

EXPLAINABLE FACT-CHECKING THROUGH QUESTION ANSWERING

Jing Yang¹, Didier Vega-Oliveros¹, Taís Seibt², and Anderson Rocha¹

¹ Artificial Intelligence Lab. (Recod.ai), Institute of Computing, University of Campinas, SP, Brazil

² Department of Journalism, Universidade do Vale do Rio dos Sinos (Unisinos), Porto Alegre, RS, Brazil

ABSTRACT

Misleading or false information has been creating chaos in some places around the world. To mitigate this issue, many researchers have proposed automated fact-checking methods to fight the spread of fake news. However, most methods cannot explain the reasoning behind their decisions, failing to build trust between machines and humans using such technology. Trust is essential for fact-checking to be applied in the real world. Here, we address fact-checking explainability through question answering. In particular, we propose generating questions and answers from **claims** and answering the same questions from **evidence**. We also propose an answer comparison model with an attention mechanism attached to each question. Leveraging question answering as a proxy, we break down automated fact-checking into several steps — this separation aids models’ explainability as it allows for more detailed analysis of their decision-making processes. Experimental results show that the proposed model can achieve state-of-the-art performance while providing reasonable explainable capabilities¹.

Index Terms— Fact-checking, question answering, model explainability, misinformation

1. INTRODUCTION

The Internet has become an extension of our physical world, whereby almost everyone is connected. As a result, one small piece of false news can spread worldwide and ruin a person, company, or even a country’s economy or reputation. To mitigate the impact of fake news, recognizing and flagging them is necessary due to the amount of news released.

In the research community, many researchers have devoted efforts to studying automatic fake news detection. Among them, automated fact-checking has attracted a great deal of attention [1, 2]. The task is to check if a claim is factually correct based on evidence retrieved from reliable sources. However, according to a recent survey [3], human fact-checkers generally do not trust automated solutions. Some works have been proposed to build the bridge between humans and machines. For example, Yang et al. [4] proposed a work to summarize claims for more scalable fact-checking and involved human-in-the-loop to evaluate summarization results.

Another popular line of work to increase trust for fact-checking is to generate explanations for the predicted results. Atanasova et al. [5] first proposed a pioneer work to generate explanations. They performed an optimization to learn together veracity prediction and explanation extraction from evidence. Subsequently, Kotonya et al. [6] proposed a joint extractive and abstractive text summarization

method for explanation generation. The authors also published a survey specifically about generating fact-checking explanations [7].

However, although generating explanations can provide more precise evidence to understand fact-checking decisions, existing systems lack a way to evaluate the explanations properly. Especially for explanations based on abstractive document summarization, researchers have shown that such models have problems of hallucination [8, 9], generating summaries factually inconsistent with their original document. To deal with this issue, several works have been proposed [10, 11, 12]. In particular, Pagnoni et al. [13] summarized different types of errors some models make and metrics used to evaluate them. Among these evaluation metrics, leveraging question answering (QA) as a proxy has been the focus of some work [11, 12]. The idea is to rely upon a question answering mechanism as an evaluation for the faithfulness of summaries. Wang et al. [12] extracted answers and questions from summaries and fine-tuned a QA model to generate answers from the documents; the answers from the document and its summary for the same questions are compared to determine the actual consistency of the summary. Recently, Nan et al. [14] proposed an improved method than [12], where instead of generating answers for both the summary and document, they model the likelihood of the summary and document conditioned on question-answer pairs generated from the summaries. Through this, the likelihood metric becomes suitable as a training objective to improve the factual consistency of summaries.

A few works have been proposed to leverage QA to help in fact-checking. For example, in PathQG [15], Wang et al. generated questions from facts. They accomplished this task in two steps: first, they identified facts from an input text to build a knowledge graph (KG) and then generated an ordered sequence as a query path; second, they utilized a seq2seq model to learn to generate questions based on the query path. The human evaluation showed that their model could generate informative questions. In another work, Fan et al. [16] generated question-answer pairs as a type of brief, along with passage brief and entity brief, and provided them to the human fact-checkers, aiming at improving their checking efficiency.

Inspired by the QA works in checking factual consistency of documents and their summaries, we believe it is suitable for the fact-checking task, where we assess if claims are factually consistent with retrieved evidence. Therefore, we propose to leverage automated QA protocols and integrate them into the traditional fact-checking pipeline. As a result, we can provide explainable fact-checking results through question answering. The answer comparison model will predict a label and pinpoint the wrong part of a claim by showing which questions are more important for the decision. In this way, humans fact-checker can easily interpret the results and correct them if necessary. Our work differs from prior works [15, 16] because we not only generate question-answer pairs but also fully integrate QA protocols in the fact-checking pipeline to automatically compare answers and predict their labels. We compare the proposed method

Research funded by the São Paulo Research Foundation (FAPESP) under the Grants DéjàVu #2017/12646-3, #2019/04053-8 and #2019/26283-5.

¹Our implementation is available at: <https://bit.ly/3mtSUE3>

with several baselines, achieving state-of-the-art results but with the critical feature of adding explainability to the fact-checking process.

We summarize our contributions as follows.

- We propose a novel pipeline for using question answering as a proxy for explainable fact-checking;
- We introduce an answer comparison model with an attention mechanism on questions to learn their importance on the claims;

2. PROPOSED METHODOLOGY

We introduce question answering (QA) in the fact-checking process. Despite previous mentions of using QA for fact-checking, no previous work has explored integrating QA protocols in its pipeline. Our proposed solution is described as follows:

- 1) Given a claim C , generate multiple questions Q_1, \dots, Q_n and answers A_1^C, \dots, A_n^C from it;
- 2) Retrieve and re-rank evidence E based on the claim (and possibly questions);
- 3) For each question generated from 1), ask retrieved evidence for answers A_1^E, \dots, A_n^E respectively;
- 4) Compare the answer pairs (A_i^C, A_i^E) and transform the result into a label of SUPPORTS or REFUTES.

Our proposed pipeline leads to more explainability as we break down the fact-checking process into more steps, allowing a more fine-grained analysis of each part of the process (e.g., question generation, question answering, or answer comparison). In addition, through answer generation from claims and evidence, we vastly reduce the information (from claims and their evidence to only answer pairs) fed to the final classification model. Thus, the model learns from more direct and precise inputs.

To focus on how question answering empowers explainability, we use gold evidence instead of retrieved evidence. It means that for step 2), we take the gold evidence directly instead of retrieving them to focus on evaluating the other three stages of the problem. Future work will be dedicated to the retrieval by itself. Therefore, we focus on steps 1), 3), and 4) of the pipeline. Next, we detail the proposed methodology steps.

2.1. Question and Answer Generation

To generate questions from a text, answers for the text are usually provided first to generate more relevant questions [11, 12]. Answers are usually extracted based on named entities and noun phrases; then, questions are generated given the claim and answers. They can also be generated in parallel with questions [14]. We adopt the approach to generate questions and answers from claims simultaneously [14]. In particular, we follow the instruction of [14] to fine-tune the BART-large model to generation question-answer pairs $(Q_1, A_1^C), \dots, (Q_n, A_n^C)$ from a given claim C . Using beam search, 64 question-answer pairs are generated, then pairs are removed if the claim does not contain the answers. For answers of evidence E , we utilize a pre-trained extractive QA model to answer the questions generated previously from the claim. The model generates multiple answers, and we choose the one with the highest score (the most likely answer).

2.2. Answer Pair Comparison

For answer comparison, the token-level F1 score is usually used to measure similarity between answer pairs; however, it does not work

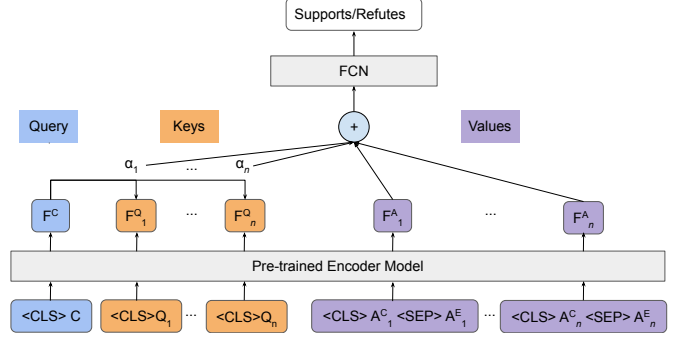


Fig. 1. Answer comparison model with attention on questions. C represents a given claim, Q_i represents i_{th} questions, and (A_i^C, A_i^E) represents i_{th} answer pairs for claim and evidence. n denotes the number of questions and answer pairs.

when the two answers have non-overlapping words but are semantically similar. We propose to fine-tune a transformer model to learn answer comparison. Considering that different questions have various purposes, thus also vary in their importance. To account for this, we add attention to each question to learn the importance weights. The structure of the model is shown in Fig. 1.

Specifically, we rely on a pre-trained masked language model to encode the claim C , questions Q_1, \dots, Q_n , and answer pairs $(A_1^C, A_1^E), \dots, (A_n^C, A_n^E)$. For the answer pairs, we add a $\langle \text{SEP} \rangle$ token between two answers of the same question. We use one encoder model for encoding all the inputs, which means the weights are shared. After the encoding, we take the representation of the $\langle \text{CLS} \rangle$ token as each sentence embedding, thus transforming the claim, questions and answer pairs into features: $F_C, F_1^Q, F_2^Q, \dots, F_n^Q$, and $F_1^A, F_2^A, \dots, F_n^A$ respectively, where n is the number of questions for each claim. Then we utilize additive attention proposed by Bahdanau et al. [17] to learn the importance of each question. We treat the claim as a query, questions as keys, and answers as values for each representation. The details are formulated as follows.

$$f_{att}(F^C, F^{Q_j}) = \mathbf{W}_3 \tanh(\mathbf{W}_1 F^C + \mathbf{W}_2 F^{Q_j}) \quad (1)$$

$$a_j = \text{softmax}(f_{att}(F^C, F^{Q_j})) \quad (2)$$

$$\mathbf{F} = \sum_j a_j F_j^A \quad (3)$$

where f_{att} calculate the attention weight between F^C and F_j^Q ($j = 1, 2, \dots, n$), $\mathbf{W}_1, \mathbf{W}_2$ and \mathbf{W}_3 are learnable parameters.

In Eq. (1) and (2), attention weights are calculated and normalized by the softmax function. Then Eq. (3) uses a weighted sum to combine all answer features to the final feature F . Feature F is then fed to a fully connected layer to have the final prediction of SUPPORTS or REFUTES. Notice that all the information present in both claim and evidence are reduced into their respective answers. The claim and generated questions take part in selecting the most relevant answers, assigning higher weights to them.

3. EXPERIMENTAL SETUP

3.1. Dataset

We adopt the Fool-Me-Twice (FM2) dataset, which comprises 12,968 claims and their associated evidence. FM2 is a recently

published dataset collected through a multi-player game. In the game, one player generates a claim and tries to fool other players. The others have to decide if the claim is true or false based on evidence retrieved by the game before a timer runs out. The game’s setting makes this dataset challenging as the players are motivated to generate claims hard to verify. FM2 a more difficult and less biased dataset than the seminal dataset FEVER [1], in which a model can exploit specific words from the claim [18] to achieve reasonable accuracy (79.1% for two classes). In contrast, FM2 is shown not to have biases, a classification based only on claims resulted in low prediction accuracy (61.9%).

3.2. Implementation details

For question-answer pairs generation, we follow the code provided in [16] to fine-tune a BART-large model based on XSUM and CN-NDM datasets². For answer generation for evidence, we use the FARM framework from deepset³ to generate answers from evidence, and the question-answering model is *deepset/electra-base-squad2*. For answer comparison, we use *microsoft/mpnet-base* model for encoding all input representation, as it has shown to perform well in question answering tasks [19]. As the question generation model does not output the same number of questions for every claim, we selected the first ten questions if the claim has more than 10; if the number of questions is less than 10, we repeat the first question until 10. We choose this quantity of questions because the average number of questions for each claim is 11.5. The hyperparameters for training the answer comparison models are: *number of epochs* = 5, *batch size* = 32, *learning rate* = 2e-5, which is the standard for fine-tuning a masked language model, and *maximum token length* = 32. For statistical significance, we run each experiment 5 times and report the average and standard deviation. As the dataset is well-balanced, we use macro average accuracy as the evaluation metric.

3.3. Baselines

We set questions and answers for the baselines to be the same, only varying different answer comparison methods.

- **Blackbox method:** we compare our results with the original proposed method in [20]. We refer to it as the *black box method* as they concatenate claim and evidence for the prediction without providing interpretability. We used the code provided by the authors⁴ and ran it five times to have an average result.
- **QUALS score:** it is an automatic metric for checking factual consistency [14]. It does not generate answers for evidence. Instead, it calculates the likelihood of the evidence given the question-answer pair from the claim, compromising explainability.
- **Token level F1-score:** a standard metric for question-answer tasks. It counts words overlap between two answers.
- **BERTscore:** a common metric for measuring the similarity of two sentences. We use the default model *roberta-large*.
- **Cosine similarity:** a metric also used for measuring sentence similarity. We use sentence transformer *all-mpnet-base-v2* to embed the answers and calculate the cosine similarities between the embeddings.

²<https://bit.ly/3iBZyqR>

³<https://github.com/deepset-ai/FARM>

⁴<https://bit.ly/2Z06Ctr>

Table 1. Fact-checking label accuracy of different methods. ‘X-AI’ denotes Explainability capabilities.

Methods	Dev Acc	Test Acc
Blackbox (No X-AI)	76.17±1.23	74.58±1.66
QUALS (th=-1.2)	56.12	56.01
BERTscore (th=0.843)	58.68	62.32
cosine sim (th=0.305)	61.16	62.75
F1-score (th=0.06)	64.07	63.77
Attention C-Q-AA (ours, X-AI)	75.44±0.52	73.43±0.83

th: threshold

Only the black box method requires training. The others are metrics to evaluate the answer pairs. These metrics calculate a score representing similarity for each answer pair, except for QUALS that outputs a score for all answers of the same claim. As each claim has several questions, we compute the average score for the claim and provide a threshold to convert the score to a binary label.

4. RESULTS AND ANALYSIS

4.1. Comparison with baselines

We show the results with different baselines in Table 1. For the metric-based methods, we do a binary search to find the highest accuracy on the development set for the threshold selection.

The results show that training an answer comparison model specifically for the fact-checking task improves accuracy compared with the methods without training. Our attention-based method achieves slightly lower accuracy than the black box method. However, our method is more suitable for real-world applications than the black box one because: 1) our method essentially reduces the input needed for prediction while remaining almost the same accuracy, 2) we enable error analysis for fact-checking with several steps, and 3) our model additionally provides more explainability by learning the importance of each question.

4.2. Attention visualization

To illustrate how attention helps explainability, we show an example of our generated questions with attention weights and their answers from claim and evidence in Fig. 2. The question with the highest weight is bold, and the second-highest underlined. Although some answers are incorrect and there are non-matching answer pairs, the model can attend more on the questions and answers more relevant to the factuality of the claim, showing our approach’s potential. We also see that because the claim is short, most questions are repetitive.

4.3. Ablation study

We carry out an ablation study to show if our attention mechanism improves performance compared with simple classification. Thus we remove the attention layer of our proposed attention model, the network structure is shown in Fig. 3. Specifically, to use all available questions, we concatenate all questions and all answers: so the model has two inputs $\langle \text{CLS} \rangle C \langle \text{SEP} \rangle Q_1 Q_2 \dots Q_n$, and $\langle \text{CLS} \rangle A_1^C A_2^C \dots A_n^C \langle \text{SEP} \rangle A_1^E A_2^E \dots A_n^E$ (note here n can be different for different claims). As the inputs are concatenated, the maximum token length here is 128. Then through the embedding model, each input is transformed into a feature vector, and the two vectors are concatenated to be fed into the classification layer.

To study the effect of different components of our proposed model, we design the inputs as follows:

Question	Answer for claim	Answer for evidence	Attention weight (%)
How old was MacPhail when he died?	92	95	12.17
Where did MacPhail die?	at his home	Delray Beach, Florida	4.29
Who died at the age of 92?	MacPhail	no_answer	1.66
When did MacPhail die?	at the age of 92	November 8, 2012	3.23
What was MacPhail's age?	92	95	13.19
<u>At what age did MacPhail die?</u>	<u>92</u>	<u>95</u>	<u>18.87</u>
Where did MacPhail die?	his home	Delray Beach, Florida	4.29
What age was MacPhail when he died?	92	95	20.67
<u>At what age did MacPhail die?</u>	<u>92</u>	<u>95</u>	<u>18.87</u>
Who died at age 92?	MacPhail	no_answer	2.75

Predicted label: REFUTES **Gold label:** REFUTES

Fig. 2. An example of our model generated questions, answer pairs, and attention weights. The question with the highest weight is in bold, and the second highest is underlined.




Fig. 3. Answer comparison model without attention on questions. C represents a given claim, Q_i represents i_{th} questions, and (A^C_i, A^E_i) represents i_{th} answer pairs for claim and evidence. n denotes the number of questions for claim C .

- C: only claims $\langle \text{CLS} \rangle C$;
- Q: only concatenated questions $\langle \text{CLS} \rangle Q_1 Q_2 \dots Q_n$;
- AA: only answer pairs $\langle \text{CLS} \rangle A^C_1 A^C_2 \dots A^C_n \langle \text{SEP} \rangle A^E_1 A^E_2 \dots A^E_n$;
- Q-AA: concatenated questions $\langle \text{CLS} \rangle Q_1 Q_2 \dots Q_n$ and answer pairs $\langle \text{CLS} \rangle A^C_1 A^C_2 \dots A^C_n \langle \text{SEP} \rangle A^E_1 A^E_2 \dots A^E_n$;
- CQ-AA: our full model without attention (shown in Fig. 3).
- Attention C-Q-AA: our full model with attention.

From the ablation study, we want to know how much each input affects the model’s performance. In Table 2, we can see that with C or Q only, the model cannot perform well, indicating that the model can not rely solely on the claim information to achieve high accuracy. Our result agrees with the original paper, in which the model only with claims achieved an accuracy of 61.9%. Also, when adding Q and CQ information to AA, Q-AA and CQ-AA perform slightly better. This indicates that the model can learn most of the information from answer pairs only. It is reasonable because the answer pairs carry most of the critical information from both claims and evidence. Comparing CQ-AA with our attention-based C-Q-AA, we see that the attention mechanism can help increase performance because it uses claims and questions to weigh up essential answer pairs.

Table 2. Ablation study of the model without attention

Inputs	Dev Acc	Test Acc
C	59.15±1.22	61.57±1.67
Q	56.22±1.37	56.90±0.90
AA	74.15±1.33	72.61±1.04
Q-AA	74.46±0.80	72.62±1.59
CQ-AA	74.88±0.81	72.89±1.14
Attention C-Q-AA	75.44±0.52	73.43±0.83

4.4. Limitations

• **Question Generation.** Generating diverse and relevant questions aiming at the factuality of a claim is challenging. Claims can be altered by changing the subject, object, time, place, actions, or even multiple editions together. In some cases, we observed that the questions have the problem of not recognizing complete phrases of the claim, and sometimes most questions of a claim are semantically similar because the claim is too short. For example in Fig. 1, we can see that most of the questions are paraphrases. Hence, better ways of generating questions and filtering less relevant and repetitive questions are needed to improve performance.

• **Question Answering.** Answering correctly giving the context is a non-trivial and crucial step in the pipeline. Unfortunately, state-of-the-art models can fail to answer correctly in some cases, as they require reasoning and logical thinking to calculate the correct answer from the context. We show a failing example here: **Evidence:** *Weber was born in Eutin, Bishopric of Lübeck, the eldest of the three children of Franz Anton von Weber and his second wife, Genovefa Weber; a Viennese singer.* **Question:** *How many siblings did Albert Weber have?* **Answer for evidence:** *three.*

In the example, the model is not able to give the correct answer *–two*, because it is an extractive QA model, which is a limitation of this type of model. Nevertheless, the explainability provide by questions and answers gives us a better idea of which part is wrong in the claim and what could help us improve the model.

Reasoning over text is a very challenging task; other ways of transforming the claim into a format like tabular data [21] may also help simplify the reasoning and thus improve performance.

5. CONCLUSION

In this paper, we proposed a novel pipeline for using QA as a proxy for fact-checking. Based on this pipeline, we proposed an answer comparison model with an attached attention mechanism, which learns to attend critical questions with interpretability capabilities.

Our ablation study showed that the model can achieve near state-of-the-art performance with only information from answer pairs. Thus, using QA, we can encourage the model to learn from more precise evidence; this can aid fact-checkers in better understanding models’ decisions. Then, when necessary, they can compare the answers and make decisions for themselves.

In future work, we plan to add the retrieval step to the pipeline instead of using gold evidence, as the retrieval is also a crucial part of fact-checking. We can also instead answer questions directly from a more extensive set of document evidence. In addition, we plan to work on more datasets to address the generalization capabilities of the method. Finally, we intend to have human evaluations on the questions and answers to improve the generation and potentially help human fact-checkers by providing high-quality QAs.

6. REFERENCES

- [1] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal, “FEVER: a large-scale dataset for fact extraction and verification,” in *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2018.
- [2] Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu, “Fine-grained fact verification with kernel graph attention network,” in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [3] Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino, “Automated fact-checking for assisting human fact-checkers,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [4] Jing Yang, Didier Vega-Oliveros, Tais Seibt, and Anderson Rocha, “Scalable fact-checking with human-in-the-loop,” *arXiv preprint*, 2021.
- [5] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein, “Generating fact checking explanations,” in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [6] Neema Kotonya and Francesca Toni, “Explainable automated fact-checking for public health claims,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [7] Neema Kotonya and Francesca Toni, “Explainable automated fact-checking: A survey,” in *International Conference on Computational Linguistics (ACL)*, 2020.
- [8] Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li, “Faithful to the original: Fact aware neural abstractive summarization,” in *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [9] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald, “On faithfulness and factuality in abstractive summarization,” in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [10] Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher, “Evaluating the factual consistency of abstractive text summarization,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [11] Esin Durmus, He He, and Mona Diab, “FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization,” in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [12] Alex Wang, Kyunghyun Cho, and Mike Lewis, “Asking and answering questions to evaluate the factual consistency of summaries,” in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [13] Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov, “Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics,” in *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2021.
- [14] Feng Nan, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Kathleen McKeown, Ramesh Nallapati, Dejiro Zhang, Zhiguo Wang, Andrew O Arnold, and Bing Xiang, “Improving factual consistency of abstractive summarization via question answering,” in *Joint Conference of Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, 2021.
- [15] Siyuan Wang, Zhongyu Wei, Zhihao Fan, Zengfeng Huang, Weijian Sun, Qi Zhang, and Xuan-Jing Huang, “PathQG: Neural question generation from facts,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [16] Angela Fan, Aleksandra Piktus, Fabio Petroni, Guillaume Wenzek, Marzieh Saeidi, Andreas Vlachos, Antoine Bordes, and Sebastian Riedel, “Generating fact checking briefs,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [17] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [18] Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay, “Towards debiasing fact verification models,” in *Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [19] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu, “MPNet: Masked and permuted pre-training for language understanding,” *arXiv preprint*, 2020.
- [20] Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger, and Jordan Boyd-Graber, “Fool me twice: Entailment from wikipedia gamification,” in *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2021.
- [21] Vivek Gupta, Riyaz A Bhat, Atreya Ghosal, Manish Srivastava, Maneesh Singh, and Vivek Srikumar, “Is my model using the right evidence? systematic probes for examining evidence-based tabular reasoning,” *arXiv preprint*, 2021.