

Overview and Insights from the SCIVER Shared Task on Scientific Claim Verification

David Wadden

University of Washington
dwadden@cs.washington.edu

Kyle Lo

Allen Institute for AI
kylel@allenai.org

Abstract

We present an overview of the SCIVER shared task, presented at the 2nd Scholarly Document Processing (SDP) workshop at NAACL 2021. In this shared task, systems were provided a scientific claim and a corpus of research abstracts, and asked to identify which articles SUPPORT or REFUTE the claim as well as provide evidentiary sentences justifying those labels. 11 teams made a total of 14 submissions to the shared task leaderboard, leading to an improvement of more than +23 F1 on the primary task evaluation metric. In addition to surveying the participating systems, we provide several insights into modeling approaches to support continued progress and future research on the important and challenging task of scientific claim verification.

1 Introduction

Due to both rapid growth in scientific publications and the proliferation of mis- and dis-information online, there is a growing need for automated systems that can assist humans in assessing the veracity of scientific claims with evidence found in the research literature. For the SCIVER shared task, systems are given a claim about a scientific entity and a corpus of abstracts from peer-reviewed research articles, and are asked to identify the articles in the corpus that SUPPORT or REFUTE the claim; each prediction is required to be accompanied with evidentiary sentences, or rationales, from the abstract that justify the labeling decision.

This task poses various challenges to systems. For example, entailment relationships between claims and evidences found in scientific writing are often complex, and understanding them to arrive at a correct SUPPORT or REFUTE labeling decision may require reasoning about numerical quantities, increases and decreases in measurements, or causal chains. Furthermore, since annotations require scientific expertise, training data for this

task are scarce. As a result, systems must employ techniques to overcome the lack of training data, such as domain adaptation or transfer learning.

Here, we report the results on the SCIVER shared task. A total of 11 teams made 14 submissions to the shared task leaderboard, leading to a collective improvement of +23 F1 on the primary evaluation metric, compared to previous baselines. The remainder of this report is organized as follows: In §2, we describe the shared task setup, including the choice of dataset, task definition, and evaluation metrics. In §3, we provide an overview of the systems submitted for the task, and highlight unique features of individual systems. In §4, we present the results from the shared task evaluation. Finally, in §5, we identify several insights into modeling approaches demonstrated by participating systems.

2 Shared task description

We briefly describe the dataset, task, and evaluation. Additional details on the data collection process and evaluation metrics can be found in [Wadden et al. \(2020\)](#).

Dataset We use the SCIFACT dataset from [Wadden et al. \(2020\)](#). SCIFACT consists of 1,409 claims with train, dev, and test splits of 809, 300, and 300 claims respectively. The full train and dev sets and test claims were publicly available six weeks prior to opening the submission portal; gold evidences and labels in the test set are not released publicly. Each claim is an assertion about a single biomedical entity or process, curated by a biomedical expert. These claims are verified by other biomedical experts against a corpus of 5,183 abstracts from peer-reviewed biomedical research articles. For each claim, relevant abstracts are those annotated with evidentiary sentences that SUPPORT or REFUTE the claim.

Task Given an input claim, the task is to (i) identify all abstracts in the corpus that are relevant to

System Name	Team affiliations	Associated paper
VerT5erini	University of Waterloo	Pradeep et al. (2021)
ParagraphJoint	UT Dallas / Chan Zuckerberg Initiative / UCLA	Li et al. (2021)
Law & Econ	ETH Zurich	Stammbach and Ash (2020)
QMUL-SDS	Queen Mary University of London	Zeng and Zubiaga (2021)
first_1	Ping An of China	-
SciKGAT	Tsinghua University / Microsoft Research	Liu et al. (2020a)
bioBert for sciFact	-	-
JC_UKP	TU Darmstadt	-
sum_rationale	UC Irvine	-
base_3_4	-	-
pasic_scibert_tfidf	Ping An International Smart City	-
VERISCI	Allen Institute for AI	Wadden et al. (2020)

Table 1: Systems that submitted to the SCIVER leaderboard. Alongside system names are affiliations of the associated team (if applicable) and any papers associated with the submission (if applicable).

the claim, (ii) label the relation of each relevant abstract to the claim as either `SUPPORTS` or `REFUTES`, and (iii) identify a set of evidentiary sentences (i.e. *rationales*) to justify each label.

Evaluation Systems are evaluated by their F1 score for correctly identifying relevant abstracts. We report two evaluation metrics: *abstract-level* F1 rewards a system for identifying and labeling abstracts correctly. A system may predict up to three evidentiary sentences for each abstract. As long as these three sentences contain a gold rationale, the system is scored as correct; this is similar to the FEVER score from (Thorne et al., 2018). In contrast, *sentence-level* F1 rewards a system for identifying and labeling *individual* evidentiary sentences correctly, similar to the “conditional score” introduced in Thorne et al. (2021). Unlike abstract-level evaluation, this metric penalizes models for over-predicting evidentiary sentences. In practice, we find that systems rank similarly in terms of sentence-level and abstract-level performance.

Model Submissions Submissions were made through a publicly-available online leaderboard¹. To prevent overfitting on the test set, teams were permitted to make one submission per week. The leaderboard was available for seven weeks before final submissions were due.

3 Overview of systems

Table 1 presents the submitted systems. As the online leaderboard is still accepting new submissions, we only include in this report the systems that were present by the shared task submission deadline.

¹<https://leaderboard.allenai.org/scifact>

3.1 Modeling approaches

All systems for which model descriptions are available use a three-stage pipeline consisting of (1) retrieval of relevant abstracts, (2) selection of evidentiary sentences from retrieved abstracts, and (3) label prediction using the identified rationales. Many teams introduced improvements to these three components, which we summarize here.

Abstract retrieval Most systems rely on “bag-of-words” approaches such as TF-IDF or BM25 (Robertson and Zaragoza, 2009) to retrieve an initial set of candidate abstracts. In contrast, PARAGRAPHJOINT computes the distance between “dense” claim and abstract representations using BioSentVec (Chen et al., 2019).

Some systems further refine the initial set of retrieved abstracts using a neural model: VERT5ERINI uses T5 (Raffel et al., 2020) while SCIKGAT uses BERT (Devlin et al., 2019) to re-rank retrieved candidates. QMUL-SDS uses a BioBERT (Lee et al., 2020) text classifier for candidate filtering.

Evidence selection Systems take one of two approaches: *single-sentence* systems predict whether a given sentence was evidence based on the claim and the sentence alone. VERT5ERINI uses T5, QMUL-SDS uses BioBERT and SCIKGAT uses BERT to produce representations for each claim-sentence pair.

The *full-document* systems encode the claim together with the entire abstract, and make predictions for each sentence based on this full-document encoding². The PARAGRAPHJOINT model inter-

²Abstracts longer than 512 tokens are shortened by truncating long sentences

leaves [SEP] tokens between sentences, encodes the entire abstract using RoBERTa (Liu et al., 2019), and finally obtains sentence representations through self-attention pooling over words within each sentence. The Law & Econ model treats rationale selection as a sequence tagging task, using a SciBERT (Beltagy et al., 2019) token-level tagger: any sentence with at least one predicted positive token is taken as evidence.

Label prediction Systems predict SUPPORT, REFUTE, or NEI (Not Enough Info) labels by concatenating the claim and all selected evidentiary sentences and feeding it through a neural three-way text classifier. Unless otherwise noted, systems tend to use the same model class for evidence selection and label prediction stages (e.g., for both pipeline stages, VERT5ERINI use T5 and PARAGRAPHJOINT use RoBERTa).

The SCIKGAT system, which uses BERT for evidence selection, switches to using a kernel graph attention network (Liu et al., 2020b) for aggregating sentences for label prediction. We note the PARAGRAPHJOINT team report experimenting with a similar approach but opt not to use it in their final system due to lack of positive results (Li et al., 2021).

The QMUL-SDS system switches from BioBERT to RoBERTa for the label prediction stage. Furthermore, the model improves classification performance using a two-stage approach – first classifying abstracts as either Containing Evidence or Not, and then classifying evidence-containing abstracts as Supported or Not.

3.2 Model training techniques

We summarize several helpful techniques that teams used to improve model training when developing their systems.

Transfer learning VERT5ERINI, Law & Econ, PARAGRAPHJOINT, and SCIKGAT models were first trained on the FEVER dataset (Thorne et al., 2018) before training on SCIFACT. In contrast, QMUL-SDS was only trained on SCIFACT. SCIKGAT performed additional language model pre-training on data from the CORD-19 dataset (Wang et al., 2020).

Negative sampling Teams devised a number of strategies to expose the model to “negative samples” – abstracts that do *not* contain evidence relevant to a given claim. The VERT5ERINI team used

non-evidence sentences from relevant abstracts as negative samples.

The PARAGRAPHJOINT team used irrelevant abstracts with high lexical similarity to “gold” relevant abstracts as additional negatives. Their system was trained using 12 negative abstracts per positive abstract, allowing it to maintain high precision while retrieving a larger number of candidate abstracts for each claim. This system also used *scheduled sampling*, in which the label predictor is given gold rationales early in training and gradually transitions to using predicted rationales; this was found to increase model robustness.

Along similar lines, the QMUL-SDS team trained their evidence selection and label prediction components with retrieved abstracts rather than just using gold abstracts.

Dev set usage Both VERT5ERINI and PARAGRAPHJOINT teams found it beneficial to perform initial hyperparameter selection on the dev set, and then train a final model on the train and dev sets combined. This is unsurprising given the moderate size of the SCIFACT training set.

4 Results

Table 2 presents performance of all submissions on the public leaderboard during the seven week shared task period. VERT5ERINI performs best on sentence-level evaluation, achieving an improvement of +23.87 F1 (+60.4%) for sentence-level evaluation relative to the VERISCI baseline. PARAGRAPHJOINT performs best on abstract-level evaluation, improving over the VERISCI by +20.66 F1 (+44.4%). Using the dev set as additional training data provides a substantial boost; this strategy alone improves VERT5ERINI performance by roughly +5 F1. While VERT5ERINI achieves higher recall, PARAGRAPHJOINT has higher precision; this is likely because PARAGRAPHJOINT was exposed to more negative samples during training.

5 Insights

A number of participating teams reported performance on individual pipeline components in publications associated with their systems (see Table 1). Based on their reports and discussions with shared task participants, we highlight findings on several modeling decisions that have had significant impact on results in the shared task.³

³Metrics reported in this section are self-reported by participants and have not been verified by the task organizers.

Submission	Sentence-level Selection+Label			Abstract-level Label+Rationale		
	P	R	F1	P	R	F1
VerT5erini (Neural) (Train+Dev)	60.59	66.49	63.40	62.85	71.62	66.95
ParagraphJoint	68.94	54.59	60.94	73.66	61.71	67.16
VerT5erini (Neural)	60.00	57.57	58.76	61.47	63.96	62.69
Law & Econ	56.63	55.41	56.01	62.80	58.56	60.61
VerT5erini (BM25)	58.33	52.97	55.52	60.28	58.11	59.17
QMUL-SDS	66.17	47.57	55.35	72.97	48.65	58.38
first_1	65.06	47.30	54.77	65.36	52.70	58.35
SciKGAT	61.15	42.97	50.48	76.09	47.30	58.33
bioBert for sciFact	54.87	45.68	49.85	58.29	49.10	53.30
JC_UKP	40.69	41.35	41.02	50.24	47.75	48.96
sum_rationale	34.45	47.30	39.86	42.38	51.35	46.44
sum_rationale	28.45	52.97	37.02	35.34	55.41	43.16
base_3_4	5.87	55.95	10.63	9.54	50.45	16.05
pasic_scibert_tfidf	7.95	14.32	10.22	10.11	16.22	12.46
VERISCI	38.56	40.54	39.53	46.61	46.40	46.50
VERISCI Zero-Shot	23.71	31.08	26.90	42.26	31.98	36.41

Table 2: System performance on the SCIVER shared task. Systems are ordered by sentence-level F1. For VERT5ERINI, (Neural) indicates that a neural re-ranker was used for retrieval, and (Train+Dev) indicates that the model was trained on both the train and dev sets as described in §3.2.

Neural refinement of abstract candidates

The original VERISCI model uses TF-IDF to retrieve the top $k = 3$ documents for each claim, but shared task participants have demonstrated substantial improvement over “bag-of-words” retrievers using neural refinement. For instance, the VERT5ERINI neural re-ranker improves Recall@3 from the TF-IDF score of 69.4 to 86.6, a +24.8% increase. The QMUL-SDS system, which uses a binary classifier to filter for relevant abstracts, achieves an F1 of 74.2, while TF-IDF with a Top-3 strategy for identifying relevant abstracts only gets 26.3 F1.

Can we simply replace the bag-of-words component with retrieval based on pre-trained dense representations of? Not yet. The PARAGRAPHJOINT team, which entirely replaced bag-of-words with dense embeddings obtained using BioSentVec (Chen et al., 2019), showed that abstract retrieval performance is slightly worse than simple TF-IDF, with a Recall@3 of 67.0. The QMUL-SDS team also informally reported negative results experimenting with pre-trained DPR (Karpukhin et al., 2020), another dense retrieval technique. These shared task findings comparing bag-of-words, neural refinement, and dense-only retrieval on SCIFACT abstract retrieval have also been demon-

strated in other work (see Table 2 in Thakur et al. (2021)). Adapting pre-trained dense retrieval models to SCIFACT remains an interesting area for future work.

Surrounding context for evidence selection

PARAGRAPHJOINT and Law & Econ demonstrate positive results when incorporating full-abstract context during evidence selection, compared to independent assessment of each sentence. In an ablative analysis using *oracle* abstracts, the PARAGRAPHJOINT team reports that incorporating full-abstract context increases evidence selection performance from 65.3 F1 to between 68.8 and 71.7 F1, relative improvements of +5.4% to +9.8%.

Reliance on larger models

These shared task results demonstrate the unreasonable effectiveness of simply using larger models to improve performance. In particular, VERT5ERINI achieves substantial improvements while preserving much of the model architecture from the original VERISCI baseline, but swapping in the larger T5 model. In an ablative experiment using *gold* evidence, the VERT5ERINI team reports that a label predictor that uses T5-3B achieves a performance of 88.2 F1, compared to 80.2 F1 when using RoBERTa-large, a +10.0% relative improvement. They observe similar improvements on evidence selection as well. In fact, Wadden et al. (2020) even reports higher performance by RoBERTa-

Furthermore, metrics reported in this section are not directly comparable to those in Table 2 as the main results evaluate end-to-end system performance while the metrics reported here are ablations with respect to a pipeline component.

large than SciBERT, where the former has 3 times more parameters, despite the latter being trained on in-domain data.

Still, PARAGRAPHJOINT demonstrates a surprising result in achieving performance competitive to VERT5ERINI using RoBERTa-large (355M parameters, $\approx 10\%$ of T5-3B size), showing that other modeling strategies (e.g. negative sampling) can have significantly benefit system performance while keeping computational burden low. First, this result suggests the need for an improved evaluation setup for the SCIVER task that accounts for model weight classes for submissions and making comparisons. Second, we postulate the need for a follow-up study investigating whether specific modeling strategies employed by smaller models would still translate to significant improvements when using larger models like T5.

6 Conclusion

The SCIVER shared task on scientific claim verification received 14 submissions from 11 teams. The submitted systems achieved substantial improvements on all three stages of the scientific claim verification pipeline – abstract retrieval, evidentiary sentence identification, and label prediction – improving on the previous state-of-the-art by 60% and 44% on the two primary evaluation metrics. We’ve surveyed the various approaches taken by participants and discussed several findings that have emerged from this shared task. Looking forward, the strong performance of systems on SCIVER suggests it may be time to explore other more ambitious challenges on the path toward building real-world scientific claim verification systems. For instance, future work could study retrieval from a much larger scientific corpus, approaches to consolidate evidence from multiple documents while taking into account the strength and credibility of different evidence sources, or techniques for building efficient systems that can approximate performance of the top submissions at a fraction of the cost.

Acknowledgements

We thank all the teams for their participation in the SCIVER shared task. We are especially grateful to Ronak Pradeep and Xueguang Ma (VERT5ERINI), Xiangci Li (PARAGRAPHJOINT), Dominik Stammbach (Law & Econ), Xia Zeng (QMUL-SDS), and Zhenghao Liu and Chenyan Xiong (SciKGAT) for

their helpfulness in providing details about their systems and for participating in workshop presentations and discussions. We also thank the organizing committee of the SDP 2021 workshop for hosting the SCIVER shared task and for their feedback and assistance in helping us organize this event. We finally thank Iz Beltagy, Arman Cohan, Hanna Hajishirzi, and Lucy Lu Wang from AI2 for their help with reviewing, shared task insights and feedback in writing this report.

References

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *EMNLP*.
- Q. Chen, Y. Peng, and Z. Lu. 2019. [BioSentVec: creating sentence embeddings for biomedical texts](#). In *ICHI*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *NAACL*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *EMNLP*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*.
- Xiangci Li, G. Burns, and Nanyun Peng. 2021. [A paragraph-level multi-task learning model for scientific fact-verification](#). In *Proceedings of the Workshop on Scientific Document Understanding*. AAAI.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *arXiv*, 1907.11692.
- Zhenghao Liu, Chenyan Xiong, Zhuyun Dai, Si Sun, Maosong Sun, and Zhiyuan Liu. 2020a. [Adapting open domain fact extraction and verification to COVID-FACT through in-domain language modeling](#). In *Findings of EMNLP*.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020b. [Fine-grained fact verification with kernel graph attention network](#). In *ACL*.
- Ronak Pradeep, Xueguang Ma, Rodrigo Nogueira, and Jimmy Lin. 2021. [Scientific claim verification with VerT5erini](#). In *Proceedings of the 12th International Workshop on Health Text Mining and Information Analysis*. ACL.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Dominik Stammach and Elliott Ash. 2020. [e-FEVER: Explanations and summaries for automated fact checking](#). In *TTO*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogenous benchmark for zero-shot evaluation of information retrieval models](#). *arXiv*, 2104.08663.
- James Thorne, Max Glockner, Gisela Vallejo, Andreas Vlachos, and Iryna Gurevych. 2021. [Evidence-based verification for real world information needs](#). *arXiv*, 2104.00640.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *NAACL*.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *EMNLP*.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Michael Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey A. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex D. Wade, Kuansan Wang, Nancy Xin Ru Wang, Christopher Wilhelm, Boya Xie, Douglas M. Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020. [CORD-19: The COVID-19 open research dataset](#). In *Proceedings of the 1st Workshop on NLP for COVID-19*. ACL.
- Xia Zeng and Arkaitz Zubiaga. 2021. [QMUL-SDS at SCIVER: Step-by-step binary classification for scientific claim verification](#). In *Proceedings of the Second Workshop on Scholarly Document Processing*. ACL.