathcal{E}_{tr} and \mathcal{E}_{te} . However, prior knowledge of test data is not always guaranteed in practice (e.g. the model’s performance gains from duplicate tweets in the training and test data are unlikely to generalize), and previous assumptions can lead to performance overestimation caused by intra-event textual similarity (see Section 3.2).

In order to eliminate these confounding event-specific correlations, we propose to study a more practical problem, namely **event-separated rumor detection**, where $\mathcal{E}_{tr} \cap \mathcal{E}_{te} = \emptyset$. This task is challenging due to the underlying event distribution shift, and it thereby provides a means to evaluate debiased rumor detection performance.

4.2 Existing Approaches

We compare the event-mixed and event-separated rumor detection performance of representative approaches on TWITTER15, TWITTER16 and PHEME datasets to investigate the impact of event-specific spurious correlations.

Propagation-Based: (1) **TD-RvNN** [20]: a recursive neural model that encodes long-distance user interactions with gated recurrent units. (2) **GLAN** [21]: a global-local attentive model based on microblog-user heterogeneous graphs. (3) **BiGCN** [6]: a GCN-based model that encodes augmented bi-directional rumor propagation patterns. (4) **SMAN** [22]: a multi-head attention model that enhances training with user credibility modeling.

Content-Based: (1) **BERT** [24]: a deep language model that encodes bidirectional context with Transformer blocks. (2) **XLNet** [27]: a generalized autoregressive approach that trains on all possible permutations of word factorizations. (3) **RoBERTa** [25]: a refined BERT-like approach that adopts dynamic masking for varied masking patterns over different epochs. (4) **DistilBERT** [26]: a distilled model that maintains 97% of BERT’s performance with 60% less parameters.

Data Splitting For all three datasets, we sample 10% instances for validation, then split the remainder 3:1 into training and test sets. Specifically, we obtain the event-separated splits based on the *publicly available event IDs* released in TWITTER15 [13], TWITTER16 [13] and PHEME [15], respectively.

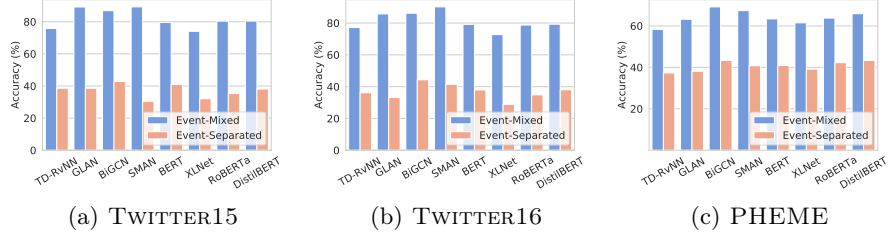


Fig. 5. Existing rumor detection approaches fail to generalize across events. Comparing between event-mixed and event-separated settings, mean accuracy based on 20 different runs of each approach demonstrates drastic performance deterioration. (Over 40% drop on TWITTER15 & 16, and over 20% on PHEME.)

4.3 SOTA Models’ Performance is Heavily Overestimated

Fig. 5 reveals a stark contrast between event-mixed and event-separated rumor detection performance. More specifically, test accuracy plummets from 74.0%~89.2% to 30.5%~42.8% on TWITTER15, from 72.7%~90.2% to 28.8%~44.3% on TWITTER16, and from 58.3%~69.2% to 37.3%~43.4% on PHEME. Furthermore, despite the consistency of best event-separated performance across all three datasets, all models achieve significantly higher event-mixed performance on TWITTER15 and TWITTER16 than on PHEME, where the former adopts event-level labeling and the latter adopts post-level labeling (see Section 3.1). This gap is in line with our hypothesis that direct event-label correlations induce additional bias.

Results imply the heavy reliance of existing methods on spurious event-specific correlations. Despite performing well under the event-mixed setting, these models cannot generalize to unseen events, resulting in poor real-world adaptability.

5 Proposed Method

To tackle the challenges of event-separated rumor detection, we propose Publisher Style Aggregation (PSA), a novel approach that learns generalizable publisher characteristics based on each publisher’s aggregated posts, as illustrated in Fig.6.

5.1 Consistency of Publisher Style

Source post publishers are highly influential users who produce claims towards newsworthy events. Therefore, each publisher’s unique credibility stance and writing style can exhibit distinctive traits that help determine the veracity of their statements. For a more intuitive view, we illustrate the TWITTER15 publisher tendency towards each class in Fig. 7. Specifically, for publisher u , we define u ’s tendency score under class c as $(\# \text{ microblogs posted by } u \text{ under class } c) / (\#$

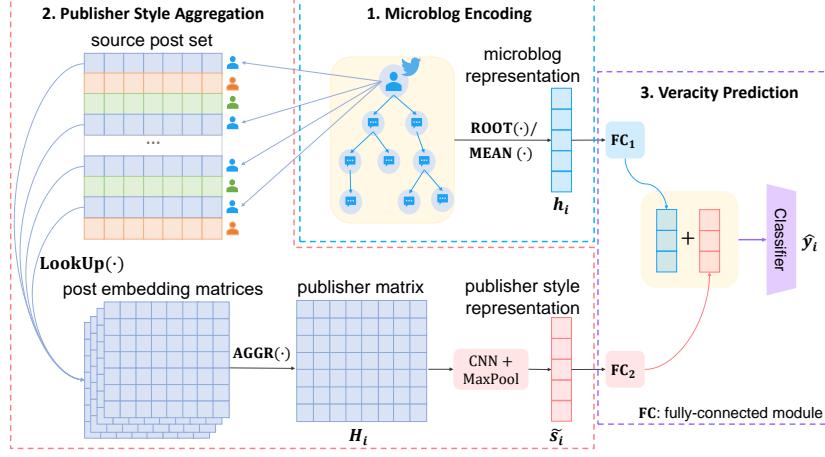


Fig. 6. Overview of our proposed PSA framework.

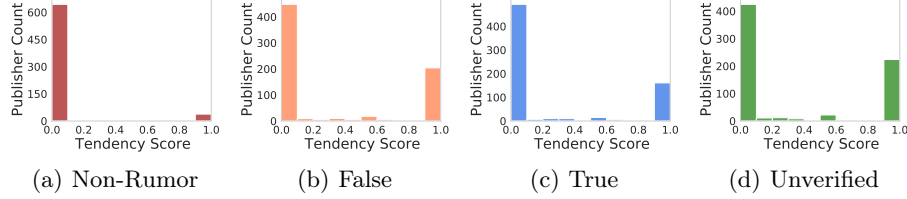


Fig. 7. Publishers tend to post tweets of the same credibility type, as exemplified by TWITTER15 publisher behavior patterns.

microblogs posted by u). Fig. 7 shows that most publishers have scores either approaching 0 or approaching 1 towards one particular class, i.e., most publishers tend to post microblogs under a single veracity label, which verifies our hypothesis of publisher style consistency.

5.2 Content-Based Microblog Encoding

We first propose two simple neural classifiers, namely RootText and MeanText, to study the importance of source post and comment features in social media rumor detection. In each dataset, all source posts and comments constitute a vocabulary of size $|V|$. Following [6,13], we formulate each source post feature $\mathbf{r}_i \in \mathbb{R}^{|V|}$ and its related comment features $\mathbf{r}_i^j \in \mathbb{R}^{|V|}$ as the sum of all one-hot word vectors within the corresponding source post or comment.

RootText: Source posts contain the claims to be fact-checked. Therefore, we propose to encode each microblog instance T_i solely based on its source post textual feature \mathbf{r}_i , i.e., $\mathbf{h}_i := \mathbf{r}_i$.

MeanText: We also propose to consider the user comments of source posts for more robust credibility measurement. Here, we adopt mean-pooling to condense source post and comment features into microblog representations:

$$\mathbf{h}_i := \frac{\mathbf{r}_i + \sum_{j=1}^k \mathbf{r}_i^j}{k+1}. \quad (1)$$

We obtain the encoding $\mathbf{h}_i \in \mathbb{R}^{|V|}$ of microblog T_i based on either RootText or MeanText, and extract high-level features $\tilde{\mathbf{h}}_i \in \mathbb{R}^n$ via a two-layer fully-connected neural network with the ReLU activation function. Then, we employ dropout to prevent overfitting before passing $\tilde{\mathbf{h}}_i$ through the final fully-connected layer with output dimensionality $|\mathcal{C}|$ for veracity prediction.

5.3 Publisher Style Aggregation

As shown in Section 5.1, the writing stance and credibility of highly influential source post publishers remain relatively stable in a fixed timeframe. Inspired by this observation, we further propose Publisher Style Aggregation (PSA), a generalizable method that jointly leverages multiple microblog instances produced by each publisher and extracts distinctive publisher traits to enhance local features learned within each individual microblog. More specifically, PSA (1) looks up a set of microblog instances produced by each publisher, (2) learns publisher style representations based on these source posts' aggregated textual features, and (3) augments the representation of each microblog (i.e. $\tilde{\mathbf{h}}_i$ learned via RootText / MeanText) with its corresponding publisher style representation $\tilde{\mathbf{s}}_i$.

Publisher Style Modeling. Assume that publisher u_i has produced $m_i \geq 1$ microblog instances, with the corresponding source posts denoted as $\mathcal{P}(u_i) = \{p_k | u_k = u_i, k = 1, \dots, N\}$; note that only accessible data are used during training. We treat the j -th post $p_i^j \in \mathcal{P}(u_i)$ as a word token sequence with maximum length L . Then, we construct an embedding matrix $\mathbf{W}_i^j \in \mathbb{R}^{L \times d}$ for p_i^j based on trainable d -dimensional word embeddings. We aggregate all post embedding matrices $\{\mathbf{W}_i^j\}_{j=1}^{m_i}$ of u_i , and obtain the corresponding publisher matrix $\mathbf{H}_i \in \mathbb{R}^{L \times d}$ as follows:

$$\mathbf{H}_i = \text{AGGR}(\{\mathbf{W}_i^j\}_{j=1}^{m_i}), \quad (2)$$

where the AGGR operator can be either MEAN or SUM.

To capture high-level publisher characteristics, we apply convolution on each \mathbf{H}_i to extract latent publisher style features. Specifically, we use three convolutional layers with different window sizes to learn features with varied granularity. Each layer consists of F filters, and each filter outputs a feature map $\mathbf{f}_* = [f_*^1, f_*^2, \dots, f_*^{L-k+1}]$, with

$$f_*^j = \text{ReLU}(\mathbf{W}_f \cdot \mathbf{H}_i[j : j+k-1] + b), \quad (3)$$

where $\mathbf{W}_f \in \mathbb{R}^{k \times d}$ the convolution kernel, k the window size and $b \in \mathbb{R}$ a bias term. We perform max-pooling to extract the most prominent value of each \mathbf{f}_* , and stack these values to form a style feature vector $\mathbf{s} \in \mathbb{R}^F$. Then, we concatenate

the \mathbf{s}_* produced by each of the three CNN layers to obtain the publisher style representation $\tilde{\mathbf{s}}_i \in \mathbb{R}^{3F}$:

$$\tilde{\mathbf{s}}_i = \text{Concat}[\mathbf{s}_1; \mathbf{s}_2; \mathbf{s}_3]. \quad (4)$$

Microblog Veracity Prediction. We augment microblog representation $\tilde{\mathbf{h}}_i \in \mathbb{R}^n$ with the corresponding publisher style representation $\tilde{\mathbf{s}}_i$. Finally, we utilize a fully connected layer to predict the microblog veracity label \hat{y}_i :

$$\hat{y}_i = \text{Softmax}(\mathbf{W}_2^T(\tilde{\mathbf{h}}_i + \mathbf{W}_1^T\tilde{\mathbf{s}}_i)), \quad (5)$$

where transformations $\mathbf{W}_1 \in \mathbb{R}^{3F \times n}$ and $\mathbf{W}_2 \in \mathbb{R}^{n \times |\mathcal{C}|}$. We also apply dropout before the final layer to prevent overfitting.

Model parameters are optimized by minimising the cross-entropy loss between \hat{y}_i and ground truth y_i .

6 Experiments

In this section, we review our experiments for answering the following questions:

Q1 (Model Performance): Does PSA outperform the existing baselines on event-separated rumor detection?

Q2 (Early Rumor Detection): Does PSA work well under temporal rumor detection deadlines?

Q3 (Model Generalization): Is PSA effective under cross-dataset settings?

6.1 Experimental Setup

We implement our proposed PSA model and its variants based on PyTorch 1.6.0 with CUDA 10.2, and train them on a server running Ubuntu 18.04 with NVIDIA RTX 2080Ti GPU and Intel(R) Xeon(R) CPU E5-2690 v4 @ 2.60GHz. We adopt an Adam optimizer with $(\beta_1, \beta_2) = (0.9, 0.999)$, learning rate of 10^{-4} (0.005), and weight decay 10^{-5} (10^{-4}) for TWITTER15/16 (PHEME). We obtain microblog representations via a 2-layer neural network with layer sizes of 128 and 64. We utilize the 300-dimensional word vectors from [21] to form publisher matrices, employ three CNN layers with the same filter number $F = 100$ but different window sizes $k \in \{3, 4, 5\}$, concatenate their outputs and use a fully-connected layer to extract publisher style representations with the size of 64. We implement AGGR in Eq. 2 as both SUM and MEAN, and report average performance over 20 different runs on the event-separated data splits presented in Section 4.2.

6.2 Q1. Model Performance

We compare PSA (base classifier: RootText / MeanText) with existing approaches in Table 2.

Importance of Textual Features: We observe that MeanText outperforms existing methods on TWITTER15&16, while RootText only achieves 0.6% lower accuracy than the best baseline on PHEME. This implies severe degradation

Table 2. PSA significantly improves event-separated rumor detection accuracy (%) and Macro F1 Score (%) (S stands for SUM and M for MEAN; averaged over 20 runs).

Method	TWITTER15		TWITTER16		PHEME	
	Acc.	F1	Acc.	F1	Acc.	F1
TD-RvNN [20]	38.62 \pm 1.85	36.40 \pm 2.38	36.15 \pm 1.90	35.66 \pm 1.89	37.30 \pm 2.54	34.17 \pm 2.56
GLAN [21]	38.56 \pm 3.38	35.52 \pm 5.31	33.13 \pm 4.54	27.93 \pm 5.53	38.10 \pm 2.85	34.60 \pm 3.04
BiGCN [6]	42.83 \pm 2.27	38.17 \pm 3.04	44.28 \pm 3.39	42.31 \pm 3.77	43.36 \pm 1.71	37.93 \pm 2.16
SMAN [22]	30.52 \pm 2.62	28.80 \pm 3.30	41.42 \pm 2.65	40.62 \pm 2.95	40.74 \pm 1.36	36.02 \pm 1.62
BERT [24]	40.95 \pm 4.80	37.47 \pm 7.56	37.89 \pm 6.68	34.76 \pm 9.22	40.90 \pm 3.22	36.33 \pm 4.02
XLNet [27]	32.05 \pm 6.78	26.00 \pm 9.20	28.82 \pm 4.08	20.59 \pm 7.06	39.14 \pm 5.14	34.35 \pm 6.65
RoBERTa [25]	35.30 \pm 5.65	28.41 \pm 7.93	34.84 \pm 6.69	27.90 \pm 9.74	42.25 \pm 4.47	37.67 \pm 5.70
DistilBERT [26]	38.09 \pm 4.48	33.74 \pm 5.21	38.02 \pm 4.24	33.98 \pm 5.56	43.33 \pm 3.62	38.37 \pm 4.16
RootText(RT)	33.80 \pm 2.74	30.56 \pm 2.92	30.54 \pm 1.72	28.87 \pm 2.32	42.75 \pm 1.25	38.74 \pm 1.47
MeanText(MT)	48.68 \pm 1.80	47.18 \pm 1.63	45.19 \pm 1.86	44.23 \pm 1.72	32.57 \pm 1.70	30.48 \pm 1.51
RT+PSA(S)	47.85 \pm 5.64	45.26 \pm 5.13	57.88 \pm 3.16	55.30 \pm 4.64	43.52 \pm 0.93	37.81 \pm 1.22
RT+PSA(M)	45.67 \pm 0.82	38.55 \pm 0.89	47.28 \pm 2.87	42.74 \pm 4.15	46.30 \pm 1.28	41.57 \pm 1.49
MT+PSA(S)	61.83 \pm 1.43	58.75 \pm 2.08	64.89 \pm 1.75	64.31 \pm 1.70	37.43 \pm 1.05	32.36 \pm 1.25
MT+PSA(M)	54.73 \pm 1.04	50.85 \pm 1.56	60.16 \pm 2.76	58.16 \pm 3.13	40.19 \pm 1.20	35.98 \pm 1.45

of overparameterized models when the spurious attribute-label correlations (i.e. event-specific cues) in the training data do not apply to the test data, in line with prior work [41]. Comparing between RootText and MeanText, we also observe that the former performs better on PHEME but otherwise on TWITTER15 and TWITTER16. Different labeling schemes may account for such differences; as PHEME labels each microblog independently, the source posts would contain the most distinctive features. While the source post content is not as distinctive in TWITTER15 and TWITTER16, both datasets exhibit more complicated propagation patterns (see Table 1). Therefore, adopting MeanText to aggregate comment features proves more effective in these cases.

Effectiveness of PSA: Our proposed PSA approach, with AGGR implemented as either SUM or MEAN, significantly enhances the base classifiers RootText and MeanText. The best PSA combinations outperform the best baseline by a large margin; they boost event-separated rumor detection accuracy by 19.00% on TWITTER15, 20.61% on TWITTER16, and 2.94% on PHEME. Unlike existing methods, PSA explicitly aggregates publisher style features across microblog instances from multiple events, thereby enhancing the model’s capability to learn event-invariant features. As a result, PSA is able to capture stance and style pertaining to distinctive publisher characteristics, leading to substantial performance improvements.

6.3 Q2. Early Rumor Detection

Accurate and timely misinformation detection is of vital importance. Given only partial propagation information, we compare the best PSA combinations and their

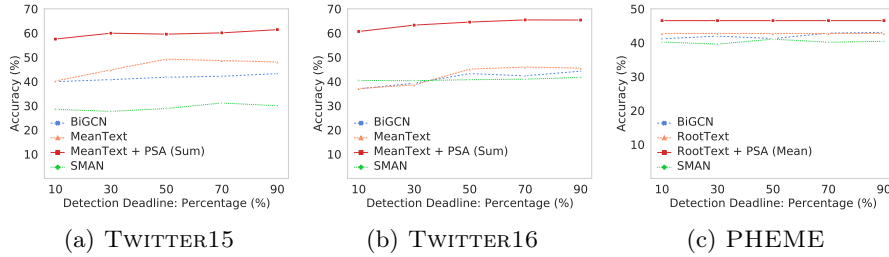


Fig. 8. Early rumor detection accuracy (%) of PSA models against the best propagation-based baselines for varying detection deadlines on TWITTER15, TWITTER16 and PHEME under the event-separated setting (averaged over 20 runs).

Table 3. Cross-dataset rumor detection accuracy (%) and Macro F1 Score (%) of PSA between T15 (TWITTER15) and T16 (TWITTER16), compared with the best propagation- and content-based baseline methods (averaged over 20 runs).

Train	Test	Method	Acc.	F1
T15	T16	DistilBERT	40.70 \pm 4.78	39.98 \pm 5.02
		BiGCN	36.80 \pm 3.79	35.76 \pm 3.92
		MeanText	46.55 \pm 1.33	44.11 \pm 1.71
		+ PSA (Sum)	63.99\pm1.53	61.95\pm1.86
T16	T15	DistilBERT	41.66 \pm 4.95	37.07 \pm 6.06
		BiGCN	44.39 \pm 1.89	42.41 \pm 2.23
		MeanText	48.00 \pm 1.70	44.83 \pm 1.63
		+ PSA (Sum)	60.82\pm1.47	57.97\pm2.07

corresponding base classifiers with their best propagation-based competitors. Fig. 8 shows the consistent superiority of PSA over baseline methods at all detection deadlines. Even with only the earliest 10% of comments, PSA achieves 57.53% accuracy on TWITTER15, 60.65% on TWITTER16, and 46.30% on PHEME. Note that RootText (+PSA) models maintain stable performance across all deadlines, as they provide instant predictions solely based on source posts. The results demonstrate that augmenting rumor detection models with publisher style representations achieves both efficiency and effectiveness.

6.4 Q3. Cross-Dataset Rumor Detection

To study the generalization ability of PSA, we conduct cross-dataset experiments on TWITTER15 and TWITTER16, where the model is trained on one dataset and tested on the other. For a fair comparison, we utilize the same event-separated data splits adopted in Sections 6.2 and 6.3. If overlapping events exist between the training set from dataset A and the test set from dataset B, we remove all instances related to these events in the training set, and replace them with the same number of non-overlapping instances randomly sampled from A’s test set.

Table 4. Classification accuracy (%) based on user interaction count and interaction time range suggests potential source-specific correlations.

Feature	TWITTER15	TWITTER16	PHEME
A: Interaction Count	43.03	43.32	23.15
B: Time Range	37.09	44.92	27.31
A + B	44.21	55.61	26.85

The cross-dataset setting is inherently more challenging, as the training and test events stem from different timeframes, which can create temporal concept shifts. However, Table 3 shows that PSA continues to excel and enhances the base classifier (our MeanText method) by 17.44% on TWITTER15 and 12.82% on TWITTER16, which further demonstrates PSA’s generality to unseen events.

6.5 Discussion: Source-Specific Spurious Cues

In Table 4, we empirically show the potential impact of source-specific spurious cues (Section 3.2). Under the event-separated setting, we construct a simple Random Forest classifier based on each microblog’s user interaction count and the time range covered by these interactions. Surprisingly, the classifier outperforms existing baseline methods on both TWITTER15 and TWITTER16, and achieves comparable performance even with only one feature. In contrast, the single-source PHEME remains unaffected. Although neither our proposed approaches nor existing methods exploit these features, we nevertheless suggest the integration of debiasing techniques in future graph- and temporal-based models.

7 Conclusion

In this paper, we systematically analyze how event-based data collection schemes create event- and source-specific spurious correlations in social media rumor detection benchmark datasets. We study the task of event-separated rumor detection to remove event-specific correlations, and empirically demonstrate severe limitations on existing methods’ generalization ability. To better address this task, we propose PSA to augment microblog representations with aggregated publisher style features. Extensive experiments on three real-world datasets show substantial improvement on cross-event, cross-dataset and early rumor detection.

For future work, we suggest (1) event-separated rumor detection performance as a major evaluation metric; (2) same-source samples and post-level expert annotations in dataset construction; and (3) integration of causal reasoning and robust learning techniques in model design, in the hope that our findings could motivate and measure further progress in this field.

Acknowledgements This work was supported in part by NUS ODPRT Grant R252-000-A81-133.

References

1. Castillo, C., Mendoza, M., Poblete, B.: Information Credibility on Twitter. In: WWW (2011)
2. Ma, J., Gao, W., Wei, Z., Lu, Y., and Wong, K.-F.: Detect Rumors Using Time Series of Social Context Information on Microblogging Websites. In: CIKM (2015)
3. Kwon, S., Cha, M., Jung, K., Chen, W., Wang, Y.: Prominent Features of Rumor Propagation in Online Social Media. In: ICDM (2013)
4. 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: IJCAI (2016)
5. Zhang, J., Dong, B., Yu, P.S.: FakeDetector: Effective Fake News Detection with Deep Diffusive Neural Network. In: ICDE (2020)
6. Bian, T., Xiao, X., Xu, T., Zhao, P., Huang, W., Rong, Y., Huang, J.: Rumor Detection on Social Media with Bi-Directional Graph Convolutional Networks. In: AAAI (2020)
7. Nguyen, V.-H., Sugiyama, K., Nakov, P., Kan, M.-Y.: FANG: Leveraging Social Context for Fake News Detection Using Graph Representation. In: CIKM (2020)
8. Shu, K., Wang, S., Liu, H.: Beyond News Contents: The Role of Social Context for Fake News Detection. In: WSDM (2019)
9. Liu, Y., Wu, Y.-F.: Early Detection of Fake News on Social Media Through Propagation Path Classification with Recurrent and Convolutional Networks. In: AAAI (2018)
10. Shu, K., Cui, L., Wang, S., Lee, D., Liu, H.: DEFEND: Explainable Fake News Detection. In: KDD (2019)
11. Pelrine, K., Danovitch, J., Rabbany, R.: The Surprising Performance of Simple Baselines for Misinformation Detection. In: WWW (2021)
12. Liu, X., Nourbakhsh, A., Li, Q., Fang, R., Shah, S.: Real-Time Rumor Debunking on Twitter. In: CIKM (2015)
13. Ma, J., Gao, W., Wong, K.-F.: Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning. In: ACL (2017)
14. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: FakeNewsNet: A Data Repository with News Content, Social Context and Dynamic Information for Studying Fake News on Social Media. arXiv preprint arXiv:1809.01286 (2018)
15. Kochkina, E., Liakata, M., Zubiaga, A.: All-in-one: Multi-task Learning for Rumour Verification. In: COLING (2018)
16. Zubiaga, A., Liakata, M., Procter, R., Wong, S.H.G., Tolmie, P.: Analysing How People Orient to and Spread Rumours in Social Media by Looking at Conversational Threads. PLOS ONE. Public Library of Science (2016)
17. Wang, Z., Culotta, A.: Identifying Spurious Correlations for Robust Text Classification. In: Findings of EMNLP (2020)
18. Niven, T., Kao, H.-Y.: Probing Neural Network Comprehension of Natural Language Arguments. In: ACL (2019)
19. Zhou, X., Elfardy, H., Christodoulopoulos, C., Butler, T., Bansal, M.: Hidden Biases in Unreliable News Detection Datasets. In: EACL (2021)
20. Ma, J., Gao, W., Wong, K.-F.: Rumor Detection on Twitter with Tree-structured Recursive Neural Networks. In: ACL (2018)
21. Yuan, C., Ma, Q., Zhou, W., Han, J., Hu, S.: Jointly Embedding the Local and Global Relations of Heterogeneous Graph for Rumor Detection. In: ICDM (2019)
22. Yuan, C., Ma, Q., Zhou, W., Han, J., Hu, S.: Early Detection of Fake News by Utilizing the Credibility of News, Publishers, and Users based on Weakly Supervised Learning. In: COLING (2020)

23. Khoo, L.M.S., Chieu, H.L., Qian, Z., Jiang, J.: Interpretable Rumor Detection in Microblogs by Attending to User Interactions. In: AAAI (2020)
24. Devlin, J., Chang, M.-W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: NAACL (2019)
25. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv preprint arXiv: 1907.11692 (2019)
26. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv preprint arXiv: 1910.01108 (2019)
27. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: XLNet: Generalized Autoregressive Pretraining for Language Understanding. Advances in Neural Information Processing Systems (2019)
28. Wolf, T., Debut, L., et al.: Transformers: State-of-the-Art Natural Language Processing. In: EMNLP (2020)
29. Huang, Y.-H., Liu, T.-W., Lee, S.-R., Alvarado, C., Henrique, F., Chen, Y.-S.: Conquering Cross-Source Failure for News Credibility: Learning Generalizable Representations beyond Content Embedding. In: WWW (2020)
30. Agarwal, V., Shetty, R., Fritz, M.: Towards Causal VQA: Revealing and Reducing Spurious Correlations by Invariant and Covariant Semantic Editing. In: CVPR (2020)
31. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel W.: ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: ICLR (2019)
32. McCoy, T., Pavlick, E., Linzen, T.: Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. In: ACL (2019)
33. Srivastava, M., Hashimoto, T.B., Liang, P.: Robustness to Spurious Correlations via Human Annotations. In: ICML (2020)
34. Li, J., Wong, Y., Zhao, Q., Kankanhalli, M.S.: Learning to Learn From Noisy Labeled Data. In: CVPR (2019)
35. Schuster, T., Shah, D., Yeo, Y.J.S., Filizzola, D., Santus, E., Barzilay, R.: Towards Debiasing Fact Verification Models. In: EMNLP-IJCNLP (2019)
36. Dou, Y., Shu, K., Xia, C., Yu, P.S. Sun, L.: User Preference-Aware Fake News Detection. In: SIGIR (2021)
37. Silva, A., Luo, L., Karunasekera, S., Leckie, C.: Embracing Domain Differences in Fake News: Cross-domain Fake News Detection using Multi-modal Data. In: AAAI (2021)
38. Wang, Y., Ma, F., Jin, Z., Yuan, Y., Xun, G., Jha, K., Su, L., Gao, J.: EANN: Event Adversarial Neural Networks for Multi-Modal Fake News Detection. In: KDD (2018)
39. Han, Y., Karunasekera, S., Leckie, C.: Graph Neural Networks with Continual Learning for Fake News Detection from Social Media. arXiv preprint arXiv: 2007.03316 (2020)
40. Huang, Q., Yu, J., Wu, J., Wang, B.: Heterogeneous Graph Attention Networks for Early Detection of Rumors on Twitter. arXiv preprint arXiv:2006.05866 (2020)
41. Sagawa, S., Raghunathan, A., Koh, P. W., Liang, P.: An Investigation of Why Overparameterization Exacerbates Spurious Correlations. In: ICML (2020)