

# Automatic Fake News Detection: Are Models Learning to Reason?

**Casper Hansen\***

University of Copenhagen  
c.hansen@di.ku.dk

**Christian Hansen\***

University of Copenhagen  
chrh@di.ku.dk

**Lucas Chaves Lima**

University of Copenhagen  
lcl@di.ku.dk

## Abstract

Most fact checking models for automatic fake news detection are based on reasoning: given a claim with associated evidence, the models aim to estimate the claim veracity based on the supporting or refuting content within the evidence. When these models perform well, it is generally assumed to be due to the models having learned to reason over the evidence with regards to the claim. In this paper, we investigate this assumption of reasoning, by exploring the relationship and importance of both claim and evidence. Surprisingly, we find on political fact checking datasets that most often the highest effectiveness is obtained by utilizing only the evidence, as the impact of including the claim is either negligible or harmful to the effectiveness. This highlights an important problem in what constitutes evidence in existing approaches for automatic fake news detection.

## 1 Introduction

Misinformation is spreading at increasing rates (Vosoughi et al., 2018), particularly online, and is considered a highly pressing issue by the World Economic Forum (Howell et al, 2013). To combat this problem, automatic fact checking, especially for estimating the veracity of potential fake news, have been extensively researched (Hassan et al., 2017; Hansen et al., 2019; Thorne and Vlachos, 2018; Elsayed et al., 2019; Allein et al., 2020; Popat et al., 2018; Augenstein et al., 2019). Given a claim, most fact checking systems are *evidence-based*, meaning they utilize external knowledge to determine the claim veracity. Such external knowledge may consist of previously fact checked claims (Shaar et al., 2020), but it typically consists of using the claim to query the web through a search API to retrieve relevant hits. While including the evidence in the model increases the effectiveness

over using only the claim, existing work has not focused on the predictive power of isolated evidence, and hence whether it assists the model in enabling better reasoning.

In this work we investigate if fact checking models learn reasoning, i.e., provided a claim and associated evidence, whether the model determines the claim veracity by reasoning over the evidence. If the model learns reasoning, we would expect the following proposition to hold: *A model using both the claim and evidence should perform better on the task of fact checking compared to a model using only the claim or evidence.* If a model is only given the claim as input, it does not necessarily have the information needed to determine the veracity. Similarly, if the model is only given the evidence, the predictive signal must come from dataset bias or the differences in the evidence obtained from claims with varying veracity, as it otherwise corresponds to being able to provide an answer to an unknown question. In our experimental evaluation on two political fact checking datasets, across multiple types of claim and evidence representations, we find the evidence provides a very strong predictive signal independent of the claim, and that the best performance is most often obtained while entirely ignoring the claim. This highlights that fact checking models may not be learning to reason, but instead exploit an inherent signal in the evidence itself, which can be used to determine factuality independent of using the claim as part of the model input. This highlights an important problem in what constitutes evidence in existing approaches for automatic fake news detection. We make our code publicly available<sup>1</sup>.

---

<sup>1</sup><https://github.com/casperhansen/fake-news-reasoning>

\*Equal contribution.

## 2 Related Work

Automatic fact checking models include deep learning approaches, based on contextual and non-contextual embeddings, which encode the claim and evidence using RNNs or Transformers (Shaar et al., 2020; Elsayed et al., 2019; Allein et al., 2020; Popat et al., 2018; Augenstein et al., 2019; Hassan et al., 2017), and non-deep learning approaches (Wang, 2017; Pérez-Rosas et al., 2018), which uses hand-crafted features or bag-of-word representations as input to traditional machine learning classifiers such as random forests, SVM, and MLP (Mihalcea and Strapparava, 2009; Pérez-Rosas et al., 2018; Baly et al., 2018; Reddy et al., 2018).

Generally, models may learn to memorize artifact and biases rather than truly learning (Gururangan et al., 2018; Moosavi and Strube, 2017; Agrawal et al., 2016), e.g., from political individuals often leaning towards one side of the truth spectrum. Additionally, language models have been shown to implicitly store world knowledge (Roberts et al., 2020), which in principle could enhance the aforementioned biases. To this end, we design our experimental setup to include representative fact checking models of varying complexity (from simple term-frequency based representations to contextual embeddings), while always evaluating each trained model on multiple different datasets to determine generalizability.

## 3 Methods

**Problem definition.** In automatic fact checking of fake news we are provided with a dataset of  $D = \{(c_1, e_1, y_1), \dots, (c_n, e_n, y_n)\}$ , where  $c_i$  corresponds to a textual claim,  $e_i$  is evidence used to support or refute the claim, and  $y_i$  is the associated truth label to be predicted based on the claim and evidence. Following current work on fact checking of fake news (Hassan et al., 2017; Thorne and Vlachos, 2018; Elsayed et al., 2019; Allein et al., 2020; Popat et al., 2018; Augenstein et al., 2019), we consider the evidence to be a list of top-10 search snippets as returned by Google search API when using the claim as the query. Note that while additional metadata may be available—such as speaker, checker, and tags—this work focuses specifically on whether models learn to reason based on the combination of claim and evidence, hence we keep the input representation to consist only of the latter.

**Overview.** In the following we describe the different models used for the experimental compari-

son (Section 4), which consists of models based on term frequency (term-frequency weighted bag-of-words (Salton and Buckley, 1988)), word embeddings (GloVe word embeddings (Pennington et al., 2014)), and contextual word embeddings (BERT (Devlin et al., 2019)). These representations are chosen as to include the typical representations most broadly used among past and current NLP models.

**Term-frequency based Random Forest.** We construct a term-frequency weighted bag-of-words representation per sample based on concatenating the text content of the claim and associated evidence snippets. We train a Random Forest (Breiman, 2001) as the classifier using the Gini impurity measure. In the setting of only using either the claim or evidence snippets as the input, only the relevant part is used for constructing the bag-of-words representation.

**GloVe-based LSTM model.** We adapt the model by Augenstein et al. (2019), which originally was proposed for multi-domain veracity prediction. Using a pretrained GloVe embedding (Pennington et al., 2014)<sup>2</sup>, claim and snippet tokens are embedded into a joint space. We encode the claim and snippets using an attention-weighted bidirectional LSTM (Hochreiter and Schmidhuber, 1997):

$$h_{c_i} = \text{attn}(\text{BiLSTM}(c_i)) \quad (1)$$

$$h_{e_{i,j}} = \text{attn}(\text{BiLSTM}(e_{i,j})) \quad (2)$$

where  $\text{attn}(\cdot)$  is a function that learns an attention score per element, which is normalized using a softmax, and returns a weighted sum. We combine the claim and snippet encodings using the matching model by Mou et al. (2016) as:

$$s_{i,j} = [h_{c_i}; h_{e_{i,j}}; h_{c_i} - h_{e_{i,j}}; h_{c_i} \cdot h_{e_{i,j}}] \quad (3)$$

where “;” denotes concatenation. The joint claim-evidence encodings are attention weighted and summed, projected through a fully connected layer into  $\mathbb{R}^L$ , where  $L$  is the number of possible labels:

$$o_i = \text{attn}([s_{i,1}; \dots; s_{i,10}]) \quad (4)$$

$$p_i = \text{softmax}(\text{FC}(o_i)) \quad (5)$$

Lastly, the model is trained using cross entropy as the loss function. In the setting of using only the claim as the input (i.e., without the evidence), then  $h_{c_i}$  is used in Eq. 5 instead of  $o_i$ . If only the

<sup>2</sup><http://nlp.stanford.edu/data/glove.840B.300d.zip>

RF (~13 seconds)	Train: Snopes				Train: PolitiFact			
	Within dataset		Out-of dataset		Within dataset		Out-of dataset	
	Eval: Snopes		Eval: PolitiFact		Eval: PolitiFact		Eval: Snopes	
	F1 <sub>micro</sub>	F1 <sub>macro</sub>	F1 <sub>micro</sub>	F1 <sub>macro</sub>	F1 <sub>micro</sub>	F1 <sub>macro</sub>	F1 <sub>micro</sub>	F1 <sub>macro</sub>
Claim	0.473	0.231	<b>0.273</b>	0.223	0.254	0.255	0.546	0.243
Evidence	0.504	<u>0.280</u>	0.244	0.195	0.301	0.299	<b>0.597</b>	0.232
Claim+Evidence	<u>0.550</u>	0.271	0.245	0.190	<u>0.310</u>	<u>0.304</u>	0.579	0.207
LSTM (~12 minutes, 888K parameters)								
Claim	0.408	0.243	0.260	<b>0.228</b>	0.237	0.237	<u>0.565</u>	0.221
Evidence	0.495	<u>0.253</u>	<u>0.262</u>	0.208	<u>0.290</u>	<u>0.295</u>	0.550	<u>0.273</u>
Claim+Evidence	<u>0.529</u>	<u>0.253</u>	0.258	0.189	0.288	0.294	0.509	0.256
BERT (~264 minutes, 109M parameters)								
Claim	0.533	0.312	<u>0.249</u>	0.209	0.275	0.282	0.550	0.273
Evidence	0.531	<b>0.321</b>	<u>0.249</u>	<u>0.224</u>	<b>0.351</b>	<b>0.359</b>	<u>0.577</u>	<b>0.286</b>
Claim+Evidence	<b>0.556</b>	0.313	0.231	0.191	0.285	0.292	0.564	0.259

Table 1: Evaluation using micro and macro F1. Per column, the best score per method is underlined and the best score across all methods is highlighted in bold. We report the training time and number of model parameters, for Claim+Evidence on PolitiFact, in the parentheses. RF is trained on 5 cores and neural models on a Titan RTX.

evidence is used, then an attention weighted sum of the evidence snippet encodings is used in Eq. 5 instead of  $o_i$ .

**BERT-based model.** In a similar fashion to the LSTM model, we construct a model based on BERT (Devlin et al., 2019)<sup>3</sup>, where the [CLS] token encoding is used for claim and evidence representations. Specifically, the claim and evidence snippets are encoded as:

$$h_{c_i} = \text{BERT}(c_i), h_{e_{i,j}} = \text{BERT}(c_i, e_{i,j}) \quad (6)$$

$$h_{e_i} = \text{attn}([h_{e_{i,1}}; \dots; h_{e_{i,10}}]) \quad (7)$$

where the claim acts as the question when encoding the evidence snippets. Similarly to Eq. 5, the prediction is obtained by concatenating the claim and evidence representations and project it through a fully connected layer into  $\mathbb{R}^L$ :

$$p_i = \text{softmax}(FC([h_{c_i}; h_{e_i}])) \quad (8)$$

where cross entropy is used as the loss function for training the model. In the setting that only the claim is used as input, then only  $h_{c_i}$  is used in Eq. 8. If only the evidence is used, then  $h_{e_{i,j}}$  is computed without including  $c_i$ , and only  $h_{e_i}$  is used in Eq. 8.

<sup>3</sup>We use bert-base-uncased from <https://huggingface.co/bert-base-uncased>.

	#Claims	Labels
PolitiFact	13,581	pants on fire! (10.6%), false (19.2%), mostly false (17.0%), half-true (19.8%), mostly true (18.8%), true (14.8%)
Snopes	5,069	false (64.3%), mostly false (7.5%), mixture (12.3%), mostly true (2.8%), true (13.0%)

Table 2: Dataset statistics.

## 4 Experimental Evaluation

### 4.1 Datasets

We focus on the domain of political fact checking, where we use claims and associated evidence from PolitiFact and Snopes, which we extract from the MultiFC dataset (Augenstein et al., 2019). Using the claim as a query, the evidence is crawled from Google search API as the search snippets of the top-10 results, and is filtered such that the website origin of a given claim does not appear as evidence. To facilitate better comparison between the datasets, we filter claims with non-veracity related labels<sup>4</sup>. The dataset statistics are shown in Table 2.

### 4.2 Experimental setup

Both datasets are split into train/val/test sets using label-stratified sampling (70/10/20% splits). We tune all models on the validation split, and use early stopping with a patience of 10 for neural

<sup>4</sup>For PolitiFact we exclude [full flop, half flip, no flip] and for Snopes we exclude [unproven, miscaptioned, legend, outdated, misattributed, scam, correct attribution].

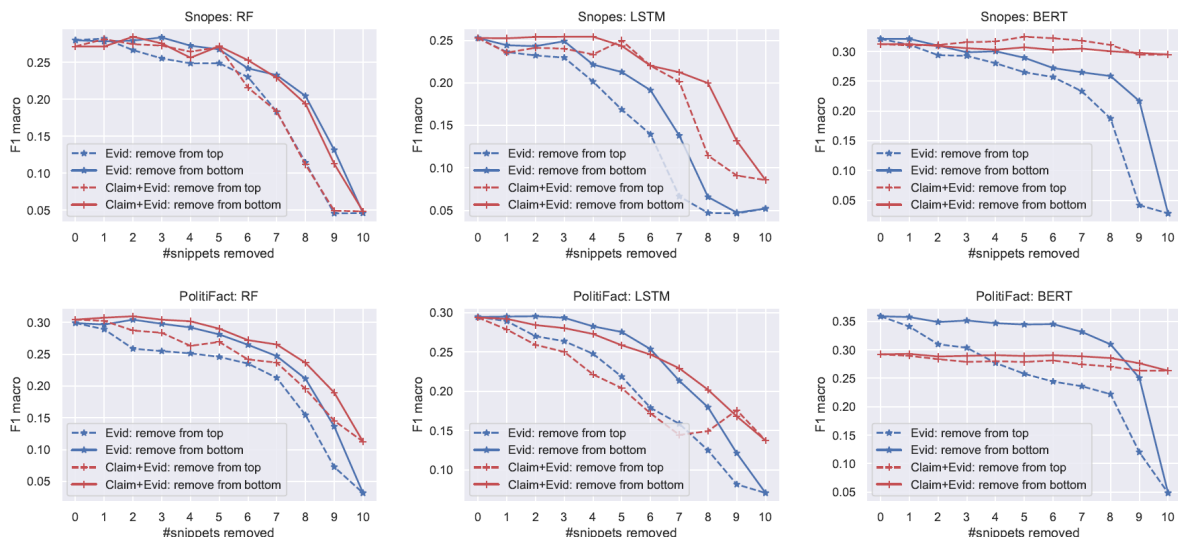


Figure 1: Macro F1 scores when removing evidence from either the top or bottom of the evidence snippet ranking.

models. Following [Augenstein et al. \(2019\)](#), we use micro and macro F1 for evaluation. The models are evaluated on both the within dataset test sets, but also out-of dataset test sets (e.g., a model trained on Snopes is evaluated on both Snopes and PolitiFact). In the out-of dataset evaluation we need the labels to be comparable, hence in that setting we merge "pants on fire!" and "false" for PolitiFact.

## 5 Tuning details

In the following, the best overall parameter configurations are underlined>. The best configuration is chosen based on the average of the micro and macro F1<sup>5</sup>. For RF, we tune the number of trees from [100,500,1000], the minimum number of samples in a leaf from [1,3,5,10], and the minimum number of samples per split from [2,5,10]. For the LSTM model, we tune the learning rate from [1e-4,5e-4,1e-5], batch size [16,32], number of LSTM layers from [1,2], dropout from [0, 0.1], and fix the number of hidden dimensions to 128. For the BERT model, we tune the learning rate from [3e-5, 3e-6, 3e-7] and fix the batch size to 8.

### 5.1 Results

The results can be seen in Table 1. Overall, we see that the BERT model trained only on Evidence obtains the best results in 4/8 columns, and, notably, in 3/4 cases the BERT model with Evidence obtains the best macro F1 score on within and out-

of dataset prediction. Random forest using term-frequency as input obtains the best out-of dataset micro F1 for both datasets (using either only Claim or only Evidence). Across all methods, the combination of Claim+Evidence only marginally obtains the best results a single time (for Snopes micro F1). For further details, in Table 3 we compute the accuracy scores for all the false labels, mixture or half-true label, and true labels.

Surprisingly, a BERT model using only the Evidence is capable of predicting the veracity of the claim used for obtaining the evidence. This shows that a strong signal must exist in the evidence itself, and the evidence found by the search engine appears to be implicitly affected by the veracity of the claim used as the query in some way<sup>6</sup>. The improvements reported in the literature by combining claim and evidence, are therefore not evident of the model learning to reason over the evidence with regards to the claim, but instead exploiting a signal inherent in the evidence itself. This highlights that the current approach for evidence gathering is problematic, as the strong signal makes it possible (and most often beneficial) for the model to entirely ignore the claim. This makes the model entirely reliant on the process behind how the evidence is generated, which is outside the scope of the model, and thereby undesirable, as any change in the search system may change the model performance significantly. It may also be problematic on a more fundamental level, e.g., to predict the

<sup>5</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.metrics.f1\\_score.html](https://scikit-learn.org/stable/modules/generated/sklearn.metrics.f1_score.html)

<sup>6</sup>Note that the claim origin website is always removed from the evidence.

RF	Train: Snopes						Train: PolitiFact					
	Within dataset Eval: Snopes			Out-of dataset Eval: PolitiFact			Within dataset Eval: PolitiFact			Out-of dataset Eval: Snopes		
	acc <sub>false</sub>	acc <sub>mix</sub>	acc <sub>true</sub>	acc <sub>false</sub>	acc <sub>mix</sub>	acc <sub>true</sub>	acc <sub>false</sub>	acc <sub>mix</sub>	acc <sub>true</sub>	acc <sub>false</sub>	acc <sub>mix</sub>	acc <sub>true</sub>
Claim	0.710	0.144	0.255	0.853	0.016	0.209	0.623	0.216	0.513	0.790	0.092	0.255
Evidence	0.705	0.152	0.441	0.829	0.006	0.117	0.654	0.248	0.510	<b>0.891</b>	0.039	0.192
Claim+Evidence	<b>0.760</b>	0.136	0.453	0.829	0.000	0.117	0.634	0.292	0.512	0.871	0.039	0.199
LSTM												
Claim	0.674	0.232	0.280	0.875	0.047	0.137	0.566	0.212	0.504	0.833	0.026	0.234
Evidence	0.721	0.272	0.267	<b>0.890</b>	0.020	0.115	0.643	0.253	0.485	0.768	0.184	0.322
Claim+Evidence	0.757	0.248	0.168	0.879	0.008	0.107	<b>0.671</b>	0.210	0.460	0.704	0.171	<b>0.378</b>
BERT												
Claim	0.746	0.256	0.379	0.854	<b>0.094</b>	0.045	0.604	0.292	0.475	0.765	0.171	0.287
Evidence	0.648	<b>0.376</b>	<b>0.559</b>	0.702	0.049	<b>0.337</b>	0.649	<b>0.326</b>	0.496	0.804	<b>0.197</b>	0.339
Claim+Evidence	0.747	0.264	0.379	0.882	0.067	0.042	0.667	0.175	<b>0.558</b>	0.790	0.092	0.367

Table 3: Accuracy scores computed on the false labels, mixture or half-true label, and true labels. All labels within a group (e.g., any false label such as false or mostly false) are considered to be the same and as such this reduces the problem to a three class classification problem.

veracity of the following two claims: ”the earth is round” and ”the earth is flat”, the evidence could be the same, but a model entirely dependent on the evidence, and not the claim, would be incapable of predicting both claims correctly.

## 5.2 Removal of evidence

We observed a strong predictive signal in the evidence alone and now consider the performance impact when gradually removing evidence snippets. The evidence is removed consecutively either from the top down or bottom up (i.e., removing the most relevant snippets first and vice versa), until no evidence is used. Figure 1 shows the macro F1 as a function of removed evidence when using the Evidence or Claim+Evidence models. We observe a distinct difference between the random forest and LSTM model compared to BERT: for random forest and LSTM, the Claim+Evidence models on both datasets drop rapidly in performance when the evidence is removed, while the BERT model only experiences a very small drop. This shows that when the Claim+Evidence is used in the BERT model, the influence of the evidence is minimal, while the evidence is vital for the Claim+Evidence RF and LSTM models. For all models, we observe that when evidence is removed from the top down, the performance drop is larger than when evidence is removed from the bottom up. Thus, the ranking of the evidence as provided by the search engine is related to its usefulness as evidence for fact checking.

## 6 Conclusion

We investigate if fact checking models for fake news detection are learning to process claim and evidence jointly in a way resembling reasoning. Across models of varying complexity and evaluated on multiple datasets, we find that the best performance can most often be obtained using only the evidence. This highlights that models using both claim and evidence are inherently not learning to reason, and points to a potential problem in how evidence is currently obtained in existing approaches for automatic fake news detection.

## References

- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. *Analyzing the behavior of visual question answering models*. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, Austin, Texas. Association for Computational Linguistics.
- Lee Howell et al. 2013. Digital wildfires in a hyperconnected world. *WEF report*, 45(3):15–94.
- Liesbeth Allein, Isabelle Augenstein, and Marie-Francine Moens. 2020. Time-aware evidence ranking for fact-checking. *arXiv preprint arXiv:2009.06402*.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. Multifaceted: A real-world multi-domain dataset for evidence-based fact checking of claims. 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), pages 4677–4691.
- Ramy Baly, Mitra Mohtarami, James R. Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. [Integrating Stance Detection and Fact Checking in a Unified Corpus](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 21–27.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- 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.
- Tamer Elsayed, Preslav Nakov, Alberto Barrón-Cedeño, Maram Hasanain, Reem Suwaileh, Giovanni Da San Martino, and Pepa Atanasova. 2019. Checkthat! at clef 2019: Automatic identification and verification of claims. In *Advances in Information Retrieval*, pages 309–315, Cham. Springer International Publishing.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural check-worthiness ranking with weak supervision: Finding sentences for fact-checking. In *Companion Proceedings of the 2019 World Wide Web Conference*, pages 994–1000.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claimbuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1803–1812.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Rada Mihalcea and Carlo Strapparava. 2009. The lie detector: Explorations in the automatic recognition of deceptive language. In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, ACLShort '09, page 309–312, USA. Association for Computational Linguistics.
- Nafise Sadat Moosavi and Michael Strube. 2017. [Lexical features in coreference resolution: To be used with caution](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 14–19, Vancouver, Canada. Association for Computational Linguistics.
- Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. 2016. Natural language inference by tree-based convolution and heuristic matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 130–136.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. [Automatic Detection of Fake News](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401. Association for Computational Linguistics.
- Kashyap Papat, Subhabrata Mukherjee, Andrew Yates, and Gerhard Weikum. 2018. Declare: Debunking fake news and false claims using evidence-aware deep learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 22–32.
- Aniketh Janardhan Reddy, Gil Rocha, and Diego Esteves. 2018. [Defactonlp: Fact verification using entity recognition, TFIDF vector comparison and decomposable attention](#). *CoRR*, abs/1809.00509.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426.
- Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523.
- Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. 2020. That is a known lie: Detecting previously fact-checked claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3607–3618.
- James Thorne and Andreas Vlachos. 2018. Automated fact checking: Task formulations, methods and future directions. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3346–3359.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.

William Yang Wang. 2017. “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426. Association for Computational Linguistics.