

Robust Information Retrieval for False Claims with Distracting Entities In Fact Extraction and Verification

Mingwen Dong Christos Christodoulopoulos* Sheng-Min Shih Xiaofei Ma
mingwd@amazon.com chrchrs@amazon.co.uk shengms@amazon.com xiaofeim@amazon.com

Abstract

Accurate evidence retrieval is essential for automated fact checking. Little previous research has focused on the differences between true and false claims and how they affect evidence retrieval. This paper shows that, compared with true claims, false claims more frequently contain irrelevant entities which can distract evidence retrieval model. A BERT-based retrieval model made more mistakes in retrieving refuting evidence for false claims than supporting evidence for true claims. When tested with adversarial false claims (synthetically generated) containing irrelevant entities, the recall of the retrieval model is significantly lower than that for original claims. These results suggest that the vanilla BERT-based retrieval model is not robust to irrelevant entities in the false claims. By augmenting the training data with synthetic false claims containing irrelevant entities, the trained model achieved higher evidence recall, including that of false claims with irrelevant entities. In addition, using separate models to retrieve refuting and supporting evidence and then aggregating them can also increase the evidence recall, including that of false claims with irrelevant entities. These results suggest that we can increase the BERT-based retrieval model’s robustness to false claims with irrelevant entities via data augmentation and model ensemble.

1 Introduction

Automated fact check has received a lot of attention recently and several datasets have been developed to facilitate relevant research (Vlachos and Riedel, 2014; Wang, 2017; Pomerleau and Rao, 2017; Hanselowski et al., 2018a). However, most of these datasets contain limited amount of examples and the human annotation is only at the

document level. Without a large amount of sentence level annotations, it is difficult to train latest neural-network based models to detect false claims. The Fact Extraction and VERification (FEVER) task was introduced to solve these data challenges (Thorne et al., 2018). The task is to fact-check sentences (claims) people wrote to be true or false by retrieving evidence sentences from Wikipedia and using them to verify the validity of the claim. True claims are validated by supporting evidence that entails them and false claims are verified by refuting evidence that contradicts them.

A lot of progress has been made on the FEVER task using classical and deep learning methods (Yoneda et al., 2018b; Hanselowski et al., 2018b; T.Y.S.S et al., 2019; Soleimani et al., 2020; Hidey et al., 2020; Lee et al., 2020). Our paper focuses on evidence retrieval and we review some related progress. Thorne et al. (2018) used TF-IDF for evidence retrieval to show the feasibility and challenge of the task. Yoneda et al. (2018b) improved the evidence retrieval accuracy by training logistic regression and MLP models using manually engineered features (e.g., token matching between claim and evidence sentence, evidence sentence length). Based on the observation that most claims focus on Wikipedia entities, Hanselowski et al. (2018b) extracted entities from claims using AllenNLP and then used those entities as queries to retrieve relevant documents leveraging MediaWiki search API (Gardner et al., 2018). Soleimani et al. (2020) combined the method from Hanselowski et al. (2018b) with BERT-based models and improved the evidence retrieval performance by 2.5%. Recently, Hidey et al. (2020) achieved state-of-the-art results using a combination of BERT models and pointer networks (Vinyals et al., 2015; Devlin et al., 2019).

Few research has reported separate results for false and true claims, possibly because they treat

Christos Christodoulopoulos is affiliated with Amazon Alexa. Other authors are affiliated with Amazon Web Services, Inc.

Table 1: A false claim with its refuting evidence from FEVER. “Stan Beeman” and “BBC” are two entities. “BBC” is an irrelevant entity that may distract an information retrieval model.

| | |
|--------------|--------------------------------------|
| Claim | Stan Beeman is only in shows on BBC. |
| Evidence | Stan Beeman acts in a US TV series. |
| Relationship | REFUTED |

evidence retrieval for FEVER as a standard information retrieval (IR) task and used the same IR model to retrieve refuting evidence for false claims and supporting evidence for true claims. Yoneda et al. (2018a) and Chakrabarty et al. (2018) reported that the accuracy is lower in detecting false claims than detecting true claims, which suggests that false claims and true claims have different characteristics. One difference we identified is that false claims are much more likely to contain irrelevant or distracting information than true claims. For example, in the false claim “Stan Beeman is only in shows on BBC”, “BBC” is not related to the actual evidence (Table 1) and may mislead a retrieval model to find evidence containing both BBC and Stan Beeman. In other words, an IR model may be susceptible to the distraction of the irrelevant entity “BBC” if it retrieves sentences with the largest lexical and semantic overlap with the query. In this paper, we show that a vanilla BERT-based retrieval model is not robust to irrelevant entities in false claims by analyzing the distribution of entities’ relationship in different types of claims and testing the model with synthetically generated claims containing irrelevant entities. We used two different methods to improve the retrieval model’s robustness. The first method is augmenting false claims with irrelevant entities so that the trained model is more robust to them. The second method is retrieving refuting and supporting evidence using separate models and then aggregating them, which allows the model for refuting evidence retrieval to be trained separately and become more robust to irrelevant entities in false claims. Our contributions are: 1) showing that a vanilla BERT-based retrieval model is not robust to irrelevant entities in false claims; 2) showing two effective methods to increase the model robustness, one by using data augmentation and the other by having separate models to retrieve refuting and supporting evidence and then aggregating.

2 FEVER Task & Models

FEVER dataset contains false, true, and unverifiable claims modified from sentences in Wikipedia (Thorne et al., 2018). The task is to classify each claim as REFUTED, SUPPORTED, or NOTE-NOUGHINFO (*NEI*) by retrieving relevant evidence sentences from Wikipedia and classifying the relationship between a claim and corresponding evidence sentences (Table 1).

The FEVER task is often decomposed into three sub tasks: document retrieval, sentence selection, and natural language inference. Our paper focuses on the sentence selection and compares the evidence retrieval recall and fact check accuracy of different sentence selection methods while using the same document retrieval and natural language inference models. Because the official FEVER evaluation code typically expects 5 evidence sentences per claim, we focus on the recall to evaluate evidence retrieval (calculated using official code @k=5 for sentence selection, k=20 for document retrieval) (Thorne et al., 2018). We also counted the number of mistakes in retrieving refuting and supporting evidence separately to measure a model’s robustness toward different types of claims. A model is said to make one mistake if the top k evidence retrieved do not contain any ground truth evidence (k=5 for sentence selection, k=20 for document retrieval). Lastly, the official FEVER score is used to measure the overall system performance.

2.1 Document Retrieval

The document retrieval model selects the top documents that potentially contain evidence sentences from millions of Wikipedia documents. The top 20 documents are passed to the sentence selection model to retrieve the sentences relevant to a claim. We adopted the method from Hanselowski et al. (2018b) because it is efficient, achieves greater than 93% recall (k=20 as in most previous research), and has been adopted by other previous research (Zhou et al., 2019; Soleimani et al., 2020). Table 3 shows that the adopted document

Table 2: Statistics of FEVER dataset & synthetic refuted claims.

| split | SUPPORTED | REFUTED | NEI | Synthetic REFUTED |
|----------|-----------|---------|--------|-------------------|
| Training | 80,035 | 29,775 | 35,639 | 7456 |
| Dev | 3,333 | 3,333 | 3,333 | 1102 |
| Test | 3,333 | 3,333 | 3,333 | - |
| Reserved | 6,666 | 6,666 | 6,666 | - |

Table 3: Document retrieval results @ $k = 20$ using method adopted from (Hanselowski et al., 2018b).

| Recall | Refuted mistakes | Supported mistakes |
|--------|------------------|--------------------|
| 0.931 | 147 | 82 |

retrieval method made more mistakes in retrieving refuting evidence than supporting evidence.

2.2 Sentence Selection

Given the documents from the document retrieval model, the sentence selection model selects refuting and supporting sentences for false and true claims, respectively. We use BERT-based binary classifiers for sentence selection because BERT is a general method and has been widely used in IR, FEVER and other NLP tasks (Devlin et al., 2019; Dai and Callan, 2019; Soleimani et al., 2020). To show that a fine-tuned BERT-based retrieval model is not robust to irrelevant entities in false claims, we compared multiple binary classification (relevant vs. irrelevant) BERT-models trained using different subsets of training data: both supported and refuted claims (Baseline), only supported claims (SUP), or only refuted claims (REF). We expect that the baseline model will make more mistakes than REF model for refuted claims (refuted mistake) and make more mistakes than SUP model for supported claims (supported mistake). To get the best of SUP and REF models (SR), we aggregated and reranked their results based on the confidence scores. Note that we also tried to combine results from the SUP and REF model based on the rank or a trained linear regression model, but the results were similar to that based on the confidence score. In the end, we used data augmentation and trained another model (DA), in which the original supported and refuted claims is augmented with synthetically generated false claims described in the next part.

Generating False Claims with Distracting Entities. We hypothesize that retrieving refuting

evidence for false claims is more challenging than retrieving supporting evidence for true claims for two reasons. Firstly, false claims more often contain irrelevant entities than true claims and the irrelevant entities can be distracting to the evidence retrieval models which often uses lexical matching as one important feature. Secondly, FEVER data set contains fewer training examples for false (refuted) claims than true (supported) claims (Table 2). We programmatically modified the true claims in the training and dev set to create additional false claims with irrelevant and distracting entities. For each true claim, its entities were programmatically recognized and linked to WikiData using the API from van Hulst et al. (2020). If two or more entities were identified, the 2nd entity was replaced with one of its sibling entities sampled from WikiData to generate one false claim (Figure 1). An entity’s sibling is any other entities that share the same parent. Notice that not all true claims can have a false claim generated from it because the true claim might not have a second entity with sampled sibling. We chose to replace the 2nd entity to generate false claims because the irrelevant entity of false claims in FEVER tends to be the 2nd one based on our manual check. The evidence of the original claim is treated as the refuting evidence for the generated false claim. In total, 7456 and 1102 false claims were generated from the training and dev set, respectively (Table 2). To check the quality of the synthetic false claims, we randomly sampled 100 synthetic claims and manually checked whether they are true or false. Among 100 sampled claims, 73 are false, 22 are hard to determine (probably false but not semantically meaningful), 5 are true. Because synthetic claims are used to train the augmented sentence selection model (DA) which does not distinguish between true and false claims during training, the accidentally generated true and hard to determine claims will not negatively influence the training results.

Training. A binary classifier classifies whether




Figure 1: An example false claim generated using entity linking and Wikidata.

a candidate evidence sentence is relevant to a given claim. The probability of being relevant is used to rank the candidate sentences. All models are fine-tuned to minimize the cross-entropy (negative log-likelihood) loss starting from a pre-trained BERT-base-based model using the transformers library (Wolf et al., 2019). Each model is fine-tuned for 2 epochs with a learning rate of $2.5e-6$. Following previous work, the Wikipedia page title is appended to each candidate sentence during training and inference for pronoun resolution (Yoneda et al., 2018b; Soleimani et al., 2020).

Negative Sampling. For each supported or refuted claim in the training set, the ground truth evidence sentences were used as positive examples. For each positive (relevant) sentence of a claim, we sampled 15 negative (irrelevant) sentences from a TF-IDF ranker: 5 from the documents that contain the positive sentences, 5 from documents that do not contain positive sentences, and another 5 such that each comes from a unique document that has not been sampled before. This sampling strategy ensures that the training set contains a diverse negative examples that better represent those encountered during inference. We find that the top negatives from TF-IDF alone often come from one or two documents and the model trained using them has lower recall than a MSMARCO fine-tuned model (Dai and Callan, 2019), possibly because the negatives do not reflect those seen during inference. During training, the hardest negatives with the highest predicted relevance score are used with the positives to train the model so that the numbers of positive and negative samples match (Soleimani et al., 2020).

2.3 Natural Language Inference

We adopted a similar method from Soleimani et al. (2020) for natural language inference to have a complete FEVER pipeline to evaluate our sentence selection methods on the blind test set. A pre-trained 3-class classification BERT model is fine-tuned to classify the relationship between

Table 4: Distribution of numbers of entities linked to WikiData in refuted and supported claims in the original FEVER dev set.

| # of entities | REFUTED | SUPPORTED |
|-----------------|---------|-----------|
| ≤ 1 entity | 4090 | 4166 |
| > 2 entities | 2576 | 2500 |

Table 5: Distribution of entities’ relationships in refuted and supported claims in the original FEVER dev set. The entities were linked to Wikidata using the API from van Hulst et al. (2020).

| Entities’ Relationship | REFUTED | SUPPORTED |
|------------------------|---------|-----------|
| Directly Related | 571 | 998 |
| Not Directly Related | 1928 | 1404 |

an evidence sentence and a claim as REFUTED, SUPPORTED, or NOTENOUGHINFO using cross-entropy loss. During inference, we classify the relationship between the top 5 evidence sentences with the claim separately and aggregate the results using majority vote. If there is a tie, we break it using the following sequence: NOTENOUGHINFO to SUPPORTED to REFUTED.

3 Results

False Claims are More Likely to Have Unrelated Entities than True Claims. For each refuted or supported claim in the original FEVER dev set, we programmatically linked its entities to WikiData using the API from van Hulst et al. (2020). If two or more entities were identified, we check whether there is any relationship between any 2 entities using WikiData SPARQL. If any 2 entities are related in any relationship, we say they are directly related. Otherwise, they are not directly related. Table 5 shows that the refuted claims are significantly more likely to have irrelevant entities than supported claims (χ^2 test, $p < 0.01$, $\chi^2(1) = 195.91$), even though refuted and supported claims have similar distributions in

| Model | Baseline | REF | SUP | SR | DA |
|--------------------|----------|-------|-----------|--------------|-------|
| Recall | 0.919 | 0.864 | 0.914 | 0.936 | 0.924 |
| REFUTED mistakes | 330 | 271 | 542 | 241 | 308 |
| SUPPORTED mistakes | 205 | 808 | 99 | 112 | 171 |

Table 6: Sentence selection results on the original FEVER dev set assuming perfect document retrieval system. REF model is trained only with refuted claims. SUP model is trained only with supported claims. SR model aggregates the results from SUP and REF model. DA model is trained with both original FEVER training data and synthetic false claims. See section 2.2 for details.

| Model | Baseline | SR | DA |
|------------------|----------|-------|--------------|
| Recall (k=5) | 0.688 | 0.767 | 0.813 |
| REFUTED mistakes | 282 | 187 | 141 |

Table 7: Sentence selection results on the 1102 artificially generated false claims from the dev set assuming perfect document retrieval system.

terms of number of entities (χ^2 test, $p > 0.1$, $\chi^2(1) = 1.79$, Table 4). This result shows that refuted claims are more likely to contain unrelated or irrelevant entities that may be distracting to a retrieval model than supported claims.

Baseline BERT-based Retrieval Model is Not Robust to Irrelevant Entities in Claims. To compare how various models differ in selecting refuting and supporting sentences without being biased by the document retrieval results (Table 3), the ground truth documents are appended to those from the document retrieval stage. This method simulates a realistic scenario where the sentence selection model has to retrieve the relevant sentences among a large pool of irrelevant sentences. Table 6 shows that the baseline model made more refuted mistakes than supported mistakes. This result, together with the fact that false claims are more likely to have irrelevant entities than true claims, showing that the baseline model is not robust to irrelevant entities in the false claims. To further test this hypothesis, we evaluated the recall of the baseline model in retrieving evidence for the artificially generated false claims with irrelevant entities from the dev set. Table 7 and 6 shows that the baseline model had lower recall on the adversarial false claims with irrelevant entities than on the original claims, which supports our hypothesis that the baseline model is not robust to irrelevant entities in false claims.

Improve the Model Robustness to Irrelevant Entities in False Claims. Table 6 shows that

SUP and REF model made fewer mistakes than the baseline model in retrieving supporting and refuting evidence, respectively, suggesting that recall may be improved by retrieving refuting and supporting evidence separately and then aggregating. Indeed, on the original dev set, the highest recall is achieved by the SR model, which combines the SUP and REF results. The SR model also made fewer mistakes in retrieving refuting evidence than the baseline model (Table 6). When tested on the artificially generated false claims with irrelevant entities, the SR model has higher recall than baseline model (Table 7). These results together show that SR model is more robust to irrelevant entities in the false claims than the baseline model (see Table 8 for an example). Trained with additional synthetic refuted claims, the DA model achieved higher recall than baseline model and made fewer refuted mistakes on the original dev set (Table 6). When tested on the artificially generated false claims with irrelevant entities, the DA model achieves the highest recall (Table 7). These results show that data augmentation is another effective method to improve the retrieval model’s robustness to irrelevant entities in false claims (Table 8). Note that our current data augmentation method is far from ideal because the irrelevant entity is always the 2nd entity in the artificially generated false claims. Consequently, the DA model has lower recall than the SR model on the original dev and blind test set (Table 6 & 9), even though the DA model has higher recall than the SR/baseline model for the artificially generated adversarial false claims with irrelevant entities (Table 7).

Full FEVER Pipeline. Table 9 shows that our robust models (SR & DA) achieves similar or higher recall and FEVER score on the blind test set compared with the baseline, showing that our methods generalize to the blind test set. Note that the gain in recall and FEVER score is smaller

Table 8: An example false claim from the original FEVER dev set and corresponding top 5 evidence sentences retrieved from the baseline, DA, and SR models. The ground truth evidence sentence is in bold. Note that the title of the wikipedia article is added to the beginning of each evidence sentence for pronoun resolution.

| Claim | | Johnny Galecki only acted in NBC shows. |
|----------|---|--|
| Baseline | 1 | The Tonight Show. The Tonight Show is an American late-night talk show currently broadcast from the NBC studios in Rockefeller Center in New York City (and previously from various studios in the Los Angeles region) and airing on NBC since 1954. |
| | 2 | The Tonight Show. It is the world’s longest-running talk show, and the longest running, regularly scheduled entertainment program in the United States. |
| | 3 | Blind Faith (miniseries). Blind Faith is a 1990 NBC miniseries based on the 1989 true crime book of the same name by Joe McGinniss. |
| | 4 | The Master Cleanse. The Master Cleanse is a 2016 American horror film written and directed by Bobby Miller. |
| | 5 | The Tonight Show Starring Johnny Carson. The Tonight Show Starring Johnny Carson is a talk show hosted by Johnny Carson under The Tonight Show franchise from October 1, 1962 through May 22, 1992. |
| DA | 1 | The Tonight Show. The Tonight Show is an American late-night talk show currently broadcast from the NBC studios in Rockefeller Center in New York City (and previously from various studios in the Los Angeles region) and airing on NBC since 1954. |
| | 2 | The Tonight Show. It is the world’s longest-running talk show, and the longest running, regularly scheduled entertainment program in the United States. |
| | 3 | AFL on NBC. The AFL on NBC is the branding used for broadcasts of Arena Football League (AFL) games produced by NBC Sports, the sports division of the NBC television network in the United States, that aired from the 2003 to 2006 seasons. |
| | 4 | Johnny Galecki. He is known for playing David Healy in the ABC sitcom Roseanne from 1992 – 1997 and Dr. Leonard Hofstadter in the CBS sitcom The Big Bang Theory since 2007. |
| | 5 | Blind Faith (miniseries). Blind Faith is a 1990 NBC miniseries based on the 1989 true crime book of the same name by Joe McGinniss. |
| SR | 1 | Johnny Galecki. Galecki also appeared in the films National Lampoon’s Christmas Vacation (1989), Prancer (1989), Suicide Kings (1997), I Know What You Did Last Summer (1997), Bookies (2003), and In Time (2011). |
| | 2 | The Tonight Show. The Tonight Show is an American late-night talk show currently broadcast from the NBC studios in Rockefeller Center in New York City (and previously from various studios in the Los Angeles region) and airing on NBC since 1954. |
| | 3 | The Master Cleanse. The Master Cleanse is a 2016 American horror film written and directed by Bobby Miller. |
| | 4 | Johnny Galecki. He is known for playing David Healy in the ABC sitcom Roseanne from 1992 – 1997 and Dr. Leonard Hofstadter in the CBS sitcom The Big Bang Theory since 2007. |
| | 5 | NBC. The National Broadcasting Company (NBC) is an American English language commercial broadcast television network that is a flagship property of NBCUniversal, a subsidiary of Comcast. |

| Model | Baseline | SR | DA | Soleimani et al. (2020)* |
|----------------|----------|--------------|--------|--------------------------|
| Recall (k=5) | 0.868 | 0.882 | 0.8721 | 0.752 |
| FEVER Score | 0.656 | 0.660 | 0.658 | 0.611 |
| Label accuracy | 0.705 | 0.707 | 0.707 | 0.678 |

Table 9: FEVER results on the blind test set. * indicates the results are from our implementation.

in the blind test set than those seen in the validation set. It may be because the document retrieval model is not robust to false claims with irrelevant entities and sentence selection model cannot recover from these errors when evaluated on the blind test set whereas we assumed a perfect document retrieval model for the validation set. In addition, DA & SR models perform significantly better on the generated false claims with irrelevant entities than the baseline model, suggesting DA & SR model will be at least more robust to these examples than the baseline model. Note that the result from [Soleimani et al. \(2020\)](#) in our implementation is lower than that reported in their paper. This difference is probably due to different pre-processing steps because [Soleimani et al. \(2020\)](#) only open-sourced the model training scripts but not the preprocessing scripts that generate the training data.

4 Conclusion & Future Directions

Our results, together with those from [Yoneda et al. \(2018a\)](#) and [Chakrabarty et al. \(2018\)](#), show that evidence retrieval is more challenging for false claims than for true claims. One possible reason is that false claims more often contain irrelevant entities than true claims and the baseline model is not robust to irrelevant entities in a claim. When the training data is augmented with additional synthetic false claims with irrelevant entities, the same model achieves higher evidence recall than when there are no augmented data. When separate models are trained to retrieve refuting and supporting evidence, the model for refuting evidence can learn to ignore irrelevant entities whereas the model for supporting evidence can learn to retrieve evidence that has the largest lexical and semantic overlap with the claim. After aggregating the results from the two models, we achieved the highest sentence selection recall. In the current paper, we only investigated how to improve the robustness of sentence selection model. In the future, we will look into improving the robustness of the document retrieval and natural language inference model. For example, instead of always replacing the 2nd entity in a claim, we can dynamically replace the entity while considering the relationship between entities. If synthetic claims of higher quality can be obtained, the synthetic data can be used to train the document retrieval and NLI model to further increase the

model robustness. Another future work is to investigate whether the state-of-the-art models (e.g., [Hidey et al. \(2020\)](#)) are robust to irrelevant entities in the false claims, which can potentially reveal additional ways to increase the model robustness. In the end, we suggest that future research to report separate metrics for false and true claims. In the real world, false claims but not true claims are harmful to the society. Therefore, it is important to know our models retrieve refuting evidence as well as (or better than) supporting evidence.

References

- Tuhin Chakrabarty, Tariq Alhindi, and S. Muresan. 2018. Robust document retrieval and individual evidence modeling for fact extraction and verification.
- Zhuyun Dai and Jamie Callan. 2019. [Deeper text understanding for ir with contextual neural language modeling](#). *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- 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.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew E. Peters, M. Schmitz, and Luke Zettlemoyer. 2018. Allennlp: A deep semantic natural language processing platform. *ArXiv*, abs/1803.07640.
- Andreas Hanselowski, Avinesh PVS, Benjamin Schiller, Felix Caspelherr, Debanjan Chaudhuri, Christian M. Meyer, and Iryna Gurevych. 2018a. [A retrospective analysis of the fake news challenge stance-detection task](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1859–1874, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018b. [Ukp-athene: Multi-sentence textual entailment for claim verification](#). *arXiv preprint arXiv:1809.01479*.
- Christopher Hidey, Tuhin Chakrabarty, Tariq Alhindi, Siddharth Varia, Kriste Krstovski, Mona Diab, and Smaranda Muresan. 2020. [Deseption: Dual sequence prediction and adversarial examples for improved fact-checking](#). *arXiv preprint arXiv:2004.12864*.

- Johannes M. van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P. de Vries. 2020. Rel: An entity linker standing on the shoulders of giants. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*. ACM.
- Nayeon Lee, Belinda Z. Li, Sinong Wang, Wen tau Yih, Hao Ma, and Madian Khabsa. 2020. Language models as fact checkers? *ArXiv*, abs/2006.04102.
- Dean Pomerleau and Delip Rao. 2017. [Fake news challenge](#).
- Amir Soleimani, Christof Monz, and Marcel Worring. 2020. Bert for evidence retrieval and claim verification. In *European Conference on Information Retrieval*, pages 359–366. Springer.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018. [The fact extraction and VERification \(FEVER\) shared task](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Santosh T.Y.S.S, G. Vishal, Avirup Saha, and Niloy Ganguly. 2019. Attentivechecker: A bi-directional attention flow mechanism for fact verification. In *NAACL-HLT*.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in neural information processing systems*, pages 2692–2700.
- Andreas Vlachos and S. Riedel. 2014. Fact checking: Task definition and dataset construction. In *LTCSS@ACL*.
- William Yang Wang. 2017. "liar, liar pants on fire": A new benchmark dataset for fake news detection. In *ACL*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenetorp, and S. Riedel. 2018a. Ucl machine reading group: Four factor framework for fact finding (hexaf).
- Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018b. Ucl machine reading group: Four factor framework for fact finding (hexaf). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 97–102.
- Jie Zhou, Xu Han, Cheng Yang, Z. Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2019. Gear: Graph-based evidence aggregating and reasoning for fact verification. *ArXiv*, abs/1908.01843.