

# Abstract, Rationale, Stance: A Joint Model for Scientific Claim Verification

Zhiwei Zhang<sup>1,2</sup>, Jiyi Li<sup>2\*</sup>, Fumiyo Fukumoto<sup>2</sup> and Yanming Ye<sup>1</sup>

Hangzhou Dianzi University, Hangzhou, China<sup>1</sup>

University of Yamanashi, Kofu, Japan<sup>2</sup>

hduxiaozhi97@gmail.com, {jyli, fukumoto}@yamanashi.ac.jp  
yeym@hdu.edu.cn

## Abstract

Scientific claim verification can help the researchers to easily find the target scientific papers with the sentence evidence from a large corpus for the given claim. Some existing works propose pipeline models on the three tasks of abstract retrieval, rationale selection and stance prediction. Such works have the problems of error propagation among the modules in the pipeline and lack of sharing valuable information among modules. We thus propose an approach, named as ARSJOINT, that jointly learns the modules for the three tasks with a machine reading comprehension framework by including claim information. In addition, we enhance the information exchanges and constraints among tasks by proposing a regularization term between the sentence attention scores of abstract retrieval and the estimated outputs of rationale selection. The experimental results on the benchmark dataset SCIFACT show that our approach outperforms the existing works.

## 1 Introduction

A system of scientific claim verification can help the researchers to easily find the target scientific papers with the sentence evidence from a large corpus for the given claim. To address this issue, Wadden et al. (2020) introduced scientific claim verification which consists of three tasks. As illustrated in Figure 1, for a given claim, the system finds the abstracts which are related to the claim from a scholarly document corpus (abstract retrieval task); it selects the sentences which are the evidences in the abstract related to the claim (rationale selection task); it also classifies whether the abstract/sentences support or refute the claims (stance prediction task). Wadden et al. (2020) also provided a dataset called SCIFACT.

Most of the existing works of *general* claim verification are based on pipeline models (Soleimani et al., 2020; Alonso-Reina et al., 2019; Liu et al.,

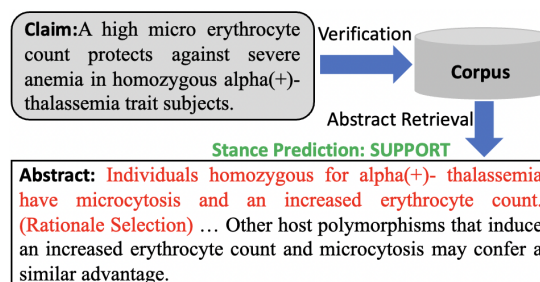


Figure 1: An example of scientific claim verification.

2020; Zhou et al., 2019; Nie et al., 2019; Lee et al., 2020b); some works utilize joint optimization strategies (Lu and Li, 2020; Yin and Roth, 2018; Hidey et al., 2020). These models attempted to jointly optimize the rationale selection and stance prediction, but did not directly link the two modules (Li et al., 2020). In the case of the *scientific* claim verification, Wadden et al. (2020) proposed a baseline model VERISCI based on a pipeline of three components for the three tasks. Pradeep et al. (2021) proposed a pipeline model called VERT5ERINI which utilized the pre-trained language model T5 (Raffel et al., 2020) and adapted a pre-trained sequence-to-sequence model. Li et al. (2020) jointly trained two tasks of rationale selection and stance prediction, and had a pipeline on abstract retrieval task and the joint module.

Above existing works on scientific claim verification are fully or partially pipeline solutions. One problem of these works is the error propagation among the modules in the pipeline. Another problem is that the module in the pipeline trained independently cannot share and leverage the valuable information among each other. Therefore, we propose an approach, named as ARSJOINT, which jointly learns the three modules for the three tasks. It has a Machine Reading Comprehension (MRC) framework which uses the claim content as the query to learn additional information. In addition, we assume that the abstract retrieval module should

have good interpretability and tend to assign high sentence-level attention scores to the evidence sentences that influence the retrieval results; it is consistent with the goal of the rationale selection module. We thus enhance the information exchanges and constraints among tasks by proposing a regularization term based on a symmetric divergence to bridge these two modules.

The experimental results on the benchmark dataset SCIFACT show that the proposed approach has better performance than the existing works. The main contributions of this paper can be summarized as follows. (1). We propose a scientific claim verification approach which jointly trains on the three tasks in a MRC framework. (2). We propose a regularization based on the divergence between the sentence attention of abstract retrieval and the outputs of rationale selection.

## 2 Our Approach

### 2.1 Notation and Definitions

We denote the query claim as  $q$  and an abstract of a scientific paper as  $a \in \mathcal{A}$ . We denote the set of sentences in abstract  $a$  as  $\mathcal{S} = \{s_i\}_{i=1}^l$  and the word sequence of  $s_i$  is  $[s_{i1}, \dots, s_{in_i}]$ . The title of the paper  $t \in \mathcal{T}$  is used as auxiliary information, the word sequence of  $t$  is  $[t_1, \dots, t_{n_t}]$ . Here,  $\mathcal{S}$ ,  $s_i$  and  $t$  are for  $a$  in default and we omit the subscripts ‘ $a$ ’ in the notations. The purpose of the *abstract retrieval task* is to detect the set of related abstracts to  $q$ ; it assigns relevance labels  $y^b \in \{0, 1\}$  to a candidate abstract  $a$ . The *rationale selection* task is to detect the decisive rationale sentences  $S^r \subseteq \mathcal{S}$  of  $a$  relevant to the claim  $q$ ; it assigns evidence labels  $y_i^r \in \{0, 1\}$  to each sentence  $s_i \in \mathcal{S}$ . The *stance prediction* task classifies  $a$  into stance labels  $y^e$  which are in  $\{\text{SUPPORTS}=0, \text{REFUTES}=1, \text{NOINFO}=2\}$ . The sentences in  $a$  have the same stance label value.

### 2.2 Pre-processing

As there are a huge amount of papers in the corpus, applying all of them to the proposed model is time-consuming. Therefore, similar to the existing works on this topic (Wadden et al., 2020; Pradeep et al., 2021; Li et al., 2020), we also utilize a lightweight method to first roughly select a set of candidate papers. We used the BioSentVec (Chen et al., 2019; Pagliardini et al., 2018) to obtain the embeddings of the claim or a scientific paper based on its title and abstract, and compute the cosine sim-

ilarity between the claim and the paper. The papers with top- $k$  similarities are used as the candidates.

### 2.3 Joint Abstract, Rationale, Stance Model

The input sequence of our model is defined as  $seq = [[\text{CLS}]q[\text{SEP}]t \cdots [\text{SEP}]s_i[\text{SEP}] \cdots]$ , which is obtained by concatenating the claim  $q$ , title  $t$  and abstract  $a$ . We compute the list of word representations  $\mathbf{H}_{seq}$  of the input sequence by a pre-trained language model (e.g., BioBERT (Lee et al., 2020a)). We obtain the word representations of the claim  $\mathbf{H}_q = [\mathbf{h}_{q1}, \dots, \mathbf{h}_{qn_q}]$ , the title  $\mathbf{H}_t = [\mathbf{h}_{t1}, \dots, \mathbf{h}_{tn_t}]$ , each sentence  $\mathbf{H}_{s_i} = [\mathbf{h}_{s_{i1}}, \dots, \mathbf{h}_{s_{in_i}}]$ , and the abstract  $\mathbf{H}_S = \mathbf{H}_a = [\cdots, \mathbf{H}_{s_i}, \cdots]$  from  $\mathbf{H}_{seq}$  and use them in our AR-SJOINT model. Figure 2 shows the framework of our model with three modules for the three tasks.

In all three modules, we use attention layer (denoted as  $g(\cdot)$ ) on word (sentence) representations to compute a sentence (document) representation. A document can be a claim, title, abstract, or their combinations. The computation is as follows (refer to (Li et al., 2020)), where the  $*$  in  $\mathbf{H}_*$  represents any type of sentence (claim  $q$ , title  $t$  or a sentence  $s$  in an abstract), the  $\star$  in  $\mathbf{H}_\star$  represents any type of document,  $\mathbf{W}$  and  $\mathbf{b}$  are trainable parameters.

$$g(\mathbf{H}_*) = \sum_i \mathbf{u}_{*i} \alpha_{*i}, \alpha_{*i} = \frac{\exp(\mathbf{W}_{w_2} \mathbf{u}_{*i} + \mathbf{b}_{w_2})}{\sum_j \exp(\mathbf{W}_{w_2} \mathbf{u}_{*j} + \mathbf{b}_{w_2})},$$

$$\mathbf{u}_{*j} = \tanh(\mathbf{W}_{w_1} \mathbf{h}_{*j} + \mathbf{b}_{w_1}) \text{ for word-level attention,}$$

$$g(\mathbf{H}_\star) = \sum_i \mathbf{U}_{\star i} \alpha_{\star i}, \alpha_{\star i} = \frac{\exp(\mathbf{W}_{c_2} \mathbf{U}_{\star i} + \mathbf{b}_{c_2})}{\sum_j \exp(\mathbf{W}_{c_2} \mathbf{U}_{\star j} + \mathbf{b}_{c_2})},$$

$$\mathbf{U}_{\star j} = \tanh(\mathbf{W}_{c_1} \mathbf{H}_{\star j} + \mathbf{b}_{c_1}) \text{ for sentence-level attention.} \quad (1)$$

**Abstract Retrieval:** In this task, a title can be regarded as an auxiliary sentence that may contain the information related to the claim for the abstract, we thus use the title with the sentences in the abstract together. We build a document  $ta = [t, a]$  and concatenate the word representations of  $t$  and  $a$  into  $\mathbf{H}_{ta} = [\mathbf{H}_t, \mathbf{H}_a]$  as the input to this module. We use a hierarchical attention network (HAN) (Yang et al., 2016) to compute document representations  $\mathbf{h}_{ta} \in \mathbb{R}^d$ ,  $\mathbf{h}_{ta} = \text{HAN}(\mathbf{H}_{ta})$ . HAN is proper for document classification by considering the hierarchical document structure (a document has sentences, a sentence has words). We also compute the sentence representation of claim  $\mathbf{h}_q \in \mathbb{R}^d$  with a word-level attention layer (denoted as  $g(\cdot)$ ),  $\mathbf{h}_q = g(\mathbf{H}_q)$ . To compute the relevance between  $\mathbf{h}_{ta}$  and  $\mathbf{h}_q$ , we use a Hadamard product on them and a Multi-Layer Perception (MLP, denoted as  $f(\cdot)$ ) with Softmax (denoted as  $\sigma(\cdot)$ ); the outputs

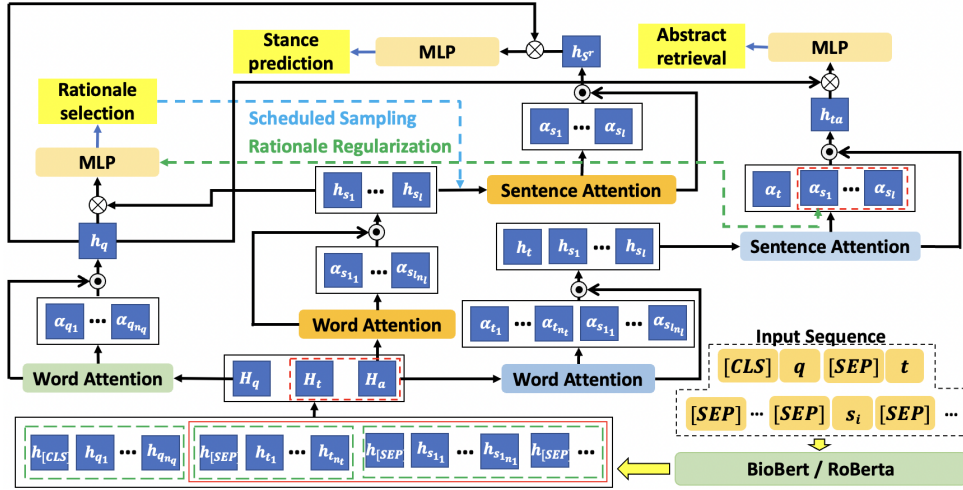


Figure 2: Framework of our ARSJOINT model which jointly learns three modules and has rationale regularization.

are the probabilities that whether the abstract is relevant to the claim,  $[p_0^b, p_1^b] = \sigma(f(\mathbf{h}_q \circ \mathbf{h}_a))$ . A cross entropy loss  $\mathcal{L}_{ret}$  is used for training.

**Rationale Selection:** This task focuses on judging whether a sentence in the abstract is a rationale one or not. For the multiple sentences in the abstract, they have same title information but have different rationale labels. Therefore, when judging each sentence in the abstract, using the title may not positively influence the performance. We thus use the word representation  $\mathbf{H}_a$  of the abstract as input. We compute the sentence representation  $\mathbf{h}_{s_i}$  by a word-level attention layer, and use a MLP with Softmax to estimate the probability  $p_{i1}^r$  and  $p_{i0}^r$  that whether  $s_i$  is the evidence of the abstract or not. The cross entropy loss is  $\mathcal{L}_{rat}$ .

**Stance Prediction:** The module first computes the sentence representation  $\mathbf{h}_{s_i}$  in a same way with that of rationale selection. After that, it only selects the sentences  $S^r$  with the true evidence label  $\hat{y}_i^r = 1$  or the estimated evidence probability  $p_{i1}^r > p_{i0}^r$ ; whether using the true label or the estimated label is decided by a scheduled sampling which will be introduced later. We then compute the estimated stance labels based on a sentence-level attention layer and a MLP with Softmax,  $\mathbf{h}_{S^r} = g(\mathbf{H}_{S^r})$  and  $[p_0^e, p_1^e, p_2^e] = \sigma(f(\mathbf{h}_q \circ \mathbf{h}_{S^r}))$ , where  $S^r = \{s_i \in S | \hat{y}_i^r = 1 \text{ or } p_{i1}^r > p_{i0}^r\}$ . The cross entropy loss is  $\mathcal{L}_{sta}$ .

**Scheduled Sampling:** Since rationale sentences  $S^r$  are used in stance prediction, the error of the rationale selection module will be propagated to the stance prediction module. To alleviate this problem, following (Li et al., 2020), we also use a scheduled sampling method (Bengio et al., 2015), which is to

	SUPPORT	NOINFO	REFUTES	ALL
Train	332 / 370	304 / 220	173 / 194	809
Dev.	124 / 138	112 / 114	64 / 71	300
ALL	456 / 508	416 / 444	237 / 265	1109

Table 1: Statistics of SCIFACT dataset. The numbers are "number of claims / number of relevant abstracts".

feed the sentences with true evidence label  $\hat{y}_i^r = 1$  to the stance prediction module at the beginning, and then gradually increase the proportion of the sentences with the estimated evidence probability  $p_{i1}^r > p_{i0}^r$ , until eventually all sentences in  $S^r$  are based on the estimated evidences. We set the sampling probability of using the estimated evidences as  $p_{sample} = \sin(\frac{\pi}{2} \times \frac{current\_epoch - 1}{total\_epoch - 1})$ .

**Rationale Regularization (RR):** The attention scores have been used for interpretability in NLP tasks (Serrano and Smith, 2019; Wiegrefe and Pinter, 2019; Sun and Lu, 2020). We assume that the abstract retrieval module should have good interpretability and tend to assign high sentence-level attention scores to the evidence sentences that influence the retrieval results; it is consistent with the goal of the rationale selection module. We thus enhance the information exchanges and constraints among tasks by proposing a regularization term based on a symmetric divergence on the sentence attention scores  $\alpha$  of abstract retrieval and the estimated outputs  $\mathbf{y}^r$  of the rationale selection to bridge these two modules. The detailed formula is as follows, where  $\mathbf{p}$  and  $\mathbf{q}$  are  $\alpha$  or  $\mathbf{y}^r$ .

$$\mathcal{D}(\mathbf{p}||\mathbf{q}) = - \sum_{i=1}^l (\mathbf{p}_i \log(\mathbf{q}_i) + (1 - \mathbf{p}_i) \log(1 - \mathbf{q}_i)),$$

$$\mathcal{L}_{RR} = \mathcal{D}(\alpha||\mathbf{y}^r) + \mathcal{D}(\mathbf{y}^r||\alpha). \quad (2)$$

**Joint Training:** We jointly train our model on abstract retrieval, rationale selection and stance prediction. The joint loss with our RR is as follows,  $\mathcal{L} = \lambda_1 \mathcal{L}_{ret} + \lambda_2 \mathcal{L}_{rat} + \lambda_3 \mathcal{L}_{sta} + \gamma \mathcal{L}_{RR}$ , where  $\lambda_1, \lambda_2, \lambda_3$  and  $\gamma$  are hyperparameters.

### 3 Experiments

#### 3.1 Experimental Settings

**Dataset:** We utilize the benchmark dataset SCIFACT<sup>1</sup>. It consists of 5,183 scientific papers with titles and abstracts and 1,109 claims in the training and development sets. Table 1 presents the statistics of the dataset.

**Experimental Settings:** For our ARSJOINT model, we use Optuna (Akiba et al., 2019) to tune the hyperparameters  $\lambda_1, \lambda_2, \lambda_3$  and  $\gamma$  of the loss  $\mathcal{L}$  on 20% of the training set and based on the performance on another 20% training set. We choose the optimal hyperparameters by the average F1-score on abstract-level and sentence-level evaluations. The search ranges of these four hyperparameters are set to [0.1, 12], and the number of search trials is set to 100. Table 2 lists the selected weight hyperparameters of our model. The other hyperparameters such as learning rate in the model refer to the ones used in exiting work (Li et al., 2020) to make a fair comparison. These hyperparameters are listed in Table 3.

We implement our ARSJOINT model<sup>2</sup> in PyTorch. Since the length of the input sequence  $seq$  is often greater than the maximum input length of a BERT-based model, we perform a tail-truncation operation on each sentence of  $seq$  that exceeds the maximum input length. For the pre-trained language model, we verify our approach by respectively using RoBERTa-large (Liu et al., 2019) and BioBERT-large (Lee et al., 2020a) trained on a biomedical corpus. We fine-tune RoBERTa-large and BioBERT-large on the SCIFACT dataset. In addition, the MLP in our model has two layers.

We compare our ARSJOINT approach with Paragraph-Joint (Li et al., 2020), VERISCI<sup>1</sup> (Wadden et al., 2020) and VERT5ERINI (Pradeep et al., 2021). We use the publicly available code<sup>2</sup> of them. The "Paragraph-Joint Pre-training" model is pre-trained on the FEVER dataset (Thorne et al., 2018) and then fine-tune on the SCIFACT dataset. The "Paragraph-Joint SCIFACT-only" is not pre-trained

<sup>1</sup><https://github.com/allenai/scifact>

<sup>2</sup>Our source code is available at: <https://github.com/ZhiweiZhang97/ARSJointModel>

Model	$\lambda_1$	$\lambda_2$	$\lambda_3$	$\gamma$
ARSJOINT w/o RR (RoBERTa)	2.7	11.7	2.2	-
ARSJOINT (RoBERTa)	0.9	11.1	2.6	2.2
ARSJOINT w/o RR (BioBERT)	0.1	10.8	4.7	-
ARSJOINT (BioBERT)	0.2	12.0	1.1	1.9

Table 2: Hyperparameters selected by Optuna for different variants of our model. The "w/o RR" means the model does not utilize rationale regularization.

Name	Value	Name	Value	Name	Value
$k_{tra}$	12	$lr_1$	$1 \times 10^{-5}$	Batch size	1
$k_{ret}$	30	$lr_2$	$5 \times 10^{-6}$	Dropout	0

Table 3: Hyperparameter settings following the existing work.  $k_{tra}$  and  $k_{ret}$  are the number of candidate abstracts for each claim in the training and testing stages.  $lr_1$  and  $lr_2$  are the learning rates of the BERT-based model and other modules of the proposed model.

on other datasets.

**Evaluation:** We evaluate the methods by using the abstract-level and sentence-level evaluation criteria given in SCIFACT<sup>1</sup>. *Abstract-level evaluation:* It evaluates the performance of a model on detecting the abstracts which support or refute the claims. For the "Label-Only" evaluation, given a claim  $q$ , the classification result of an abstract  $a$  is correct if the estimated relevance label  $\hat{y}^b$  is correct and the estimated stance label  $\hat{y}^e$  is correct. For the "Label+Rationale" evaluation, the abstract is correctly rationalized, in addition, if the estimated rationale sentences contain a gold rationale. *Sentence-level evaluation:* It evaluates the performance of a model on detecting rationale sentences. For the "Selection-Only" evaluation, an estimated rationale sentence  $s_i$  of an abstract  $a$  is correctly selected if the estimated rationale label  $\hat{y}_i^r$  is correct and the estimated stance label  $\hat{y}^e$  is not "NOINFO". Especially, if consecutive multiple sentences are gold rationales, then all these sentences should be estimated as rationales. For the "Selection+Label", the estimated rationale sentences are correctly labeled, in addition, if the estimated stance label  $\hat{y}^e$  of this abstract is correct. The evaluation metrics F1-score (F1), Precision (P), and Recall (R) are used. We train the model using all training data, and since Wadden et al. (2020) does not publish the labels on the test set, we evaluate the approaches on the development set following (Li et al., 2020).

#### 3.2 Experimental Results

Table 4 shows the main experimental results. First, the proposed method ARSJOINT (BioBERT) out-

Models	Sentence-level						Abstract-level					
	Selection-Only			Selection+Label			Label-Only			Label+Rationale		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
VERISCI	54.3	43.4	48.3	48.5	38.8	43.1	56.4	48.3	52.1	54.2	46.4	50.0
Paragraph-Joint SCIFACT-only	69.3	50.0	58.1	59.8	43.2	50.2	69.9	52.1	59.7	64.7	48.3	55.3
Paragraph-Joint Pre-training	74.2	57.4	64.7	63.3	48.9	55.2	71.4	59.8	65.1	65.7	55.0	59.9
VERT5ERINI (BM25)	67.7	53.8	60.0	63.9	50.8	56.6	70.9	61.7	66.0	67.0	58.4	62.4
VERT5ERINI (T5)	64.8	57.4	60.9	60.8	<b>53.8</b>	57.1	65.1	<b>65.1</b>	65.1	61.7	<b>61.7</b>	61.7
ARSJOINT w/o RR (RoBERTa)	70.9	56.6	62.9	56.8	45.4	50.5	66.1	56.0	60.6	61.0	51.7	56.0
ARSJOINT (RoBERTa)	67.9	57.1	62.0	55.5	46.7	50.7	64.5	57.4	60.8	59.1	52.6	55.7
ARSJOINT w/o RR (BioBERT)	75.4	57.7	65.3	63.6	48.6	55.1	72.7	57.4	64.2	67.9	53.6	59.9
ARSJOINT (BioBERT)	<b>76.2</b>	<b>58.5</b>	<b>66.2</b>	<b>66.5</b>	51.1	<b>57.8</b>	<b>75.3</b>	59.8	<b>66.7</b>	<b>70.5</b>	56.0	<b>62.4</b>

Table 4: Main experimental results.

Claim: Ly6C hi monocytes have a lower inflammatory capacity than Ly6C lo monocytes.	$\alpha_i$	$\hat{y}_i^r$	$y_i^r$
Blood monocytes are well-characterized precursors for macrophages and dendritic cells.	0.0745	0	0
.....			
Under inflammatory conditions elicited either by acute infection with <i>Listeria monocytogenes</i> or chronic 1,0,0 infection with <i>Leishmania major</i> , there was a significant increase in immature Ly-6C(high) monocytes, resembling the inflammatory left shift of granulocytes.	0.0936	1	1
In addition, acute peritoneal inflammation recruited preferentially Ly-6C(med-high) monocytes.	0.1613	1	1
Taken together, these data identify distinct subpopulations of mouse blood monocytes that differ in maturation stage and capacity to become recruited to inflammatory sites.	0.0745	0	0

Table 5: Result example of Rationale Regularization. Given a claim, it lists the sentences from an abstract.  $\alpha_i$  is sentence attention score in the abstract retrieval task;  $\hat{y}_i^r$  is estimated rationale label;  $y_i^r$  is true rationale label.

performs the existing works with fully or partially pipelines. VERISCI and VERT5ERINI are pipeline models and Paragraph-Joint is a partially pipeline model with a joint model on two tasks. It shows that the proposed model which jointly learns the three tasks is effective to improve the performance.

Second, when using the same pre-trained model RoBERTa-large, comparing our method and the paragraph-joint model, ARSJOINT (RoBERTa) and ARSJOINT w/o RR (RoBERTa) have better performance than "Paragraph-Joint SciFact Only", especially on Recall. It shows that jointly learning with the abstract retrieval task can improve performance. For the Paragraph-Joint method, "Paragraph-Joint Pre-training" with pre-training on another FEVER dataset has much better performance than "Paragraph-Joint SCIFACT-only" without pre-training on other datasets. Similarly, we replace RoBERTa-large with BioBERT-large which contains biological knowledge; ARSJOINT (BioBERT) achieves better performance over "Paragraph-Joint Pre-training".

Third, as an ablation study of the proposed RR, in the case of using BioBERT-large, there is a significant difference between the model with and without RR. Although only a small difference in the case of using RoBERTa-large, there is still an improvement on Recall. This indicates that ratio-

nale regularization can effectively improve the performance of the model. Table 5 shows an example of the results with RR. In this example, it lists a claim and the sentences from an abstract. The attention scores of the sentences in the abstract retrieval task are consistent with the true rationale labels (as well as the estimated rationale labels). The abstract retrieval module thus has good interpretability.

## 4 Conclusion

In this paper, we propose a joint model named as ARSJOINT on three tasks of abstract retrieval, rationale selection and stance prediction for scientific claim verification in a MRC framework by including claim. We also propose a regularization based on the divergence between the sentence attention of the abstract retrieval task and the outputs of the rationale selection task. The experimental results illustrate that our method achieves better results on the benchmark dataset SCIFACT. In future work, we will try to pre-train the model on other general claim verification datasets such as FEVER (Thorne et al., 2018) to improve the performance.

## Acknowledgments

This work was partially supported by KDDI Foundation Research Grant Program.

## References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. OpTuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631.
- Aimée Alonso-Reina, Robiert Sepúlveda-Torres, Estela Saquete, and Manuel Palomar. 2019. Team gplsi. approach for automated fact checking. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 110–114.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS’15*, page 1171–1179, Cambridge, MA, USA. MIT Press.
- Qingyu Chen, Yifan Peng, and Zhiyong Lu. 2019. Biosentvec: creating sentence embeddings for biomedical texts. In *2019 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 1–5. IEEE.
- Christopher Hidey, Tuhin Chakrabarty, Tariq Alhindi, Siddharth Varia, Kriste Krstovski, Mona Diab, and Smaranda Muresan. 2020. DeSePtion: Dual sequence prediction and adversarial examples for improved fact-checking. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8593–8606, Online. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020a. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2020b. Misinformation has high perplexity. *arXiv preprint arXiv:2006.04666*.
- Xiangci Li, Gully Burns, and Nanyun Peng. 2020. A paragraph-level multi-task learning model for scientific fact-verification. *arXiv preprint arXiv:2012.14500*.
- 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 preprint arXiv:1907.11692*.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. Fine-grained fact verification with kernel graph attention network. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7342–7351.
- Yi-Ju Lu and Cheng-Te Li. 2020. GCAN: Graph-aware co-attention networks for explainable fake news detection on social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online. Association for Computational Linguistics.
- Yixin Nie, Haonan Chen, and Mohit Bansal. 2019. Combining fact extraction and verification with neural semantic matching networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6859–6866.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised learning of sentence embeddings using compositional n-gram features. 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 528–540, New Orleans, Louisiana. Association for Computational Linguistics.
- 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*, pages 94–103, online. Association for Computational Linguistics.
- 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.
- Sofia Serrano and Noah A. Smith. 2019. Is attention interpretable? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Amir Soleimani, Christof Monz, and Marcel Worring. 2020. Bert for evidence retrieval and claim verification. In *European Conference on Information Retrieval*, pages 359–366. Springer.
- Xiaobing Sun and Wei Lu. 2020. Understanding attention for text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3418–3428, Online. Association for Computational Linguistics.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.

- 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 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). 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 11–20, Hong Kong, China. Association for Computational Linguistics.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489.
- Wenpeng Yin and Dan Roth. 2018. [TwoWingOS: A two-wing optimization strategy for evidential claim verification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 105–114, Brussels, Belgium. Association for Computational Linguistics.
- Jie Zhou, Xu Han, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2019. Gear: Graph-based evidence aggregating and reasoning for fact verification. In *Proceedings of ACL 2019*.