

It’s about Time: Rethinking Evaluation on Rumor Detection Benchmarks using Chronological Splits

Yida Mu, Kalina Bontcheva, Nikolaos Aletras

Department of Computer Science, The University of Sheffield
{y.mu, k.bontcheva, n.aletras}@sheffield.ac.uk

Abstract

New events emerge over time influencing the topics of rumors in social media. Current rumor detection benchmarks use random splits as training, development and test sets which typically results in topical overlaps. Consequently, models trained on random splits may not perform well on rumor classification on previously unseen topics due to the temporal concept drift. In this paper, we provide a re-evaluation of classification models on four popular rumor detection benchmarks considering chronological instead of random splits. Our experimental results show that the use of random splits can significantly overestimate predictive performance across all datasets and models. Therefore, we suggest that rumor detection models should always be evaluated using chronological splits for minimizing topical overlaps.

1 Introduction

Unverified false rumors can spread faster than news from mainstream media, and often can disrupt the democratic process and increase hate speech (Vosoughi et al., 2018; Zubiaga et al., 2018). Automatic detection of rumors is an important task in computational social science, as it helps prevent the spread of false rumors at an early stage (Ma et al., 2017; Zhou et al., 2019; Karmakharm et al., 2019; Bian et al., 2020).

Current rumor detection approaches typically rely on existing annotated benchmarks consisting of social media data, e.g., Twitter 15 (Ma et al., 2017), Twitter 16 (Ma et al., 2017), Weibo (Ma et al., 2016), and PHEME (Zubiaga et al., 2016) that cover a wide range of time periods. These benchmarks use random splits for train, development and test sets which entail some topical overlap among them (see Table 1 for recent previous work). However, the distribution of topics in various NLP benchmarks (e.g., news, reviews, and biomedical) can be significantly affected by time (Huang and

Paper	Twitter 15	Twitter 16	PHEME	Weibo
Tian et al. (2022)	✓	✓	-	✓
Zeng and Gao (2022)	-	✓	✓	-
Sheng et al. (2022)	-	-	-	✓
Mukherjee et al. (2022)	-	-	✓	-
Sun et al. (2022)	✓	✓	✓	-
de Silva and Dou (2021)	✓	✓	-	-
Ren et al. (2021)	-	-	✓	-
Wei et al. (2021)	✓	✓	✓	-
Li et al. (2021)	-	-	✓	-
Rao et al. (2021)	✓	✓	-	✓
Lin et al. (2021)	✓	✓	✓	-
Farinneya et al. (2021)	-	-	✓	-
Sun et al. (2021)	-	-	✓	-
Qian et al. (2021)	-	-	✓	-
Song et al. (2021)	✓	✓	✓	-
Kochkina and Liakata (2020)	✓	✓	✓	-
Yu et al. (2020)	-	-	✓	-
Xia et al. (2020)	-	✓	-	✓
Bian et al. (2020)	✓	✓	-	✓
Lu and Li (2020)	✓	✓	-	-

Table 1: Recent work on rumor detection using random splits.

Paul, 2018, 2019). This is the phenomenon of temporal concept drift which can be induced by the changes in real-world events. Specifically, this also affects benchmarks on social media with new events such as elections, emergencies, pandemics, constantly creating new topics for discussion.

Gorman and Bedrick (2019) and Søggaard et al. (2021) have showed that using different data split strategies affects model performance in NLP downstream tasks. Previous work has demonstrated that text classifiers performance significantly drops in settings where chronological data splits are used instead of random splits in various domains, e.g., hate speech, legal, politics, sentiment analysis, and biomedical (Huang and Paul, 2018; Lukes and Søggaard, 2018; Huang and Paul, 2019; Florio et al., 2020; Chalkidis and Søggaard, 2022; Agarwal and Nenkova, 2022; Zhao et al., 2022). To minimize topical overlaps, a Leave-One-Out (LOO) evaluation protocol has been proposed (Lukasik et al., 2015, 2016). While this topic split strategy could potentially mitigate temporal concept drift, it still yields temporal overlaps between each subset and is practically not applicable to most common ru-

Dataset	id	Post	Label	Leven
Twitter 15	407231*	r.i.p to the driver who died with paul walker that no one cares about because he wasn't famous.	Rumor	3
	407236*	r.i.p to the driver that died with paul walker that no one cares about because he wasn't famous.	Rumor	
Twitter 16	594687*	the kissing islands, greenland. URL	Non-Rumor	0
	604628*	the kissing islands, greenland. URL	Non-Rumor	
PHEME	498483*	happening now in #ferguson URL	Non-Rumor	9
	499402*	Right now in #ferguson URL	Non-Rumor	
Weibo	349863*	【喝易拉罐一定要吸管】一妇女喝了罐饮料，被送进医院，离开了世界。研究显示罐上面的毒菌很多 请转给你关心的朋友 。 Translation: Please forward to your friends you care about.	Rumor	10
	350023*	【喝易拉罐一定要吸管】一妇女喝了罐饮料，被送进医院，离开了世界。研究显示罐上面的毒菌很多!! 这些你知道么 Translation: Do you know about this?	Rumor	

Table 2: Four pairs of posts from train and test data with similar or identical text content sampled from four rumor detection benchmarks. Post ids with close values indicate that two posts are published in the same period. **Leven** denotes the Levenshtein distance (Levenshtein et al., 1966) on character-level between the two posts with the same label (i.e., lower values indicate higher text similarity and vice versa).

rumor detection benchmarks with a large number of topics (e.g., Twitter 15, Twitter 16, Weibo, etc.). We observe that the LOO protocol can be used for a few specific rumor detection benchmarks, such as (PHEME (Zubiaga et al., 2016)), where each post is associated with a corresponding event, e.g. *Ottawa Shooting* and *Charlie Hebdo shooting*.

Using random splits also results into posts with almost identical textual content shared during the same period. Table 2 displays four pairs of posts with **similar or identical** text content sampled from four different rumor detection benchmarks. This potential information leakage, results in classifying data almost identical to ones already being present in the training set. For practical application reasons, we believe that in order to evaluate a rumor detection system, it is necessary to detect not only long-standing rumors, but also emerging ones.

In this paper, we design a battery of controlled experiments to explore the hypothesis that whether temporality affects the predictive performance of rumor classifiers. To this end, we re-evaluate models on popular rumor detection benchmarks using chronological data splits i.e., by training the model with earlier posts and evaluating the model performance with the latest posts. Results show that the performance of rumor detection approaches trained with random data splits is significantly overestimated than chronological splits due to temporal concept drift. This suggests that rumor detection approaches should be evaluated with chronological data for real-world applications, i.e., to automatically detect emerging rumors.

2 Methodology

2.1 Data

We use four most popular rumor detection benchmarks, three in English and one in Chinese. Note

that most related work is currently evaluating their rumor detection systems on two or three of these four benchmarks. (see Table 1).

Twitter 15 and Twitter 16: These datasets contain 1,490 and 818 tweets labeled into four categories including Non-rumor (NR), False Rumor (FR), True Rumor (TR), and Unverified Rumor (UR) introduced by Ma et al. (2017).

PHEME: This benchmark contains 5,802 verified tweets collected from 9 real-world breaking news events (e.g., Ottawa Shooting, Ferguson Unrest, etc.) associated with two labels, i.e., 1,972 Rumor and 3,830 Non-Rumor (Zubiaga et al., 2016).

Weibo: This dataset includes 4,664 verified posts in Chinese including 2,313 rumors debunked by the Weibo Rumor Debunk Platform¹ and 2,351 non-Rumors from Chinese media (Ma et al., 2016).

Data Pre-processing We opt for the binary setup (i.e., re-frame all benchmarks as rumor detection) to distinguish true/false information following Lu and Li (2020); Rao et al. (2021). We pre-process the posts by replacing @mention and hyperlinks with @USER and URL respectively. We also lowercase the tweets from three Twitter benchmarks.

2.2 Data Splits

Standard Chronological Splits For Twitter 15 and PHEME, we first sort all posts chronologically and then divide them into three subsets including a training set (70% of the earliest data), a development set (10% of data after train and before test), and a test set (20% of the latest data). There is no temporal overlap between the three subsets.

¹<https://service.account.weibo.com/?type=5&status=4>

Splits	Benchmarks Subsets	Twitter 15			Twitter 16			PHEME			Weibo		
		Train	Dev	Test	Train	Dev	Test	Train	Dev	Test	Train	Dev	Test
Standard Chronological	# of Rumors	285	35	52	-	-	-	1,420	72	480	-	-	-
	# of Non-Rumors	234	40	96	-	-	-	2,641	508	681	-	-	-
Stratified Chronological	# of Rumors	260	37	75	144	21	40	1,380	197	394	1,645	235	470
	# of Non-Rumors	259	37	74	144	21	40	2,681	383	766	1,619	231	463
Random Splits	# of Rumors	260	37	75	144	21	40	1,380	197	394	1,645	235	470
	# of Non-Rumors	259	37	74	144	21	40	2,681	383	766	1,619	231	463

Table 3: Statistics of subsets. Note that using random splitting yields the same percentage of examples in each category as in the stratified chronological splits.

Stratified Chronological Splits On the other hand, we observe that there is no temporal overlap between rumors and non-rumors in Twitter 16 and Weibo datasets. This suggests that it is not possible to use standard chronological splits as in Twitter 15 and PHEME.

Therefore, we apply a **stratified chronological split** strategy for all benchmarks. We first split rumors and non-rumors separately in chronological order. We then divide them into three subsets (a total of six subsets), i.g., all rumors are split into a training set (70% of the earliest rumors), a development set (10% of data after train and before test), and a test set (20% of the latest rumors). Finally, we merge the six subsets into the final three train, development and test sets. Note that this approach will result in no temporal overlap for **each label (i.e., rumor or non-rumor)** among the three final sets. We show the number of each split in Table 3.

Random Splits Following standard practice (e.g., Bian et al. 2020; Lin et al. 2021; Rao et al. 2021), we **randomly** split data using a 5-fold cross-validation. Note that these splits are made by preserving the percentage of posts in each category. Each split contains a training set (70%), development set (10%) and a test set (20%) with the same ratio as in our chronological splits.

Leave-One-Out (LOO) Splits For reference, we also provide the results of using the LOO evaluation protocol on PHEME dataset (see Table 5).

2.3 Models

The main purpose of our experiments is to improve model evaluation by investigating the effects of temporal drifts in rumor detection by providing an extensive empirical study. Therefore, we opted using strong text classifiers that are generic and can be applied to all of our benchmarks:

- **LR** We train a LR classifier using BOW to represent posts weighted by TF-IDF using a vocabulary of 5,000 n-grams.

- **BERT** We directly fine-tune the BERT base model by adding a linear prediction layer on the top of the 12-layer transformer architecture following (Devlin et al., 2019).
- **BERT+ (BERTweet and ERNIE)** We also experiment with two domain specific models: BERTweet (Nguyen et al., 2020) and ERNIE (Sun et al., 2020) pre-trained on social media data using the same fine-tune strategy as the original BERT model.

2.4 Hyperparameters and Implementation Details

We train the model on the training set, perform model tuning and selecting on the development set, and evaluate performance on the test set. To evaluate the chronological data splits, we run the model five times with different random seeds for consistency. All chronological splits are available for reproducibility.²

For logistic regression, we use word-level and character-level tokenizers for Twitter and Weibo datasets respectively and only consider uni-gram, bi-grams, and tri-grams that appear in more than two posts for each dataset. For BERT, we set learning rate $lr = 2e - 5$, batch size $bs = 32$, and maximum input length as 256 covering the max tokens of all posts. All BERT-style models are trained for 10 epochs using the early stopping method based on the loss on the development set. The best checkpoint model is saved for evaluation on the test set. The average run time of 10 epochs for the BERT model is less than 2 minutes. We employ Bert-Base-Uncased, Bertweet-Base and Chinese-Bert-WWM, Ernie-1.0 models from the HuggingFace library (Wolf et al., 2020). All experiments are conducted on a single NVIDIA V100 GPU with 32GB memory.

²https://github.com/YIDAMU/Rumor_Benchmarks_Temporality

Model	Strategy	Twitter15			PHEME		
		P	R	F1	P	R	F1
LR	Random	86.7 ± 2.1	85.2 ± 1.8	85.0 ± 1.8	84.1 ± 1.2	79.3 ± 1.0	80.9 ± 1.0
	Standard Chronological	56.6 ± 0.8	56.3 ± 0.7	56.4 ± 0.7	67.3 ± 0.1	64.0 ± 0.1	63.9 ± 0.1
	Stratified Chronological	56.3 ± 2.5	51.9 ± 0.7	41.4 ± 0.4	64.5 ± 0.2	63.0 ± 0.3	63.5 ± 0.3
BERT	Random	88.2 ± 2.4	87.9 ± 2.2	87.9 ± 2.2	84.8 ± 0.5	84.8 ± 1.2	84.8 ± 0.8
	Standard Chronological	54.8 ± 4.0	55.1 ± 4.3	52.9 ± 3.6	74.8 ± 1.1	75.1 ± 0.8	73.7 ± 0.4
	Stratified Chronological	58.2 ± 7.3	56.1 ± 4.5	52.8 ± 5.6	75.5 ± 0.6	77.7 ± 0.5	75.7 ± 1.1
BERT+	Random	90.8 ± 1.2	90.4 ± 1.2	90.4 ± 1.2	84.6 ± 1.0	85.5 ± 0.9	85.0 ± 0.8
	Standard Chronological	58.6 ± 1.9	58.8 ± 2.1	57.4 ± 2.5	76.1 ± 1.1	74.8 ± 1.5	71.6 ± 2.2
	Stratified Chronological	61.8 ± 6.5	57.9 ± 2.4	55.2 ± 1.5	75.3 ± 0.9	76.9 ± 2.1	71.0 ± 3.5
Model	Strategy	Twitter16			Weibo		
LR	Random	89.9 ± 1.2	89.3 ± 1.5	89.3 ± 1.5	90.1 ± 0.9	90.1 ± 0.9	90.1 ± 0.9
	Stratified Chronological	62.1 ± 6.9	55.8 ± 4.7	48.7 ± 11.4	79.1 ± 0.1	78.1 ± 0.1	77.9 ± 0.1
	BERT	Random	91.9 ± 1.0	91.5 ± 0.8	91.5 ± 0.8	92.3 ± 1.2	92.2 ± 1.2
BERT	Stratified Chronological	61.0 ± 11.2	54.3 ± 4.3	47.2 ± 3.5	89.0 ± 2.5	87.6 ± 2.6	87.5 ± 2.6
	BERT+	Random	89.8 ± 2.8	89.3 ± 3.2	89.3 ± 3.3	92.5 ± .4	92.5 ± .4
BERT+	Stratified Chronological	49.8 ± 1.7	49.9 ± 0.9	45.1 ± 2.9	88.1 ± 2.5	87.6 ± 1.4	88.5 ± 1.5

Table 4: Rumor detection prediction results across different data split methods. Green cells indicate that the model trained on random splits performs significantly better than both standard chronological splits and stratified chronological splits ($p < 0.05$, t-test).

Model	PHEME		
	P	R	F1
LR	68.3 ± 3.8	65.1 ± 6.3	63.2 ± 6.3
BERT	73.4 ± 3.1	71.9 ± 6.1	70.7 ± 4.9
BERT+	75.3 ± 2.2	72.6 ± 8.1	71.4 ± 7.0

Table 5: Leave-One-Out evaluation protocol on PHEME dataset.

2.5 Evaluation Metrics

For all tasks, we report the averaged macro Precision, Recall and F1 values across five runs using different random seeds.

3 Results

Random Splits vs. Chronological Splits Table 4 shows the experimental results across all models and rumor detection benchmarks using **chronological splits** and random **5-fold cross-validation**. Overall, we observe that the use of random splits always leads to a significant overestimation of performance compared to chronological splits (t-test, $p < 0.05$) across all models. Our results corroborate findings from previous work on studying temporal concept drift (Huang and Paul, 2018; Chalkidis and Søgaard, 2022). This suggests that chronological splits are necessary to more realistically evaluate rumor detection models.

We also note that the effect of temporality varies in datasets of different size. For both data splitting strategies, we observe that the difference in performance is 50% higher for the two datasets with hundreds of posts (e.g., Twitter 15 and Twitter 16) and around 10% in ones with thousands of posts (e.g., PHEME and Weibo). For rumor detection

tasks, temporality may have a greater impact on small-scale benchmarks than on large-scale benchmarks. For Twitter 16 and Weibo, the use of stratified chronological splits demonstrates significant performance drops compared to random splits due to the temporal concept drift.

For chronological splits, we observe that pre-trained language models (i.e., BERT and BERT+) significantly outperform (t-test, $p < 0.05$) logistic regression in all benchmarks. This is due to the fact that BERT-style models (i) outperform simpler linear models by a large margin in various NLP tasks (Devlin et al., 2019); and (ii) have been trained after the development of these four benchmarks implying some information leakage.

Standard vs. Stratified Chronological Splits

Note that dividing the datasets into standard chronological splits results in subsets that do not preserve the sample percentages for each category (see Table 3). The upper part of Table 4 displays the difference in model performance between two types of chronological splits on Twitter 15 and PHEME. We observe that using both standard and stratified chronological splits results in similar model predictive performance (t-test, $p > 0.05$). Even though stratified chronological splits contain temporal overlap, it is still not sufficient to improve model performance compared to random splits. This suggests that the temporal drift affects particular classes rather than the entire data set.

4 Error Analysis

Finally, we perform an error analysis to further investigate the type of errors made by BERT us-

Benchmark		Twitter 15			Twitter 16			PHEME			Weibo		
Splits	Test set	total	#	%	total	#	%	total	#	%	total	#	%
Chrono.	all posts	148	3	2%	82	6	7%	1161	39	3%	933	41	4%
	# of wrong predictions	63	2	1%	34	2	2%	301	5	<1%	99	7	<1%
	# of correct predictions	85	1	1%	48	4	5%	860	34	3%	834	34	4%
Random	all posts	149	35	23%	83	26	30%	1161	181	16%	933	129	14%
	# of wrong predictions	12	0	0%	5	1	1%	150	14	1%	65	4	<1%
	# of correct predictions	137	35	23%	78	25	30%	1011	167	14%	868	125	14%

Table 6: Error Analysis for all benchmarks. # denotes the number of posts that are similar to posts from training set, i.e., known data. % denote the percentage of similar posts in the test set. We set the threshold value to 20, which indicates that there are two or three different words between the two tweets.

	Example	Test	Train	Correct	Wrong
Twitter 15	#rip to the driver who died with #paulwalker that no one cares about because he wasn't famous.	4	6	4	0
Twitter 16	steve jobs was adopted. his biological father was abdufatah jandali, a syrian muslim	2	13	2	0
PHEME	Police are leaving now . #ferguson HTTPURL	4	11	4	0
Weibo	【交通新规】2013年1月1日施行:1... 扩散给大家! [广州日报] Translation: [New driving laws] From 1 Jan 2013: Running a red light will result in a fine of 100 RMB and 6 points. ... Spread the news to everyone! [Guangzhou Daily]	2	6	2	0

Table 7: Four examples of correct predictions using random splits, which artificially removes temporal concept drift. For example, in Twitter 15, there are 4 and 6 similar posts about rumors related to Paul Walker in the test set and the training set respectively.

ing both random and chronological splits. Table 6 shows the number of correct and wrong predictions for each of the two data splitting strategies. We also use the Levenshtein distance³ to calculate the quantity of posts in the test set that are similar to posts in the corresponding train set.

- We first observe that the temporal concept drift is evident in all rumor detection benchmarks. Most of the rumors on the same topic are posted in a very short time span.
- In addition, long-standing rumors are only a small part of the data (less than 5%). Second, we note that using random splits leads to topical overlap between the training and test sets (see Table 7) resulting in higher model performance.
- Finally, for both random and chronological splits, most of the posts in the test set with overlapping topics in the training set are predicted correctly. In contrast, wrong predictions are often posts with emerging or different topics compared to the posts in the train set.

5 Conclusion

We have shed light on the impact of temporal drift on computational rumor detection. Results from our controlled experiments show that the use of chronological splits causes substantially drops in predictive performance across widely-used rumor

³We set the threshold value to 20.

detection benchmarks. This suggests that random splits rather overestimate the model predictive performance. We argue that the temporal concept drift needs to be considered when developing real-world rumor detection approaches. In the future, we plan to study the impact of temporal concept drift on other NLP tasks, such as detecting user reactions to untrustworthy posts on social media (Glenski et al., 2018; Mu and Aletras, 2020; Mu et al., 2022).

Limitations

We provide the first re-evaluation of four standard rumor detection benchmarks in two languages (English and Chinese) from two platforms (Twitter and Weibo). We acknowledge that further investigation is needed in rumor detection datasets in other languages. We provide an error analysis in Section 4.

Acknowledgments

We would like to thank Ahmed Alajrami, Danae Sánchez Villegas, Mali Jin, Xutan Peng and all the anonymous reviewers for their valuable feedback.

References

- Oshin Agarwal and Ani Nenkova. 2022. Temporal effects on pre-trained models for language processing tasks. *Transactions of the Association for Computational Linguistics*, 10:904–921.
- Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 549–556.

- Ilias Chalkidis and Anders Søgaard. 2022. [Improved multi-label classification under temporal concept drift: Rethinking group-robust algorithms in a label-wise setting](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2441–2454, Dublin, Ireland. Association for Computational Linguistics.
- Nisansa de Silva and Dejing Dou. 2021. [Semantic oppositionness assisted deep contextual modeling for automatic rumor detection in social networks](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 405–415, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Parsa Farinneya, Mohammad Mahdi Abdollah Pour, Sardar Hamidian, and Mona Diab. 2021. [Active learning for rumor identification on social media](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4556–4565.
- Komal Florio, Valerio Basile, Marco Polignano, Pierpaolo Basile, and Viviana Patti. 2020. [Time of your hate: The challenge of time in hate speech detection on social media](#). *Applied Sciences*, 10(12):4180.
- Maria Glenski, Tim Weninger, and Svitlana Volkova. 2018. [Identifying and understanding user reactions to deceptive and trusted social news sources](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 176–181, Melbourne, Australia. Association for Computational Linguistics.
- Kyle Gorman and Steven Bedrick. 2019. [We need to talk about standard splits](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.
- Xiaolei Huang and Michael J. Paul. 2018. [Examining temporality in document classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 694–699, Melbourne, Australia. Association for Computational Linguistics.
- Xiaolei Huang and Michael J. Paul. 2019. [Neural temporality adaptation for document classification: Diachronic word embeddings and domain adaptation models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4113–4123, Florence, Italy. Association for Computational Linguistics.
- Twin Karmakharm, Nikolaos Aletras, and Kalina Bontcheva. 2019. [Journalist-in-the-loop: Continuous learning as a service for rumour analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 115–120, Hong Kong, China. Association for Computational Linguistics.
- Elena Kochkina and Maria Liakata. 2020. [Estimating predictive uncertainty for rumour verification models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6964–6981, Online. Association for Computational Linguistics.
- Vladimir I Levenshtein et al. 1966. [Binary codes capable of correcting deletions, insertions, and reversals](#). In *Soviet physics doklady*, pages 707–710. Soviet Union.
- Jiawen Li, Shiwen Ni, and Hung-Yu Kao. 2021. [Meet the truth: Leverage objective facts and subjective views for interpretable rumor detection](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 705–715, Online. Association for Computational Linguistics.
- Hongzhan Lin, Jing Ma, Mingfei Cheng, Zhiwei Yang, Liangliang Chen, and Guang Chen. 2021. [Rumor detection on Twitter with claim-guided hierarchical graph attention networks](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10035–10047, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yi-Ju Lu and Cheng-Te Li. 2020. [GCAN: Graph-aware co-attention networks for explainable fake news detection on social media](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online. Association for Computational Linguistics.
- Michal Lukasik, Trevor Cohn, and Kalina Bontcheva. 2015. [Classifying tweet level judgements of rumours in social media](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2590–2595, Lisbon, Portugal. Association for Computational Linguistics.
- Michal Lukasik, P. K. Srijith, Duy Vu, Kalina Bontcheva, Arkaitz Zubiaga, and Trevor Cohn. 2016. [Hawkes processes for continuous time sequence classification: an application to rumour stance classification in Twitter](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 393–398, Berlin, Germany. Association for Computational Linguistics.
- Jan Lukes and Anders Søgaard. 2018. [Sentiment analysis under temporal shift](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages

- 65–71, Brussels, Belgium. Association for Computational Linguistics.
- Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 3818–3824.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. [Detect rumors in microblog posts using propagation structure via kernel learning](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–717, Vancouver, Canada. Association for Computational Linguistics.
- Yida Mu and Nikolaos Aletras. 2020. Identifying twitter users who repost unreliable news sources with linguistic information. *PeerJ Computer Science*, 6:e325.
- Yida Mu, Pu Niu, and Nikolaos Aletras. 2022. [Identifying and characterizing active citizens who refute misinformation in social media](#). In *14th ACM Web Science Conference 2022, WebSci '22*, page 401–410, New York, NY, USA. Association for Computing Machinery.
- Rajdeep Mukherjee, Uppada Vishnu, Hari Chandana Peruri, Sourangshu Bhattacharya, Koustav Rudra, Pawan Goyal, and Niloy Ganguly. 2022. [Mtlts: A multi-task framework to obtain trustworthy summaries from crisis-related microblogs](#). In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 755–763.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Shengsheng Qian, Jinguang Wang, Jun Hu, Quan Fang, and Changsheng Xu. 2021. Hierarchical multi-modal contextual attention network for fake news detection. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 153–162.
- Dongning Rao, Xin Miao, Zhihua Jiang, and Ran Li. 2021. [STANKER: Stacking network based on level-grained attention-masked BERT for rumor detection on social media](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3347–3363, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiaoying Ren, Jing Jiang, Ling Min Serena Khoo, and Hai Leong Chieu. 2021. [Cross-topic rumor detection using topic-mixtures](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1534–1538, Online. Association for Computational Linguistics.
- Qiang Sheng, Juan Cao, Xueyao Zhang, Rundong Li, Danding Wang, and Yongchun Zhu. 2022. Zoom out and observe: News environment perception for fake news detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4543–4556.
- Anders Søgaard, Sebastian Ebert, Jasmijn Bastings, and Katja Filippova. 2021. [We need to talk about random splits](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1823–1832, Online. Association for Computational Linguistics.
- Yun-Zhu Song, Yi-Syuan Chen, Yi-Ting Chang, Shao-Yu Weng, and Hong-Han Shuai. 2021. [Adversary-aware rumor detection](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1371–1382, Online. Association for Computational Linguistics.
- Mengzhu Sun, Xi Zhang, Jianqiang Ma, and Yazheng Liu. 2021. [Inconsistency matters: A knowledge-guided dual-inconsistency network for multi-modal rumor detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1412–1423, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tiening Sun, Zhong Qian, Sujun Dong, Peifeng Li, and Qiaoming Zhu. 2022. [Rumor detection on social media with graph adversarial contrastive learning](#). In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 2789–2797, New York, NY, USA. Association for Computing Machinery.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. 2020. Ernie 2.0: A continual pre-training framework for language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8968–8975.
- Lin Tian, Xiuzhen Zhang, and Jey Han Lau. 2022. [DUCK: Rumour detection on social media by modelling user and comment propagation networks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4939–4949, Seattle, United States. Association for Computational Linguistics.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Lingwei Wei, Dou Hu, Wei Zhou, Zhaojuan Yue, and Songlin Hu. 2021. [Towards propagation uncertainty: Edge-enhanced Bayesian graph convolutional networks for rumor detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*

Conference on Natural Language Processing (Volume 1: Long Papers), pages 3845–3854, Online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Rui Xia, Kaizhou Xuan, and Jianfei Yu. 2020. [A state-independent and time-evolving network for early rumor detection in social media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9042–9051, Online. Association for Computational Linguistics.

Jianfei Yu, Jing Jiang, Ling Min Serena Khoo, Hai Leong Chieu, and Rui Xia. 2020. [Coupled hierarchical transformer for stance-aware rumor verification in social media conversations](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1392–1401, Online. Association for Computational Linguistics.

Fengzhu Zeng and Wei Gao. 2022. Early rumor detection using neural hawkes process with a new benchmark dataset. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4105–4117.

Zhixue Zhao, George Chrysostomou, Kalina Bontcheva, and Nikolaos Aletras. 2022. [On the impact of temporal concept drift on model explanations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4039–4054, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Kaimin Zhou, Chang Shu, Binyang Li, and Jey Han Lau. 2019. [Early rumour detection](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1614–1623, Minneapolis, Minnesota. Association for Computational Linguistics.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys*, 51(2):1–36.

Arkaitz Zubiaga, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Peter Tolmie. 2016. Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PloS one*, 11(3):e0150989.