

Don't Discard All the Biased Instances: Investigating a Core Assumption in Dataset Bias Mitigation Techniques

Hossein Amirkhani

University of Qom
Tehran Institute for Advanced Studies
amirkhani@qom.ac.ir

Mohammad Taher Pilehvar

Tehran Institute for Advanced Studies,
Khatam University, Iran
mp792@cam.ac.uk

Abstract

Existing techniques for mitigating dataset bias often leverage a *biased* model to identify biased instances. The role of these biased instances is then reduced during the training of the *main* model to enhance its robustness to out-of-distribution data. A common core assumption of these techniques is that the main model handles biased instances similarly to the biased model, in that it will resort to biases whenever available. In this paper, we show that this assumption does not hold in general. We carry out a critical investigation on two well-known datasets in the domain, MNLI and FEVER, along with two biased instance detection methods, partial-input and limited-capacity models. Our experiments show that in around a third to a half of instances, the biased model is unable to predict the main model's behavior, highlighted by the significantly different parts of the input on which they base their decisions. Based on a manual validation, we also show that this estimate is highly in line with human interpretation. Our findings suggest that down-weighting of instances detected by bias detection methods, which is a widely-practiced procedure, is an unnecessary waste of training data. We release our code to facilitate reproducibility and future research.¹

1 Introduction

Several studies suggest that the impressive performance of the recent natural language understanding models might not be fully indicative of the status quo in language understanding. This is partially due to the artificial nature of the evaluation datasets which usually contain simple superficial cues, such as specific keywords or sentence structures, that spuriously correlate with the gold labels (Jabri et al., 2016; Jia and Liang, 2017; Gururangan et al., 2018; McCoy et al., 2019; Wiegand et al., 2019;

Schuster et al., 2019). Such shallow patterns can be easily exploited by the model, resulting in an overestimated performance on the specific dataset and usually poor performance on other differently-constructed datasets.

Many proposals have been put forward to enhance robustness to such dataset-specific biases. Some techniques rely on the sources of bias which are known a-priori for each dataset (He et al., 2019; Clark et al., 2019; Mahabadi et al., 2020). Others alleviate this requirement by identifying biased examples without explicitly modeling them, for instance by training a low-capacity model that does not go much beyond the surface patterns (Sanh et al., 2020; Utama et al., 2020; Clark et al., 2020). The *main* model is then trained in a way not to rely much on the instances detected by the *biased* model. This is done either implicitly using the popular *product of experts* approach (Clark et al., 2019; Sanh et al., 2020; Utama et al., 2020; Mahabadi et al., 2020) or explicitly using methods such as *debaised focal loss* (Mahabadi et al., 2020) or *example reweighting* (Utama et al., 2020).

Irrespective of the approach for identifying biased instances, existing bias mitigation techniques share the assumption that the main model treats the biased instances similarly to the biased model, in that it will resort to superficial biased features in case they exist. We carry out a critical investigation of the validity of this assumption. Through a set of experiments we show that for a significant subset of instances, the biased model is unable to predict the main model's behavior. Specifically, we find that in around a third to a half of instances in two well-known datasets in the domain, MultiNLI (Williams et al., 2018, MNLI) and (Thorne et al., 2018, FEVER), the main model handles the input differently from two popular biased instance detection methods, partial-input and limited-capacity models. We further support this estimate through a manual validation. This highlights the need for

¹<https://github.com/h-amirkhani/debiasing-assumption>

re-thinking the discarding of instances detected by biased models, which is a widely-adopted approach in the current dataset bias mitigation techniques.

2 Methodology

Following the terminology used in the literature, we refer to the model used for detecting biased instances as *biased* model, the final intended model (for which we are mitigating dataset biases) as the *main* model, and those instances which are correctly classified by both the main and biased models as *easy* instances. Existing bias mitigation techniques try to weaken the role of biased instances in the learning process based on the assumption that on these instances the main model behaves similarly to the biased model. We investigate this assumption by comparing the parts of input (tokens) on which the two models base their decisions. This analysis is carried out on the *easy* instances where the biased model is deemed to exploit superficial features. Specifically, we compare the role of individual input tokens in the two models to check if the main model bases its decision on the same tokens (features) as the biased model.

Datasets. We experimented with two widely-used datasets in the domain: MNLI and FEVER. In the former, the goal is to assign each *premise-hypothesis* pair to one of three classes: *entailment*, *contradiction*, and *neutral*; whereas in the latter, each *evidence* either *supports* or *refutes* the corresponding *claim*, or there is *not enough info* (NEI) for a definite decision. For the MNLI data, we used *validation-matched* as our validation set following Clark et al. (2019); Utama et al. (2020); Yaghoobzadeh et al. (2021). For the FEVER data, we experimented with the version of Schuster et al. (2019) following Mahabadi et al. (2020); Yaghoobzadeh et al. (2021). The dataset statistics are presented in Table A1.

Main model. For all the experiments, we fine-tuned the pre-trained BERT-base-uncased model from the Hugging Face Transformers library (Wolf et al., 2019) as our main model. The hyper-parameters are chosen according to the literature (Sanh et al., 2020).

Biased models. We experimented with two biased instance detection (also loosely called *bias-only*) methods. The first one is the widely-used partial-

input model (Gururangan et al., 2018; Poliak et al., 2018; Schuster et al., 2019), which takes an instance as biased if an incomplete part of it is enough for correct classification. We used the *hypothesis* and *claim* parts as partial inputs in the MNLI and FEVER datasets, respectively. The BERT-base-uncased model was trained with the same hyper-parameters as the main model on this partial-input data. The second method, put forward by Sanh et al. (2020), is based on the observation that models with limited capacity learn to exploit dataset biases. For this, we fine-tuned the pre-trained TinyBERT model (Turc et al., 2019) with the same hyper-parameters used by Sanh et al. (2020).

Comparing biased and main models. We compare the two models in terms of the input tokens on which they base their decisions. The role of input tokens is measured using word omission, similarly to Kádár et al. (2017). We consider the two models as behaving differently if their dominating input tokens differ significantly.

Consider an *easy* instance (\mathbf{x}_i, y_i) which is correctly classified by both the biased f_b and the main f_m models. We denote the tokens of \mathbf{x}_i by $\{x_{ij}\}_{j=1}^{m_i}$, where m_i is the number of tokens in \mathbf{x}_i (BERT tokenizer was used in our experiments). The role of token x_{ij} for the decision made by a classifier f was computed as the change in the true class logit upon excluding that token from the input. More precisely, if $e_f(x_{ij})$ is the estimated effect of the token x_{ij} on f 's decision², $e_f(x_{ij}) = f(\mathbf{x}_i)_{y_i} - f(\mathbf{x}_i \setminus \{x_{ij}\})_{y_i}$, where $f(\mathbf{x})_y$ is the logit of \mathbf{x} for the class y produced by f . We represent the way that the biased model f_b treats an easy instance (\mathbf{x}_i, y_i) as $e_{f_b}(\mathbf{x}_i) = (e_{f_b}(x_{i1}), e_{f_b}(x_{i2}), \dots, e_{f_b}(x_{im_i}))$. The same is done to obtain $e_{f_m}(\mathbf{x}_i)$.³ Finally, cosine similarity between $e_{f_b}(\mathbf{x}_i)$ and $e_{f_m}(\mathbf{x}_i)$ is computed as an estimate for the similarity between the two models on instance i . We will show in Section 3.1 that this estimate is highly in line with human interpretation.

3 Experiments

With two datasets, MNLI and FEVER, and two biased models, partial-input and TinyBERT, there are four experimental settings in total. We refer to

²A better notation would be $e_f(x_{ij}; \mathbf{x}_i, y_i)$, but we omit (\mathbf{x}_i, y_i) for simplicity.

³Note that in the partial-input experiments, only the partial input tokens are considered in the representations.




Figure 1: The distribution of inter-model (cosine) similarities. Instances below the threshold are taken as differently handled by the two models.

the two biased models as *Partial* and *Tiny*.

3.1 Manual validation

Using the methodology explained in the previous section, for each easy instance, we obtain two vectors showing the role of individual input tokens in the final decision made by the main and biased models. We hypothesize that the cosine similarity between these two vectors is correlated with the similarity of the models in treating this instance. To check this assumption, we asked two human annotators to label 250 randomly picked easy instances for each experimental setting (1000 in total). Given the omission-based vectors, the annotators’ task was to judge whether the main and biased models had exploited similar evidences for their decisions or not (binary decision). An equal-width approach with 20 bins was used to sample instances uniformly across the cosine similarity scale (shown in Figure 1). Forty of the instances were shared among the annotators to assess their agreements.

Table 1 shows the results of this manual validation: area under the ROC curve (AUC) as well as Inter-annotator Agreement (IAA). AUC is the probability that a randomly chosen *positive* instance obtains a higher cosine similarity than a randomly chosen *negative* one, where the positive class denotes those instances which are treated similarly by the two models according to the annotators. High AUC values for all the four settings show that the cosine similarity between omission-based vectors of main and biased models is highly in line with

| Dataset | Biased model | IAA | AUC |
|---------|--------------|-------|-------|
| MNL | Partial | 0.900 | 0.974 |
| | Tiny | 0.875 | 0.965 |
| FEVER | Partial | 0.900 | 0.969 |
| | Tiny | 0.875 | 0.925 |

Table 1: Manual validation of the hypothesis that cosine similarity between omission-based vectors obtained from the main and biased models can represent how similarly they treat instances. We report Inter-annotator Agreement (IAA, in terms of accuracy) along with Area Under Curve (AUC).

| Dataset | Bias | Easy | F1 | Different |
|---------|---------|---------------|------|--------------|
| MNL | Partial | 5,381 (54.8) | 93.4 | 1,723 (32.0) |
| | Tiny | 6,090 (62.0) | 94.0 | 1,973 (32.4) |
| FEVER | Partial | 12,627 (63.1) | 94.5 | 3,794 (30.0) |
| | Tiny | 13,328 (66.6) | 90.9 | 6,599 (49.5) |

Table 2: The total number of *easy* instances (% in parentheses) and the subset identified as being treated differently by the two models (*different*). We also report the classifier’s F1 on the gold data.

the human perception of the similarity between behavior of the two models. The high inter-annotator agreement across all the four settings confirms the reliability of this manual validation.

3.2 Quantitative analysis

To have an estimate on the number of easy instances that were handled differently by the two




Figure 2: Examples of *different* instances from MNLI (top) and FEVER (bottom) datasets.

models, we used a simple threshold-based binary classifier tuned based on F1-score on the negative instances in the manually labeled data. The computed thresholds are shown in Figure 1 as red (dashed) lines. Instances with an inter-model similarity below this threshold are regarded as being differently treated by the two models.

The results of the quantitative analysis are presented in Table 2. The *easy* column shows the number (and the percentage in parentheses) of easy instances, i.e., those that are correctly classified by both the main and biased models. The subset of easy instances which are identified by the threshold-based classifier as being differently treated by the two models is shown under column *different*. We also report the F1-score for the corresponding classifier. The high F1-scores indicate that the threshold-based classifier is accurate in predicting the gold labels assigned by the annotators.

The results show that the main model handles a considerable subset of easy instances differently from the biased model. This undermines the soundness of a core assumption made by many bias mitigation techniques. Specifically, for three of the four settings, this comprises around a third of easy instances, the subset which is often unnecessarily discarded by the bias mitigation techniques. The FEVER-Tiny setting is an exception for which the estimate is close to 50% of the easy instances. This can explain the particular brittleness of TinyBERT-based bias detection approach on the FEVER dataset which explains the lower performance of this approach reported by Sanh et al. (2020) compared to partial-input.

3.3 Discussion

We showed that the main model does not follow the biased models for a significant number of instances.

Figure 2 shows two such instances. On top, the biased model seems to have based its decision mostly on the negative word *never* in the hypothesis to make a *contradiction* prediction. Despite having access to this keyword, the main model utilizes a wider range of relevant words such as *worked* in the premise and the negated phrase *never worked* in the hypothesis. The same is true for the second example where the biased model uses the negative word *only* whereas the main model takes additional evidence such as the contradiction between *film* in the premise and *podcast* in the hypothesis. More examples are shown in the appendix.

Figure 3 shows the distribution of *easy* and *different* instances over the three dataset labels. In FEVER, a significant proportion of *different* samples belong to the *support* class, while the *refute* class has the smallest share.⁴ The same holds for the *MNLI-Partial* setting. This shows that in most cases, it is much more likely for the main model to find evidences that are different from those exploited by the biased model when the *premise* entails *hypothesis*. On the other hand, when the *premise* contradicts the *hypothesis*, the main model highly resorts to the same cues detected by the biased model.

This observation suggests that class labels are helpful information which can be exploited by bias mitigation techniques. In other words, instance weights can be determined not only based on the biased model’s loss during training, but also by incorporating the corresponding classes. For example, it is better to reduce the down-weighting for those instances in the FEVER dataset that belong to the *support* class.

⁴In this section, we use the MNLI and FEVER terminology interchangeably.




Figure 3: Label distribution in the validation data (total), easy instances (Easy-), and the subset identified as being handled differently by the two models (Diff-) and for the two bias detection models (-Part and -Tiny).

4 Related Work

The presence of dataset-specific superficial cues have been shown in different NLP tasks including visual question answering (Jabri et al., 2016; Manjunatha et al., 2019), machine reading comprehension (Jia and Liang, 2017), natural language inference (Gururangan et al., 2018; Poliak et al., 2018), abusive language detection (Wiegand et al., 2019), and fact verification (Schuster et al., 2019). Models that rely on these dataset biases usually have low performance in out-of-distribution settings. To measure the reliance on non-generalizable patterns, various challenging sets have been proposed (Jia and Liang, 2017; Agrawal et al., 2018; McCoy et al., 2019; Zhang et al., 2019; Sakaguchi et al., 2020). There have also been some attempts to systematically reduce dataset biases during construction (Zellers et al., 2018; Le Bras et al., 2020). Despite some bias reduction reported, the datasets may still contain hidden biased patterns (Sharma et al., 2018); therefore, it is crucial to empower the learning algorithms to be robust against biases.

A popular approach to mitigate dataset biases is to encourage the main model to pay less attention to the instances that are correctly classified by a biased model. To train a biased model, some methods use a-priori known sources of biases. For instance, good performance given insufficient semantics of the input is attributed to bias exploitation (He et al., 2019; Clark et al., 2019; Cadene et al., 2019; Mahabadi et al., 2020). Others try to identify biased instances without explicitly modeling them, such as by training a limited capacity model (Sanh et al., 2020; Clark et al., 2020) or exposing the model to only a small number of training instances (Utama et al., 2020). To decrease the reliance of the model on the (likely) biased instances, some techniques implicitly reduce the updates of main model’s pa-

rameters for biased instances (Clark et al., 2019; Cadene et al., 2019; Sanh et al., 2020; Utama et al., 2020; Mahabadi et al., 2020). Others explicitly downweight the biased instances, for instance using debiased focal loss (Mahabadi et al., 2020) or example reweighting (Utama et al., 2020).

5 Conclusions

Through a set of experiments, we showed that a common core assumption of dataset bias mitigation methods does not hold for a significant portion of two widely-used benchmarks. Specifically, we observed that two widely-used bias detection approaches, partial-input and low-capacity model, are unable to accurately predict model’s handling of biased instances. We carried out extensive analysis and manual validation to attest the reliability of this observation. We infer that what identifies a biased instance is not the instance itself, but the way the model treats it. In other words, an instance that has evident biased patterns is not necessarily useless as long as the main model does not base its decision on these biases.

It is worth noting that the dissimilarity in the handling of input between the biased and main models does not imply that the latter necessarily adopts an unbiased strategy. In other words, it is possible for the main model to treat the instance differently from the biased model but still exploits a (different) bias. As immediate future work, we plan to build on the findings of this analysis for enhancing dataset bias mitigation techniques.

References

Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Confer-*

- ence on *Computer Vision and Pattern Recognition*, pages 4971–4980.
- Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. 2019. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32:841–852.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4060–4073.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2020. Learning to model and ignore dataset bias with mixed capacity ensembles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3031–3045.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. 2018. Annotation artifacts in natural language inference data. In *2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL HLT 2018*, pages 107–112.
- He He, Sheng Zha, and Haohan Wang. 2019. Unlearn dataset bias in natural language inference by fitting the residual. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 132–142.
- Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. 2016. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031.
- Akos Kádár, Grzegorz Chrupała, and Afra Alishahi. 2017. Representation of linguistic form and function in recurrent neural networks. *Computational Linguistics*, 43(4):761–780.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *International Conference on Machine Learning*, pages 1078–1088.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. End-to-end bias mitigation by modelling biases in corpora. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8706–8716.
- Varun Manjunatha, Nirat Saini, and Larry S Davis. 2019. Explicit bias discovery in visual question answering models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9562–9571.
- R Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *57th Annual Meeting of the Association for Computational Linguistics, ACL 2019*, pages 3428–3448.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis only baselines in natural language inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740.
- Victor Sanh, Thomas Wolf, Yonatan Belinkov, and Alexander M Rush. 2020. Learning from others’ mistakes: Avoiding dataset biases without modeling them. *arXiv preprint arXiv:2012.01300*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. Get your vitamin C! robust fact verification with contrastive evidence. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online.
- Tal Schuster, Darsh Shah, Yun Jie Serene Yeo, Daniel Roberto Filizzola Ortiz, Enrico Santus, and Regina Barzilay. 2019. Towards debiasing fact verification models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3410–3416.
- Rishi Sharma, James Allen, Omid Bakhshandeh, and Nasrin Mostafazadeh. 2018. Tackling the story ending biases in the story cloze test. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 752–757.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and verification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*.

- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020. Towards debiasing NLU models from unknown biases. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7597–7610.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. Detection of abusive language: the problem of biased datasets. 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 602–608.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yadollah Yaghoobzadeh, Soroush Mehri, Remi Tachet des Combes, Timothy J Hazen, and Alessandro Sordani. 2021. Increasing robustness to spurious correlations using forgettable examples. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3319–3332.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. PAWS: Paraphrase adversaries from word scrambling. 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 1298–1308.

A Appendix

A.1 Experimental setup

Table A1 shows the statistics of datasets used in experiments. For the FEVER data, we experimented with the version of Schuster et al. (2019)⁵ augmented with the NEI samples used in Schuster et al. (2021)⁶.

Following Sanh et al. (2020), the models are trained for three epochs with a learning rate of 2e-5 and a batch size of 8. The weight decay rate used for MNLI and FEVER are 0.01 and 0.1, respectively. For FEVER, the learning rate is linearly increased for 1000 warming steps and linearly decreased to 0 afterward. Other hyper-parameters are left as default.

For manual validation, the authors labeled 250 randomly picked easy instances for each experimental setting. They were shown the word omission-based role of different tokens in the main and biased models (as in Figure A1) to decide whether their dominating input tokens differed significantly. As the control check, the Inter-annotator Agreement (IAA) was monitored.

| | | # | Entail | Neutral | Contradict |
|-------|--------|---------|--------|---------|------------|
| MNLI | train | 392,702 | 33.3% | 33.3% | 33.3% |
| | valid. | 9,815 | 35.5% | 31.8% | 32.7% |
| FEVER | train | 242,911 | 41.4% | 41.4% | 17.2% |
| | valid. | 19,997 | 39.9% | 16.7% | 43.4% |

Table A1: The datasets used in the experiments.

A.2 Additional examples

Figure A1 shows some more examples of the instances which are treated differently by the biased and main models.

⁵<https://www.dropbox.com/s/v1a0depfg7jp90f/fever.train.jsonl>
<https://www.dropbox.com/s/bdwf46sa2gcuf6j/fever.dev.jsonl>

⁶https://github.com/TalSchuster/talschuster.github.io/raw/master/static/vitaminc_baselines/fever.zip

