

# Integrating Stance Detection and Fact Checking in a Unified Corpus

Ramy Baly<sup>1</sup>, Mitra Mohtarami<sup>1</sup>, James Glass<sup>1</sup>

Lluís Màrquez<sup>3\*</sup>, Alessandro Moschitti<sup>3\*</sup>, Preslav Nakov<sup>2</sup>

<sup>1</sup>MIT Computer Science and Artificial Intelligence Laboratory, MA, USA

<sup>2</sup>Qatar Computing Research Institute, HBKU, Qatar; <sup>3</sup>Amazon

{baly, mitram, glass}@mit.edu

{lluismv, amosch}@amazon.com; pnakov@qf.org.qa

## Abstract

A reasonable approach for fact checking a claim involves retrieving potentially relevant documents from different sources (e.g., news websites, social media, etc.), determining the stance of each document with respect to the claim, and finally making a prediction about the claim’s factuality by aggregating the strength of the stances, while taking the reliability of the source into account. Moreover, a fact checking system should be able to explain its decision by providing relevant extracts (rationales) from the documents. Yet, this setup is not directly supported by existing datasets, which treat fact checking, document retrieval, source credibility, stance detection and rationale extraction as independent tasks. In this paper, we support the interdependencies between these tasks as annotations in the same corpus. We implement this setup on an Arabic fact checking corpus, the first of its kind.

## 1 Introduction

Fact checking has recently emerged as an important research topic due to the unprecedented amount of fake news and rumors that are flooding the Internet in order to manipulate people’s opinions (Darwish et al., 2017a; Mihaylov et al., 2015a,b; Mihaylov and Nakov, 2016) or to influence the outcome of major events such as political elections (Lazer et al., 2018; Vosoughi et al., 2018). While the number of organizations performing fact checking is growing, these efforts cannot keep up with the pace at which false claims are being produced, including also clickbait (Karadzhov et al., 2017a), hoaxes (Rashkin et al., 2017), and satire (Hardalov et al., 2016). Hence, there is need for automatic fact checking.

\*This work was carried out when the authors were scientists at QCRI, HBKU.

While most previous research has focused on English, here we target Arabic. Moreover, we propose some guidelines, which we believe should be taken into account when designing fact-checking corpora, irrespective of the target language.

Automatic fact checking typically involves retrieving potentially relevant documents (news articles, tweets, etc.), determining the stance of each document with respect to the claim, and finally predicting the claim’s factuality by aggregating the strength of the different stances, taking into consideration the reliability of the documents’ sources (news medium, Twitter account, etc.). Despite the interdependency between *fact checking* and *stance detection*, research on these two problems has not been previously supported by an integrated corpus. This is a gap we aim to bridge by retrieving documents for each claim and annotating them for stance, thus ensuring a natural distribution of the stance labels.

Moreover, in order to be trusted by users, a fact-checking system should be able to explain the reasoning that led to its decisions. This is best supported by showing extracts (such as sentences or phrases) from the retrieved documents that illustrate the detected stance (Lei et al., 2016). Unfortunately, existing datasets do not offer manual annotation of sentence- or phrase-level supporting evidence. While deep neural networks with attention mechanisms can infer and extract such evidence automatically in an unsupervised way (Parikh et al., 2016), potentially better results can be achieved when having the target sentence provided in advance, which enables supervised or semi-supervised training of the attention. This would allow not only more reliable evidence extraction, but also better stance prediction, and ultimately better factuality prediction. Following this idea, our corpus also identifies the most relevant stance-marking sentences.

## 2 Related Work

The connection between fact checking and stance detection has been argued for by Vlachos and Riedel (2014), who envisioned a system that (i) identifies factual statements (Hassan et al., 2015; Gencheva et al., 2017; Jaradat et al., 2018), (ii) generates questions or queries (Karadzhov et al., 2017b), (iii) creates a knowledge base using information extraction and question answering (Ba et al., 2016; Shiralkar et al., 2017), and (iv) infers the statements’ veracity using text analysis (Banerjee and Han, 2009; Castillo et al., 2011; Rashkin et al., 2017) or information from external sources (Popat et al., 2016; Karadzhov et al., 2017b; Popat et al., 2017). This connection has been also used in practice, e.g., by Popat et al. (2017); however, different datasets had to be used for stance detection vs. fact checking, as no dataset so far has targeted both.

*Fact checking* is very time-consuming, and thus most datasets focus on claims that have been already checked by experts on specialized sites such as *Snopes* (Ma et al., 2016; Popat et al., 2016, 2017), *PolitiFact* (Wang, 2017), or *Wikipedia hoaxes* (Popat et al., 2016).<sup>1</sup> As fact checking is mainly done for English, non-English datasets are rare and often unnatural, e.g., translated from English, and focusing on US politics.<sup>2</sup> In contrast, we start with claims that are not only relevant to the Arab world, but that were also originally made in Arabic, thus producing the first publicly available Arabic fact-checking dataset.

*Stance detection* has been studied so far disjointly from fact checking. While there exist some datasets for Arabic (Darwish et al., 2017b), the most popular ones are for English, e.g., from SemEval-2016 Task 6 (Mohammad et al., 2016) and from the Fake News Challenge (FNC).<sup>3</sup> Despite its name, the latter has no annotations for factuality, but consists of article-claim pairs labeled for stance: *agrees*, *disagrees*, *discusses*, and *unrelated*. In contrast, we retrieve documents for each claim, which yields an arguably more natural distribution of stance labels compared to FNC.

<sup>1</sup>Annotating from scratch is needed in some cases, e.g., in the context of question answering (Mihaylova et al., 2018), or when targeting credibility (Castillo et al., 2011).

<sup>2</sup>See for example the CLEF-2018 lab on Automatic Identification and Verification of Claims in Political Debates, which features US political debates translated to Arabic: <http://alt.qcri.org/clef2018-factcheck/>

<sup>3</sup><http://www.fakenewschallenge.org/>

*Evidence extraction*. Finally, an important characteristic of our dataset is that it provides evidence, in terms of text fragments, for the *agree* and *disagree* labels. Having such supporting evidence annotated enables both better learning for supervised systems performing stance detection or fact checking, and also the ability for such systems to learn to explain their decisions to users. Having this latter ability has been recognized in previous work on rationalizing neural predictions (Lei et al., 2016). This is also at the core of recent research on machine comprehension, e.g., using the SQuAD dataset (Rajpurkar et al., 2016). However, such annotations have not been done for stance detection or fact checking before.

Finally, while preparing the camera-ready version of the present paper, we came to know about a new dataset for Fact Extraction and VERification, or FEVER (Thorne et al., 2018), which is somewhat similar to ours as it is about both factuality and stance, and it has annotation for evidence. Yet, it is also different as (i) the claims are artificially generated by manually altering Wikipedia text, (ii) the knowledge base is restricted to Wikipedia articles, and (iii) the stance and the factuality labels are identical, assuming that Wikipedia articles are reliable to be able to decide a claim’s veracity. In contrast, we use real claims from news outlets, we retrieve articles from the entire Web, and we keep stance and factuality as separate labels.

## 3 The Corpus

Our corpus contains claims labeled for factuality (*true* vs. *false*). We associate each claim with several documents, where each claim-document pair is labeled for stance (*agree*, *disagree*, *discuss*, or *unrelated*) similar to the FakeNewsChallenge (FNC) dataset. Overall, the process of corpus creation went through several stages – *claim extraction*, *evidence extraction* and *stance annotation* –, which we describe below.

**Claim Extraction** We consider two websites as the source of our claims. VERIFY<sup>4</sup> is a project that was established to expose false claims made about the war in Syria and other related Middle Eastern issues. It is an independent platform that debunks claims made by all parties to the conflict. To the best of our knowledge, this is the only platform that publishes fact-checked claims in Arabic.

<sup>4</sup><http://www.verify-sy.com>

It is worth noting that the VERIFY website only shows claims that were debunked as false and misleading, and hence we used it to extract only the *false* claims for our corpus (we extracted the *true* claims from a different source; see below).

We thoroughly preprocessed the original claims. First, we manually identified and excluded all claims discussing falsified multimedia (images or video), which cannot be verified using textual information and NLP techniques only, e.g.

(1) *Pro-regime pages have circulated pictures of fighters fleeing an explosion.*

نشرت صفحات مَوالية لِلنظام صور  
لِقَاتِلين يهربون من انفجار

Note that the claims in VERIFY were written in a form that presents the *corrected* information after debunking the original false claim. For instance, the original false claim in example 2a is corrected and published in VERIFY as shown in example 2b. We manually rendered these corrected claims to their original false form, which we used for our corpus.

(2a) (original false claim) *FIFA intends to investigate the game between Syria and Australia.*

الفيفا تعتزم التحقيق في المباراة بين  
سوريا وأستراليا

(2b) (corrected claim in VERIFY) *FIFA does not intend to investigate the game between Syria and Australia, as pro-regime pages claim.*

الفيفا لا تعتزم التحقيق في المباراة بين  
سوريا وأستراليا، كما تدعي صفحات  
مَوالية لِلنظام.

After extracting the false claims from VERIFY, we collected the true claims of our corpus from REUTERS<sup>5</sup> by extracting headlines of news documents. We used a list of manually selected keywords to extract claims with the same topics as those extracted from VERIFY.

<sup>5</sup><http://ara.reuters.com>

Then, we manually excluded claims that contained political rhetorical statements (see example 3 below), multiple facts, accusations or denials, and ultimately we only kept those claims that discuss factual events, i.e., that can be verified.

(3) *Presidents Vladimir Putin and Recep Tayyip Erdogan hope that Astana talks will lead to peace.*

الرئيسان فلاديمير بوتين ورجب طيب  
اردوغان يأملان بأن محادثات أستانة  
سوف تؤدي إلى السلام.

Overall, starting with 1,381 claims, we ended up with 422 worth-checking claims: 219 *false* claims from VERIFY, and 203 *true* claims from REUTERS.

**Evidence Extraction** Following the assumption that identifying stance towards claims can help predict their veracity, we want to associate each claim with supporting and opposing pieces of textual evidence. We used the Google custom search API for document retrieval, and we performed the following steps to increase the likelihood of retrieving relevant documents. First, as in (Karadzhov et al., 2017b), we transformed each claim into sub-queries by selecting named entities, adjectives, nouns and verbs with the highest TF.DF score, calculated on a collection of documents from the claims' sources. Then, we used these sub-queries with the claim itself as input to the search API and retrieved the first 20 returned links, from which we excluded those directing to VERIFY and REUTERS, and social media websites that are mostly opinionated. Finally, we calculated two similarity measures between the links' content (documents) and the claims: the tri-gram containment (Lyon et al., 2001) and the cosine distance between average word embeddings of both texts.<sup>6</sup> We only kept documents with non-zero values for both measures, yielding 3,042 documents: 1,239 for false claims and 1,803 for true claims.

<sup>6</sup>Word embeddings were generated by training the GloVe (Pennington et al., 2014) model on the Arabic Gigaword (Parker et al., 2011)

**Stance Annotation:** We used CrowdFlower to recruit Arabic speakers to annotate the claim-document pairs for stance. Each pair was assigned to 3–5 annotators, who were asked to assign one of the following standard labels (also used at FNC): *agree*, *disagree*, *discuss* and *unrelated*. First, we conducted small-scale pilot tasks to fine-tune the guidelines and to ensure their clarity. The annotators were also asked to focus on the stance of the document towards the claim, regardless of the factuality of either text. This ensures that stance is captured without bias, so it can be used later with other information (e.g., time, website’s credibility, author reliability) to predict factuality. Finally, the annotators were asked to specify segments in the documents representing the *rationales* that made them assign *agree* or *disagree* as labels. For quality control purposes, we further created a small hidden test set by annotating 50 pairs ourselves, and we used it to monitor the annotators’ performance, keeping only those who maintained an accuracy of over 75%.

Ultimately, we used majority voting to aggregate stance labels for each pair, using the annotators’ performance scores to break ties. On average, 77% of the annotators for each claim-document pair agreed on its label, thus allowing proper majority aggregation for most pairs. A total of 133 pairs with significant annotation disagreement required us to manually check and correct the proposed annotations. We further automatically refined the documents by (i) excluding sentences with more than 200 words, and (ii) limiting the size of a document to 100 sentences. Such extralong documents tend to originate from crawling ill-structured websites, or from parsing some specific types of websites such as web forums.

Table 1 shows the distribution over the stance labels,<sup>7</sup> which turns out to be very similar to that for the FNC dataset. We can see that there are very few documents disagreeing with *true* claims (about 0.5%), which suggests that stance is positively correlated with factuality. However, the number of documents agreeing with *false* documents is larger than the number of documents disagreeing with them, which illustrates one of the main challenges when trying to predict the factuality of news based on stance.

<sup>7</sup>The corpus is available at <http://groups.csail.mit.edu/sls/downloads/> and also at <http://alt.qcri.org/resources/>

Claims	Annotated Documents	Stance (document-to-claim)			
		Agree	Disagree	Discuss	Unrelated
False: 219	1,239	103	82	159	895
True: 203	1,803	371	5	250	1,177
Total: 402	3,042	474	87	409	2,072

Table 1: Statistics about stance and factuality labels.

## 4 Experiments and Evaluation

We experimented with our Arabic corpus, after preprocessing it with ATB-style segmentation using MADAMIRA (Pasha et al., 2014), using the following systems:

- FNC BASELINE SYSTEM. This is the FNC organizers’ system, which trains a gradient boosting classifier using hand-crafted features reflecting polarity, refute, similarity and overlap between the document and the claim.
- ATHENE. It was second at FNC (Hanselowski et al., 2017), and was based on a multi-layer perceptron with the baseline system’s features, word  $n$ -grams, and features generated using latent semantic analysis and other factorization techniques.
- UCL. It was third at FNC (Riedel et al., 2017), training a softmax layer using similarity features.
- MEMORY NETWORK. We also experimented with an end-to-end memory network that showed state-of-the-art results on the FNC data (Mohtarami et al., 2018).

The evaluation results are shown in Table 2. We use 5-fold cross-validation, where all claim-document pairs for the same claim are assigned to the same fold. We report *accuracy*, macro-average  $F_1$ -score, and *weighted accuracy*, which is the official evaluation metric of FNC.

Overall, our corpus appears to be much harder than FNC. For instance, the FNC baseline system achieves weighted accuracy of 75.2 on FNC vs. 55.6 (up to 64.8) on our corpus. We believe that this is because we used a realistic information retrieval approach (see Section 3), whereas the FNC corpus contains a significant number of totally unrelated document–claim pairs, e.g., about 40% of the *unrelated* examples have no word overlap with the claim (even after stemming!), which makes it much easier to correctly predict the *unrelated* class (and this class is also by far the largest).

Model	document Content Used	Weigh.	Acc.	F <sub>1</sub> (macro)	F <sub>1</sub> ( <i>agree, disagree, discuss, unrelated</i> )
Majority class	—	34.8	68.1	20.3	0 / 0 / 0 / 81
FNC baseline system	full document (default)	55.6	72.4	41.0	60.4 / 9.0 / 10.4 / 84
	best sentence	50.5	70.6	37.2	50.3 / 5.4 / 10.3 / 82.9
	best sentence + <b>rationale</b>	60.6	75.6	45.9	73.5 / 13.2 / 11.3 / 85.5
	full document + <b>rationale</b>	64.8	78.4	53.2	84.4 / 32.5 / 8.4 / 87.5
UCL (#3rd in FNC)	full document (default)	49.3	66.0	37.1	47.0 / 7.8 / 13.4 / 80
	best sentence	46.8	66.7	34.7	44.3 / 3.5 / 11.4 / 79.8
	best sentence + <b>rationale</b>	58.5	71.9	44.8	71.6 / 12.6 / 12.4 / 82.6
	full document + <b>rationale</b>	63.7	76.3	51.6	84.2 / 21.4 / 15.3 / 85.3
Athene (#2rd in FNC)	full document (default)	55.1	70.5	41.3	59.1 / 9.2 / 14.1 / 82.3
	best sentence	48.0	67.5	36.1	43.9 / 4.00 / 15.7 / 80.7
	best sentence + <b>rationale</b>	60.6	74.3	48.0	73.5 / 18.2 / 15.9 / 84.6
	full document + <b>rationale</b>	65.5	80.2	55.8	85.0 / 36.6 / 12.8 / 88.8
Memory Network	full document (default)	55.3	70.9	41.6	60.0 / 15.0 / 8.5 / 83.1
	best sentence	52.4	71.0	38.2	58.1 / 8.1 / 4.1 / 82.6
	best sentence + <b>rationale</b>	60.1	75.5	46.4	72.5 / 23.1 / 4.1 / 85.7
	full document + <b>rationale</b>	65.8	79.7	55.2	86.9 / 31.3 / 14.9 / 87.6

Table 2: Performance of some stance detection models from FNC when applied to our Arabic corpus.

Table 2 allows us to study the utility of having gold rationales for the stance (for the *agree* and *disagree* classes only) under different scenarios. First, we show the results when using the full document along with the claim, which is the default representation. Then, we use the best sentence from the document, i.e., the one that is most similar to the claim as measured by the cosine of their average word embeddings. This performs worse, which can be attributed to sometimes selecting the wrong sentence. Next, we experiment with using the rationale instead of the best sentence when applicable (i.e., for *agree* and *disagree*), while still using the best sentence for *discuss* and *unrelated*. This yields sizable improvements on all evaluation metrics, compared to using the best sentence (5-12 point absolute) or the full document (3-9 points absolute). We further evaluate the impact of using the rationales, when applicable, but using the full document otherwise. This setting performed best (80.2% accuracy with ATHENE, and 3-8 points of improvement over best+rationale), as it has access to most information: full document + rationale.

Overall, the above experiments demonstrate that having a gold rationale can enable better learning. However, the results should be considered as a kind of upper bound on the expected performance improvement, since here we used gold rationales at *test* time, which would not be available in a real-world scenario. Still, we believe that sizable improvements would still be possible when using the gold rationales for training only.

Finally, we built a simple fact-checker, where the factuality of a claim is determined based on aggregating the predicted stances (using FNC’s baseline system) of the documents we retrieved for it. This yielded an accuracy of 56.2 when using the full documents, and 59.7 when using the best sentence + rationale (majority baseline of 50.5), thus confirming once again the utility of having a rationale, this time for a downstream task.

## 5 Conclusion and Future Work

We have described a novel corpus that unifies stance detection, stance rationale, relevant document retrieval, and fact checking. This is the first corpus to offer such a combination, not only for Arabic but in general. We further demonstrated experimentally that these unified annotations, and the gold rationales in particular, are beneficial both for stance detection and for fact checking.

In future work, we plan to extend the annotations to cover other important aspects of fact checking such as source reliability, language style, and temporal information, which have been shown useful in previous research (Castillo et al., 2011; Lukasik et al., 2015; Ma et al., 2016; Mukherjee and Weikum, 2015; Popat et al., 2017).

## Acknowledgment

This research was carried out in collaboration between the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) and the HBKU Qatar Computing Research Institute (QCRI).

## References

- Mouhamadou Lamine Ba, Laure Berti-Equille, Kushal Shah, and Hossam M. Hammady. 2016. VERA: A platform for veracity estimation over web data. In *Proceedings of the 25th International Conference Companion on World Wide Web*. Montréal, Canada, WWW '16, pages 159–162.
- Protima Banerjee and Hyoil Han. 2009. Answer credibility: A language modeling approach to answer validation. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Boulder, CO, USA, NAACL-HLT '09, pages 157–160.
- Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on Twitter. In *Proceedings of the 20th International Conference on World Wide Web*. Hyderabad, India, WWW '11, pages 675–684.
- Kareem Darwish, Dimitar Alexandrov, Preslav Nakov, and Yelena Mejova. 2017a. Seminar users in the Arabic Twitter sphere. In *Proceedings of the 9th International Conference on Social Informatics*. Oxford, UK, SocInfo '17, pages 91–108.
- Kareem Darwish, Walid Magdy, and Tahar Zanouda. 2017b. Improved stance prediction in a user similarity feature space. In *Proceedings of the Conference on Advances in Social Networks Analysis and Mining*. Sydney, Australia, ASONAM '17, pages 145–148.
- Pepa Gencheva, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017. A context-aware approach for detecting worth-checking claims in political debates. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Varna, Bulgaria, RANLP '17, pages 267–276.
- Andreas Hanselowski, Avinash PVS, Benjamin Schiller, and Felix Caspelherr. 2017. Team Athene on the fake news challenge. <https://medium.com/@andre134679/team-athene-on-the-fake-news-challenge-28a5cf5e017b>.
- Momchil Hardalov, Ivan Koychev, and Preslav Nakov. 2016. In search of credible news. In *Proceedings of the 17th International Conference on Artificial Intelligence: Methodology, Systems, and Applications*. Varna, Bulgaria, AIMSAS '16, pages 172–180.
- Naeemul Hassan, Chengkai Li, and Mark Tremayne. 2015. Detecting check-worthy factual claims in presidential debates. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. Melbourne, Australia, CIKM '15, pages 1835–1838.
- Israa Jaradat, Pepa Gencheva, Alberto Barrón-Cedeño, Lluís Màrquez, and Preslav Nakov. 2018. Claim-Rank: Detecting check-worthy claims in Arabic and English. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18.
- Georgi Karadzhov, Pepa Gencheva, Preslav Nakov, and Ivan Koychev. 2017a. We built a fake news & click-bait filter: What happened next will blow your mind! In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Varna, Bulgaria, RANLP '17, pages 334–343.
- Georgi Karadzhov, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017b. Fully automated fact checking using external sources. In *Proceedings of the Conference on Recent Advances in Natural Language Processing*. Varna, Bulgaria, RANLP '17, pages 344–353.
- David M.J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science* 359(6380):1094–1096.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Austin, TX, USA, EMNLP '16, pages 107–117.
- Michal Lukasik, Trevor Cohn, and Kalina Bontcheva. 2015. Point process modelling of rumour dynamics in social media. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. Beijing, China, ACL-IJCNLP '15, pages 518–523.
- Caroline Lyon, James Malcolm, and Bob Dickerson. 2001. Detecting short passages of similar text in large document collections. In *Proceedings of Conference on Empirical Methods in Natural Language Processing*. Pittsburgh, PA, USA, EMNLP '01, pages 118–125.
- 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*. New York, NY, USA, IJCAI '16, pages 3818–3824.
- Todor Mihaylov, Georgi Georgiev, and Preslav Nakov. 2015a. Finding opinion manipulation trolls in news community forums. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*. Beijing, China, CoNLL '15, pages 310–314.

- Todor Mihaylov, Ivan Koychev, Georgi Georgiev, and Preslav Nakov. 2015b. Exposing paid opinion manipulation trolls. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*. Hissar, Bulgaria, RANLP '15, pages 443–450.
- Todor Mihaylov and Preslav Nakov. 2016. Hunting for troll comments in news community forums. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. Berlin, Germany, ACL '16, pages 399–405.
- Tsvetomila Mihaylova, Preslav Nakov, Lluís Marquez, Alberto Barron-Cedeno, Mitra Mohtarami, Georgi Karadzhov, and James Glass. 2018. Fact checking in community forums. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. New Orleans, LA, USA, AAAI '18.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiao-Dan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of the International Workshop on Semantic Evaluation*. Berlin, Germany, SemEval '16, pages 31–41.
- Mitra Mohtarami, Ramy Baly, James Glass, Preslav Nakov, Lluís Marquez, and Alessandro Moschitti. 2018. Automatic stance detection using end-to-end memory networks. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18.
- Subhabrata Mukherjee and Gerhard Weikum. 2015. Leveraging joint interactions for credibility analysis in news communities. In *Proceedings of the 24th ACM Conference on Information and Knowledge Management*. Melbourne, Australia, CIKM '15, pages 353–362.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Austin, TX, USA, EMNLP '16, pages 2249–2255.
- Robert Parker, David Graff, Ke Chen, Junbo Kong, and Kazuaki Maeda. 2011. Arabic Gigaword Fifth Edition, LDC2011T11. Web Download. Philadelphia: Linguistic Data Consortium.
- Arfath Pasha, Mohamed Al-Badrashiny, Mona T Diab, Ahmed El Kholly, Ramy Eskander, Nizar Habash, Manoj Pooleery, Owen Rambow, and Ryan Roth. 2014. MADAMIRA: A Fast, Comprehensive Tool for Morphological Analysis and Disambiguation of Arabic. In *Proceedings of the Conference on Language Resources and Evaluation*. Reykjavik, Iceland, LREC '14, pages 1094–1101.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the conference on empirical methods in natural language processing*. Doha, Qatar, EMNLP '14, pages 1532–1543.
- Kashyap Papat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2016. Credibility assessment of textual claims on the web. In *Proceedings of the International on Conference on Information and Knowledge Management*. Indianapolis, IN, USA, CIKM '16, pages 2173–2178.
- Kashyap Papat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the truth lies: Explaining the credibility of emerging claims on the web and social media. In *Proceedings of the Conference on World Wide Web Companion*. Perth, Australia, WWW '17, pages 1003–1012.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, TX, USA, EMNLP '16, pages 2383–2392.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. EMNLP '17, pages 2931–2937.
- Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the Fake News Challenge stance detection task. *ArXiv:1707.03264*.
- Prashant Shiralkar, Alessandro Flammini, Filippo Menczer, and Giovanni Luca Ciampaglia. 2017. Finding streams in knowledge graphs to support fact checking. In *Proceedings of the IEEE International Conference on Data Mining*. New Orleans, LA, USA, ICDM '17.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: A large-scale dataset for fact extraction and verification. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18.
- Andreas Vlachos and Sebastian Riedel. 2014. Fact checking: Task definition and dataset construction. In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*. Baltimore, MD, USA, pages 18–22.
- 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*. Vancouver, Canada, ACL '17, pages 422–426.