

# Scientific Discourse Tagging for Evidence Extraction

Xiangci Li<sup>1\*</sup> Gully Burns<sup>2</sup> Nanyun Peng<sup>3</sup>

<sup>1</sup> University of Texas at Dallas, <sup>2</sup> Chan Zuckerberg Initiative

<sup>3</sup> University of California Los Angeles

lixiangci8@gmail.com, gully.burns@chanzuckerburg.com

violetpeng@cs.ucla.edu

## Abstract

Evidence plays a crucial role in any biomedical research narrative, providing justification for some claims and refutation for others. We seek to build models of scientific argument using information extraction methods from full-text papers. We present the capability of automatically extracting text fragments from primary research papers that describe the evidence presented in that paper’s figures, which arguably provides the raw material of any scientific argument made within the paper. We apply richly contextualized deep representation learning pre-trained on biomedical domain corpus to the analysis of scientific discourse structures and the extraction of “evidence fragments” (i.e., the text in the results section describing data presented in a specified subfigure) from a set of biomedical experimental research articles. We first demonstrate our state-of-the-art scientific discourse tagger on two scientific discourse tagging datasets and its transferability to new datasets. We then show the benefit of leveraging scientific discourse tags for downstream tasks such as claim-extraction and evidence fragment detection. Our work demonstrates the potential of using evidence fragments derived from figure spans for improving the quality of scientific claims by cataloging, indexing and reusing evidence fragments as independent documents.

## 1 Introduction

Primary experimental articles (i.e., papers that describe original experimental work) provide the crucial raw material for all other subsequent scientific research. However, the drastically growing number of scientific literature makes it increasingly difficult for domain experts to efficiently utilize them. Automatic information extraction from biomedical

\* Work performed while the author is interning at the Information Sciences Institute, University of Southern California

lc3 , the mammalian atg8 homolog , undergoes a set of modifications resulting in conversion from lc3i to lc3ii during autophagy 42 . [fact] to further test the function of rag in autophagy [goal] we examined the lc3 modification in hek293 cells . [method] expression of raga ql and raga sn inhibited lc3 conversion in response to amino acid starvation ( fig. 7e ) . [result] furthermore , expression of raga tn and raga ql enhanced lc3 conversion even in the presence of amino acids . [result] these results are consistent with the data observed in drosophila and further demonstrate a role of the rag gtpases in autophagy regulation in response to nutrient signals [implication]

Figure 1: An example paragraph tagged with scientific discourse tags on each clause in SciDT dataset (Dasigi et al., 2017). The text is tokenized and converted to lower case.

literature is a crucial step to help researchers to achieve this goal.

Extracting important information from biomedical literature to facilitate and accelerate scientific discovery has been a goal for computational linguistics for some time (Hobbs, 2002), with the focus of identifying relevant entities, relations, and events from text to populate a knowledge base. However, these methods do not take into account the fact that scientific work involves attempting to provide explanations for evidence derived from experiments and is therefore driven principally by authors attempting to convince expert readers that their claims are the “correct” explanations for the experimental evidence. Thus, an important aspect of building machines capable of understanding scientific literature is first recognizing different *rhetorical components* of scientific discourse, with which we will then be able to distinguish the *observations* made in experiments from their *implications* and distinguish between *claims* supported by evidence and *hypotheses* put forward to prompt further research. It is this goal, of being able to distinguish between the different rhetorical components of scientific discourses so that we can build AI systems to facilitate more accurate analysis and understand-

Type	Definition
Goal	Research goal
Fact	A known fact, a statement taken to be true by the author
Result	The outcome of an experiment
Hypothesis	A claim proposed by the author
Method	Experimental method
Problem	An unresolved or contradictory issue
Implication	An interpretation of the results
None	Anything else

Table 1: Eight label taxonomy defined by De Waard and Maat (2012).

ing of scientific literature, that motivates our work.

Scientific discourse tagging is a task that tags clauses or sentences in a scientific article with different rhetorical components of scientific discourses. Figure 1 shows an example of a paragraph with discourse tags. In this work, we leverage a state-of-the-art contextualized word embedding and a novel word-to-sentence attention mechanism to develop a model for scientific discourse tagging that achieves the state-of-the-art performances on two benchmark datasets SciDT (Dasigi et al., 2017) and PubMed-20k-RCT (Jin and Szolovits, 2018) by 6.9% and 2.3% absolute F1 respectively.<sup>1</sup> More importantly, we show the strong transferability of our scientific discourse tagger to new datasets by beating the baseline (Huang et al., 2020) via zero-shot prediction on CODA-19 dataset (Huang et al., 2020). Furthermore, we demonstrate the effectiveness of scientific discourse tagging on two downstream scientific literature understanding tasks: claim-extraction and evidence fragment detection, and demonstrate the benefit of leveraging scientific discourse tags information. In particular, we outperform the state-of-the-art claim extraction model (Achakulvisut et al., 2019) by 3.8% F1, and outperform figure span detection baseline (Burns et al., 2017) by 5% F1.

## 2 Background and Related Works

**Problem Formulation.** We define scientific discourse tagging as a task that labels sentences in a scientific article based on its rhetorical elements of scientific discourse. Formally, a paragraph can be represented as an ordered collection of sequences

<sup>1</sup><https://github.com/jacklxc/ScientificDiscourseTagging>

$\mathbf{S} = [S_1, S_2, \dots, S_n]$ , and each element  $S_i$  is annotated with a discourse label  $L_i \in \{L_1, L_2, \dots, L_k\}$ . Note that  $S_i$  may be defined differently in different datasets – e.g., sentences in the PubMed-RCT dataset (Dernoncourt and Lee, 2017), clauses in the SciDT dataset composed by Burns et al. (2016) and Dasigi et al. (2017), and sentence fragments in CODA-19 dataset (Huang et al., 2020). For conciseness, we refer all these variations as *sentences*. The labels also can be slightly different. For example, in PubMed-RCT,  $\mathbf{L} = \{objective, background, methods, results, conclusions\}$ , in CODA-19 (Huang et al., 2020),  $\mathbf{L} = \{background, purpose, method, finding/contribution, other\}$  while in SciDT dataset (Burns et al., 2016; Dasigi et al., 2017), the labels  $\mathbf{L} = \{goal, fact, hypothesis, problem, method, result, implication, none\}$  as defined by De Waard and Maat (2012). Table 1 gives more details about the definitions of the tags.

### 2.1 Prior Works on Scientific Discourse Tagging

#### Feature-based Scientific Discourse Tagging.

There has been a significant amount of work aimed at understanding types of scientific discourse. Teufel and Moens (1999) and Teufel and Moens (2002) described argumentative zoning, which groups sentences into a few rhetorical zones highlighted by important clauses such as “in this paper we develop a method for”. Hirohata et al. (2008) used conditional random field (CRF) (Lafferty et al., 2001) with handcrafted features to classify sentences in abstracts into 4 categories: *objective, methods, results, and conclusions*. Liakata (2010) defined “zone of conceptualization” which classifies sentences into 11 categories in scientific papers and Liakata et al. (2012) used CRF and LibSVM to identify these “zone of conceptualization”. Guo et al. (2010) used Naive Bayes and Support Vector Machine (SVM) (Cortes and Vapnik, 1995) to compare three schema: section names, argumentative zones and conceptual structure of documents. Burns et al. (2016) studied the problem of scientific discourse tagging, which identifies the discourse type of each clause in a biomedical experiment paragraph and composed a dataset for it. They adopted the discourse type taxonomy for biomedical papers proposed by De Waard and Maat (2012). The taxonomy contains eight types including *goal, fact, result, hypothesis, method, problem, implication* and *none* as Table 1 shows. Most recently,

Cox et al. (2017) used the same schema (De Waard and Maat, 2012) by exploring a variety of methods for balancing classes before applying classification algorithms.

### Deep Learning for Scientific Discourse Tagging.

Due to the prevalence of deep learning, neural sequence labeling approach using bidirectional LSTM (Hochreiter and Schmidhuber, 1997) and CRF (BiLSTM-CRF) (Huang et al., 2015) has been prevailing for classic word-level sequence tagging problems such as named entity recognition (NER), part of speech tagging (POS), and word segmentation (Huang et al., 2015; Peng and Dredze, 2015, 2016; Ma and Hovy, 2016; Chiu and Nichols, 2016; Peng and Dredze, 2017; Wang et al., 2017; Huang et al., 2019). Since scientific discourse tagging, which is a sentence-level sequence tagging problem, has one additional dimension of input comparing to word-level sequence tagging problems, an encoder is required to encode word-level representations to clause/sentence-level representations. While one simple way is to pre-compute sentence embeddings from word embeddings (Arora et al., 2016), there are more sophisticated methods to compute sentence-level embeddings on-the-fly using BiLSTM (Jin and Szolovits, 2018; Srivastava et al., 2019) or attention (Dasigi et al., 2017), before feeding them into a clause/sentence-level sequence tagger. Alternatively, as BERT (Devlin et al., 2018) prevails among various natural language processing (NLP) tasks, a simple baseline method is to directly use a BERT-like model’s (e.g. SciBERT (Beltagy et al., 2019)) prefix token ( $[CLS]$ ) representation of each sentence as the sentence representation for classification task (Huang et al., 2020). In this work, we combine these methods to present a state-of-the-art scientific discourse tagger.

## 2.2 Downstream Applications

### Claim Extraction.

Claim extraction has been extensively studied in various domains. In addition to scientific articles (Stab et al., 2014), previous work has analyzed social media (Dusmanu et al., 2017), news (Habernal et al., 2014; Sardianos et al., 2015) and Wikipedia (Thorne et al., 2018; Fréard et al., 2010) for a task called *Argumentation Mining* to extract claims and premises. However, there are less attention and dataset available in the biomedical domain. Achakulvisut et al. (2019) composed a

Panax notoginseng saponins (PNS) are components derived from Chinese herb panax notoginseng and play important roles in the cure of wounds. However, how PNS plays this function is still unclear. In this study, we used MTT assay, wound healing assay, western blot, quantitative real time PCR and enzyme-linked immunosorbent assay to detect the effects of PNS on the proliferation, migration and expression of collagen and fibronectin of anterior cruciate ligament (ACL) fibroblasts as well as the underlying mechanism. We found that PNS promoted the proliferation and migration of ACL fibroblasts and increased the expression levels of collagen and fibronectin. **Further mechanism study indicates that PNS might play its function through the phosphorylation of PI3K, AKT and ERK. This study provides a possible mechanism for the function of PNS and lays foundation for further study on the function of panax notoginseng.** Figure 2: An example abstract with claim sentences highlighted in claim-extraction dataset (Achakulvisut et al., 2019).

To ensure equal loading of protein in all lanes  
[goal, O] ..... cell supernatants from  
uninfected and ... [method, O] As seen in  
Figure 1C , SHIP- / [result, B] BMMs  
produced significantly elevated levels of ... Fragment  
[result, I] Similar results were obtained with for Fig. 1C  
peritoneal [result, I] that SHIP-deficient  
macrophages make more ... [implication, I]  
Protein-matched cell lysates from SHIP+/+ ...  
[method, B] BMMs were analyzed by ... Fragment  
(Figure 1D) . [goal, I] These data indicate for Fig. 1D  
[implication, I] that SHIP is a negative  
regulator of IL-12 , IL-6 , ... [implication, I]  
The regulatory influence of ... (MOIs ; Figure  
2A-2C ) , [result, B] suggesting that the  
negative regulatory effect of ... [implication, I]  
In addition , the viability of the bacteria is  
not ... [implication, B] as macrophages Fragment for  
infected with ... ( Figure 2D-2F ) . [result, I] Fig. 2D, 2F  
Likewise , SHIP also down-regulated [result, I]  
F. tularensis LVS-induced ... ( Figure 2D-  
2F ) . [result, I]

Figure 3: An example paragraph of evidence fragment detection. The explicit mention of subfigure codes are underlined. The red lines indicate the borders of the evidence fragments. For each clause, the discourse type as well as the BIO tags indicating “blocks” (see Section 3.2.2) are provided.

claim-extraction dataset derived from MEDLINE<sup>2</sup> paper abstracts, and proposed a neural network model that significantly outperformed the rule-based method proposed by Sateli and Witte (2015). Figure 2 shows an example abstract with the last two sentences annotated as claims.

In this work, we formulate claim extraction (Achakulvisut et al., 2019) similarly as scientific discourse tagging:  $S$  contains sentences and  $L_i \in \{0, 1\}$  indicates whether the corresponding sentence is a claim or not.

**Evidence Fragment detection.** Burns et al. (2017) coined the concept of “evidence fragments” as the text section in narrative surrounding a figure reference that directly describes the experimental

<sup>2</sup><https://www.nlm.nih.gov/bsd/medline.html>

figure. They composed an evidence fragment detection dataset, and proposed the evidence fragment detection task that tags each clause with semantically referred subfigure codes. They further proposed a rule-based method of using these subfigure codes as anchors to link evidence fragments to European Bioinformatics Institute’s INTACT (Orchard et al., 2013) data records. As a result, INTACT’s preexisting, manually-curated structured interaction data can serve as a gold standard for machine reading experiments.

Burns et al. (2017) formulated the problem into a clause-level tagging problem. Formally, each clause  $S_i$  in a paragraph  $\mathbf{S} = [S_1, S_2, \dots, S_n]$  is annotated with a set of subfigure codes  $f^i = \{f_1^i, f_2^i, \dots, f_m^i\}$  that each clause is semantically referring to, where the length  $m$  can be any non-negative integer. Figure 3 shows an illustration of a paragraph of evidence fragment detection annotation. Each clause in the paragraph is associated with a set of semantically relevant subfigures.

### 3 Approaches

#### 3.1 Scientific Discourse Tagger

**Model Overview.** We formulate scientific tagging as a sentence level sequence tagging problem. We develop a deep structured model extending Dasigi et al. (2017), which consists of a contextualized word embedding layer, an attention layer that summarizes word embeddings into sentence embeddings, and a BiLSTM-CRF sequence tagger (Huang et al., 2015) on top of the sentence embeddings for discourse type tagging. Figure 4 gives an overview of the architecture. We detail each component in this section.

**Embeddings.** We explore pre-trained BioGloVe embedding (Burns et al., 2019), BioBERT (Lee et al., 2019) and SciBERT (Beltagy et al., 2019) embedding, which are GloVe and BERT embeddings trained on the text in biomedical domain.

**Sentence Representations via Attention.** We observe that only keywords are essential to determine the discourse types, and attention is an appropriate mechanism for emphasizing certain inputs and ignoring others. Dasigi et al. (2017) also explored using an attention mechanism to summarize word representations to sentence representations, however, we propose a new variation of attention mechanism using an LSTM. Specifically, we first encode the sentence using an LSTM to get contextualized hidden vectors of each word  $h_i$ , and use them to

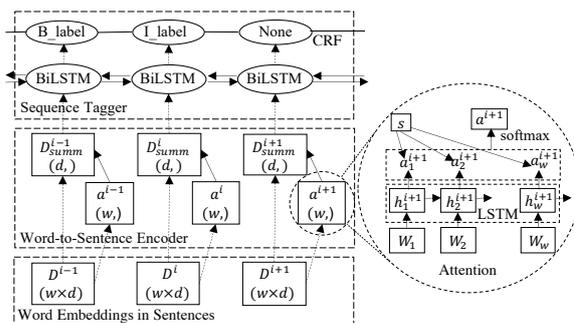


Figure 4: High-level overview of our scientific discourse tagger. Dashed arrows indicate we may apply dropout in those connections.

learn attentions by introducing another trainable vector  $s$  of the same dimension of  $h_i$ . We then apply the attention to summarize the word embeddings into a clause embedding. Detailed equations are provided in section A.1. The dashed circle in Figure 4 illustrates our LSTM-Attention based clause encoder.

**Sentence-level Sequence Tagging.** We observe that the discourse labels have a clear transition of logic flow (e.g. *result* usually followed by *implication*, and *method* usually followed by *hypothesis*). Therefore, we extend LSTM sequence tagger used by Dasigi et al. (2017) to BiLSTM-CRF sequence tagger (Huang et al., 2015) to label discourse types for each sentence in a paragraph.

**Labels in BIO Scheme.** We use the BIO scheme (Sang and Veenstra, 1999) to train all of our models (Baseline models for SciDT dataset do not use BIO2 scheme). Specifically, we convert the labels into BIO scheme where *none* label represents *O* and all other labels are converted into *B\_label* when the previous label type is different from the current label and *I\_label* when the previous label is the same as the current label.

#### 3.2 Downstream Applications

##### 3.2.1 Claim Extractor

Due to the similar problem formulation of evidence extraction task (Achakulvisut et al., 2019), we directly employ the discourse tagging model for claim extraction.

##### 3.2.2 Evidence Fragment Detector

**Problem Reduction.** As Figure 3 shows, since each clause in evidence fragment detection task may refer to more than one subfigure codes, we cannot directly solve it as a standard classification task. Instead, we reduce it to a clause-level sequence

tagging problem under a block-based assumption. We treat each paragraph as a single input. During training, we encode the subfigure code reference sequences of the clauses in each paragraph into a single BIO (Sang and Veenstra, 1999) sequence (where  $B$  indicates the clause is the beginning of a block,  $I$  indicates the clause is in the same block as the previous clause, and  $O$  indicates that no subfigure code is being referred to) as demonstrated at the end of each clause in Figure 3. For prediction, we decode the semantic subfigure code references of all clauses from the BIO sequence for each paragraph following the same block-based assumption.

**Block-Based Assumption.** Most subfigure code reference labels are block-based. We call contiguous clauses that share the same subfigure code reference labels as a block, which is segmented by red lines in Figure 3. We further observe that most blocks explicitly mention all of the semantically referred subfigure codes at least once. Therefore, assuming this property is true for all blocks, we can reconstruct a sequence of semantic subfigure code references for all clauses in a paragraph. We use the explicitly mentioned subfigure codes for each block and a BIO sequence indicating where each block starts and ends for the reconstruction. Consequently, during encoding, we convert annotated semantically referenced subfigure code labels into BIO scheme. During decoding, we first localize the start and end position of each block using BIO predicted tags, then fill each block with all explicitly mentioned subfigure codes.

**Clause-level Sequence Tagger.** The key part of our sequence tagging-based solution for evidence fragment detection is to determine where a block starts and ends. We apply a clause-level sequence tagger to tag each clause in a paragraph. Due to the small size of the evidence fragment detection dataset, we empirically observe that feature-based CRF sequence taggers outperform neural-network based sequence taggers, we thus adopt the feature-based model. In addition to the scientific discourse tags, we use explicitly mentioned subfigure codes as well as unigram, bigram and trigram words as features. For each clause, we use all features described previously from the current clause in addition to the same sets of features from the adjacent previous and next clauses.

SciDT		CODA-19		PubMed-RCT	
result	2053	finding / contribution	50018	methods	79214
method	1542	background	46082	results	77507
implication	797	method	16216	conclusions	36321
fact	732	purpose	8252	background	28797
hypothesis	514	other	2486	objective	18548
goal	225				
problem	193				
none	68				

Figure 5: Count of each label in three datasets. The lines correspond to the mappings from SciDT dataset (Burns et al., 2016; Dasigi et al., 2017) and PubMed 20k RCT dataset (Dernoncourt and Lee, 2017) to CODA-19 dataset (Huang et al., 2020) for zero-shot predictions (Section 5.2).

## 4 Experimental Setup

We evaluate the performance of our scientific discourse tagger on PubMed-RCT dataset (Dernoncourt and Lee, 2017) and SciDT dataset (Burns et al., 2016; Dasigi et al., 2017) (Section 5.1). We also examine the transferability of our scientific discourse tagger to new datasets using CODA-19 dataset (Huang et al., 2020) (Section 5.2). We further study the efficiency of scientific discourse tags on claim-extraction task via transfer learning as well as evidence fragment detection task in a pipeline fashion (Section 6).

### 4.1 Datasets

Figure 5 shows the distribution of the labels in the three datasets introduced below as well as their mappings used for zero-shot predictions in Section 5.2.

**PubMed-RCT Dataset.** We use PubMed-RCT (Dernoncourt and Lee, 2017) as the standard dataset to evaluate our scientific discourse tagger against other strong baselines. PubMed-RCT is derived from PubMed for sequential sentence classification. It has two versions – a smaller PubMed 20k RCT, and a 10 times larger PubMed 200k RCT. Due to our limited availability of computational resources, we only consider PubMed 20k RCT in this work. PubMed 20k RCT is a large dataset that consists of 20k abstracts of randomized controlled trials (RCTs), with vocabulary of 68k across 240k sentences. Each sentence of an abstract is labeled with one of the following roles (section heads) in the abstract: *background*, *objective*, *method*, *result* or *conclusion*.

**SciDT Dataset.** Similar to PubMed-RCT (Dernoncourt and Lee, 2017), SciDT dataset (Burns et al., 2016; Dasigi et al., 2017) is a clause-based

dataset with more fine-grained taxonomy. We further expand SciDT dataset by applying the same clause parsing and annotation pipeline described by Dasigi et al. (2017). This dataset is derived from the Pathway Logic (Eker et al., 2002) and INTACT databases (Orchard et al., 2013). Texts from all sections of each of those papers were pre-processed by parsing each sentence to generate a sequence of main and subordinate clauses using Stanford Parser (Socher et al., 2013). Domain experts were asked to label each of the clauses using the 7-label taxonomy proposed by De Waard and Maat (2012) whose distributions are shown in Figure 5. We apply sequential methods to sequences of clauses in individual paragraphs.

Overall, SciDT dataset has a total of 634 paragraphs and 6124 clauses. We randomly split 570 paragraphs as the training and validation set and the rest as the test set. Each paragraph contains up to 30 clauses and the number of word per clause has a mean of 17.7 and a standard deviation of 12.5. The total vocabulary size is 8563, which is a small dataset for an NLP task. However, we note the difficulties of obtaining such dataset. We further perform a quality assessment of the dataset by re-annotating the test set. We obtain Cohen’s kappa coefficient  $\kappa = 0.823$ , which indicates a high quality of the dataset.

**CODA-19 Dataset.** CODA-19 (Huang et al., 2020) is a human-annotated dataset on a subset of the abstracts of CORD-19 (Wang et al., 2020), which is a corpus of scholarly articles about COVID-19. Wang et al. (2020) segmented each abstract into sentence fragments by comma (,), semicolon (;), and period (.). Each sentence fragment is labeled with one of the research aspects: *background*, *purpose*, *method*, *finding/contribution* or *other*, which is similar to the label sets of PubMed-RCT (Dernoncourt and Lee, 2017). There are 10966 abstracts in total. We use this dataset to further examine our scientific discourse tagger architecture’s applicability to new datasets as well as the transferability of our trained scientific discourse tagger to new datasets.

## 4.2 Baseline Models

**PubMed-RCT Dataset.** We compare our discourse tagger against two strong baselines on the PubMed 20k RCT dataset: (1) a hierarchical sequential labeling network (HSLN) proposed by Jin and Szolovits (2018) and (2) the state-of-the-

Model		RCT	SciDT
CRF			0.679
SVM			0.737
Dasigi et al. (2017)			0.791
HSLN-CNN		0.922	
HSLN-RNN		0.926	
Srivastava et al. (2019)		0.928	
Embedding	Attention		
BioGloVe	No Context	0.901	0.745
BioGloVe	RNN	0.909	0.763
BioGloVe	LSTM	0.913	0.794
BioBERT	No Context	0.909	0.794
BioBERT	RNN	0.915	0.775
BioBERT	LSTM	0.927	0.794
SciBERT	No Context	0.918	0.806
SciBERT	RNN	0.922	0.817
<b>SciBERT</b>	<b>LSTM</b>	<b>0.951</b>	<b>0.841</b>

Table 2: Scientific discourse tagging performance measured by test F1 score on PubMed 20k RCT and SciDT dataset.

art model (Srivastava et al., 2019) on this dataset. HSLN (Jin and Szolovits, 2018) used bio-word2vec (Moen and Ananiadou, 2013), a word2vec embedding (Mikolov et al., 2013) trained on corpora of Wikipedia, PubMed, and PMC, a convolutional neural network (CNN) (LeCun et al., 2015) (HSLN-CNN) or a BiLSTM (HSLN-RNN) as a sentence encoder, followed by a BiLSTM-CRF architecture (Huang et al., 2015) as a sentence-level sequence tagger. Srivastava et al. (2019) used a similar architecture: bio-word2vec (Moen and Ananiadou, 2013) as word embedding, BiLSTM layer with a special dilation mechanism and a capsule layer (Hinton et al., 2011) as the sentence encoder and BiLSTM-CRF (Huang et al., 2015) as the sentence-level sequence tagger.

**SciDT Dataset.** In addition to the model of Dasigi et al. (2017) trained on our expanded SciDT dataset, we also compare with feature based CRF and SVM with unigram, bigram and trigram words in the previous, current and next clauses as features.

**CODA-19 Dataset.** Huang et al. (2020) composed the CODA-19 dataset and studied a few baselines for scientific discourse tagging. Their best model is a fine-tuned SciBERT (Beltagy et al., 2019).

## 5 Experimental Results

### 5.1 Supervised Learning Results

Table 2 reports the test F1 score of our scientific discourse tagger and its variations against baseline models on PubMed 20k RCT dataset and SciDT dataset. Our best scientific discourse tagger outperforms the state-of-the-art model (Srivastava et al.,

2019) on PubMed 20k RCT dataset by more than 2 % absolute F1 score. Given the large size of PubMed 20k RCT, this result robustly demonstrates the strength of our model. Our model also significantly outperforms Dasigi et al. (2017) with 5% absolute F1 score (per McNemar’s test,  $p < 0.01$ ). Based on these performance, we claim our scientific discourse tagger as state-of-the-art. Note that for scientific discourse tagging, the micro F1 performance is equivalent to accuracy.

**Ablation Studies.** We also perform ablation studies to compare the effect of different word embeddings and attention mechanisms to the performance of our scientific discourse tagger on PubMed-RCT and SciDT dataset in Table 2. All neural network based models discussed for scientific discourse tagging tasks, including ours consist of a word embedding, a word-to-sentence encoder and a sentence-level sequence tagger. As we introduce in Section 3.1, our best model has SciBERT (Beltagy et al., 2019) as our contextualized word embedding, an LSTM-attention structure as our word-to-sentence encoder and BiLSTM-CRF (Huang et al., 2015) as our sentence sequence tagger. Comparing to other baseline models, we improve the model design by adopting the state-of-the-art BERT (Devlin et al., 2018) based language model as our contextualized embedding. Instead of bidirectional LSTM as word-to-sentence encoder used by Jin and Szolovits (2018) and Srivastava et al. (2019), we improve the attention structure proposed by Dasigi et al. (2017). We compare the effect of different embeddings and attention types used in scientific discourse tagger. As Table 2 indicates, our main improvement comes from SciBERT (Beltagy et al., 2019). In addition to BioBERT (Lee et al., 2019) which trains BERT (Devlin et al., 2018) on biomedical domain corpus, SciBERT uses a domain specific vocabulary. BERT as a contextualized embedding also contributes partially to the performance improvement as the BioBERT embedding globally outperforms BioGloVe (Burns et al., 2019), which is a static embedding trained on biomedical domain corpus, on PubMed-RCT dataset. Another source of improvement comes from the attention structure. Our LSTM-attention outperforms the RNN-attention that Dasigi et al. (2017) used.

**Error Analysis.** Figure 6 compares the confusion matrices of Dasigi et al. (2017) and our best scientific discourse tagger on SciDT test set. As

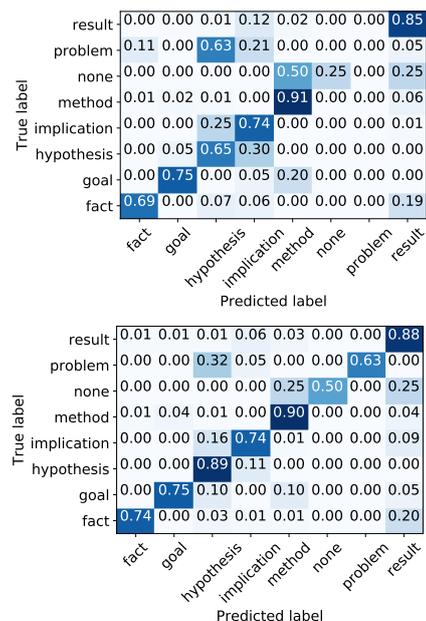


Figure 6: Confusion matrix on SciDT test data. Up: Dasigi et al. (2017). Down: Our scientific discourse tagger.

suggested by the overall performance, our model globally predicts the discourse tags more precisely than Dasigi et al. (2017). Specifically, Dasigi et al. (2017) failed to predict *problem* tag, but our model achieved 0.63 accuracy on predicting *problem* tag. Figure 6 also indicates the different difficulties of predicting different discourse labels due to the imbalance of the label distributions, as Table 5 shows.

## 5.2 Transfer Learning on CODA-19 Dataset

We further demonstrate the strong performance of our scientific discourse tagger by training it on CODA-19 dataset (Huang et al., 2020). As Table 3 shows, our model outperforms the baseline from Huang et al. (2020) by 14.6% absolute F1 on the test set.

More importantly, we use these results as baselines and CODA-19 dataset as an example dataset to show the transferability of our model to new datasets. We first perform zero-shot prediction using our best *trained* scientific discourse taggers on PubMed-RCT (Dernoncourt and Lee, 2017) or SciDT dataset (Dasigi et al., 2017). We map the labels from the original datasets to the target CODA-19 dataset by applying majority vote to the predicted labels on the training set as the lines in Figure 5 show. Then we perform predictions using the best trained scientific discourse taggers on CODA-19 test set and convert the predicted labels from the original label sets to the target CODA-

Model	Test F1
Huang et al. (2020)	0.749
<b>Ours</b>	<b>0.885</b>
Zero-shot Prediction from RCT	0.760
<b>Zero-shot Prediction from SciDT</b>	<b>0.761</b>
<b>PubMed-RCT pre-train</b>	<b>0.909</b>

Table 3: Transfer Learning Performance on CODA-19 Dataset.

Model	Test F1
Achakulvisut et al. (2019)	0.790
Ours (No pre-train)	0.791
<b>Ours (PubMed-RCT pre-train)</b>	<b>0.828</b>

Table 4: Claim extraction performance measured by binary F1 score, which regards 0 as negative label.

19 label set. As a result, as Table 3 shows, our zero-shot prediction results are even higher than the baseline from Huang et al. (2020) which was directly trained on the CODA-19 dataset. This result indicates the strong transferability of our trained scientific discourse tagger as a useful tool on new datasets.

Furthermore, we separately perform a standard transfer learning by taking the scientific discourse tagger pre-trained on PubMed-RCT dataset and fine-tuning it on CODA-19 dataset. We replace the last CRF layer with a new one to match the labels of CODA-19 dataset. As a result, we achieved 0.909 test F1, which is another 2.4% absolute F1 improvement on our model directly trained on CODA-19 dataset. This is likely due to the similar label structures between PubMed-RCT and CODA-19 dataset.

## 6 Downstream Applications

### 6.1 Claim Extraction

**Dataset.** Achakulvisut et al. (2019) introduced an expertly annotated dataset for extracting claim sentences from biomedical paper abstracts. They followed the definitions by Sateli and Witte (2015) to annotate a claim as a statement that either declares something is better, proposes something new, or describes a new finding or a causal relationship. Each sentence is tagged with a binary label indicating it is a claim or not. Each abstract may contain multiple claims as Figure 2 shows. The dataset contains 1500 abstracts sampled from MEDLINE database.

**Baseline Model.** Achakulvisut et al. (2019) constructed claim-extraction dataset and proposed a

model using the sentence classification technique presented by Arora et al. (2016) as sentence encoding method, and the standard BiLSTM-CRF (Huang et al., 2015) as the sentence-level sequence tagger. Their best model was pre-trained on PubMed 200k RCT (Dernoncourt and Lee, 2017) for transfer learning and used GloVe (Pennington et al., 2014) as their word embedding.

**Model Performance.** Table 4 compares the test binary F1 performance of Achakulvisut et al. (2019) with our test performance. We first train our scientific discourse tagger model directly on the claim-extraction dataset. We obtain test binary F1 score of 0.791, which is already higher than Achakulvisut et al. (2019). Then as Achakulvisut et al. (2019) suggested, we pre-train the scientific discourse tagger on PubMed 20k RCT (Dernoncourt and Lee, 2017) and fine-tune it on the claim-extraction dataset. We replace the last CRF layer with a new one to match the binary label structure of claim-extraction dataset. As a result, we obtain test binary F1 score of 0.828, which is another 3.7% absolute F1 improvement on our model without transfer learning. This result demonstrates the benefit of transfer learning from scientific discourse tagging task to its downstream-tasks.

### 6.2 Evidence Fragment Detection

**Dataset.** Burns et al. (2017) introduced evidence fragment detection dataset, which shares the same format and source of clause-based paragraphs with SciDT dataset (Dasigi et al., 2017). As Figure 3 shows, each clause was annotated with subfigure codes that it is semantically referring to. Each clause may not refer to any subfigure code, or simultaneously refer to multiple subfigure codes. The explicit mentions of the subfigure codes were also annotated. All paragraphs are from *Results* section of experimental papers, and most of the paragraphs are from a subset of SciDT dataset (Burns et al., 2016; Dasigi et al., 2017). We further expand evidence fragment detection training set by annotating extra *Results* section paragraphs from SciDT dataset. Overall this small dataset consists of 191 paragraphs as training data and 19 paragraphs as test data.

**Baseline Model.** Burns et al. (2017) proposed a rule-based method for evidence fragment detection task. The key steps are determining where each evidence fragment begins and ends based on the discourse tags of each clause. They treat *hypothe-*

Model	BIO F1	Test F1
Burns et al. (2017)	N/A	0.75
Ours (W/O Discourse Tags)	0.750	0.742
<b>Ours (W/ Discourse Tags)</b>	<b>0.821</b>	<b>0.807</b>

Table 5: Evidence fragment detection performance measured by micro F1 score. Our block-based decoding method achieves 0.94 F1 using ground truth BIO sequences.

*sis*, *problem* and *fact* as indicators of beginning of a evidence fragment, and *result* and *implication* as indicators of the end of a evidence fragment. They also used other features including section headings and whether the references to subfigures are entirely disjoint. Note that their document-level rule-based tagging is across multiple paragraphs in the *Results* section.

**Model Performances.** We use a feature-based CRF with block-based encoding-decoding method to solve this task as a sequence tagging problem. The decoding method described in Section 3.2.2 achieves 0.94 F1 score given the ground truth BIO sequences. Table 5 compares our feature-based CRF model performance with Burns et al. (2017). We also compare our feature-based CRF model performances trained with or without scientific discourse tags from SciDT dataset. Our feature-based CRF model without scientific discourse tags as inputs does not outperform Burns et al. (2017). However, by adding the scientific discourse tag as a feature, we obtain 5.7% absolute F1 improvement over Burns et al. (2017), reaching 0.807 test F1. This improvement is because of the improvement of the CRF sequence tagger. This result shows the strong benefit of scientific discourse tags as the upstream task of evidence fragment detection.

## 7 Discussion

We use the claim-extraction task and the evidence fragment detection task as two examples to demonstrate the benefit of leveraging pre-trained scientific discourse taggers and scientific discourse tags to improve the downstream-task performance via transfer learning or in a pipeline fashion. As Burns et al. (2017) proposed, given the output of evidence fragment detection system, we can link subfigure codes with INTACT (Orchard et al., 2013) records to obtain evidence fragments for each experimental figure.

We further suggest that the evidence fragment detection task can help biocurators delineate evidence

fragments as independent documents so they can be cataloged, indexed, and reused. Traditionally scientists’ arguments are based on relationships between claims and evidences within the same paper and possibly a limited number of cited papers. With the help of evidence fragments, we are able to discard the convention of only linking claims to evidence from a single paper or of following citations, which are often based on linking separate claims from different papers. As a future work, we can surface the evidence fragments combined with figures and captions across multiple papers. Clark et al. (2014) proposed the “Micropublications” semantic model, which is an abstract framework that integrates scientific argument and evidence from scientific documents. Our scientific discourse tagger, claim extractor and evidence fragment detector may serve as the actual implementation of the modules in such a framework. Ultimately, we hope to dramatically increase the amount of primary evidence used to generate individual claims and therefore improve the quality of those claims.

## 8 Conclusions

We develop a state-of-the-art model for scientific discourse tagging and demonstrate its strong performance on PubMed-RCT dataset (Dernoncourt and Lee, 2017) and SciDT dataset (Burns et al., 2016; Dasigi et al., 2017) as well as its strong transferability on new datasets such as CODA-19 dataset (Huang et al., 2020). We then demonstrate the benefit of leveraging the scientific discourse tags on downstream-tasks by providing claim-extraction task and evidence fragment detection task as two show cases. We further propose a future direction that scientific discourse tagging helps delineate evidence fragments as independent documents so they can be cataloged, indexed, and reused. As a result, we can dramatically increase the amount of primary evidence used to generate individual claims and therefore improve the quality of those claims.

## Acknowledgement

We thank the anonymous reviewers for their useful comments, and the PlusLab members for their initial feedback. This work is supported by a National Institutes of Health (NIH) R01 grant (LM012592). The views and conclusions of this paper are those of the authors and do not reflect the official policy or position of NIH.

## References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. **TensorFlow: Large-scale machine learning on heterogeneous systems**. Software available from tensorflow.org.
- Titipat Achakulvisut, Chandra Bhagavatula, Daniel E. Acuña, and Konrad P. Körding. 2019. **Claim extraction in biomedical publications using deep discourse model and transfer learning**. *CoRR*, abs/1907.00962.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2016. A simple but tough-to-beat baseline for sentence embeddings.
- Iz Beltagy, Arman Cohan, and Kyle Lo. 2019. Scibert: Pretrained contextualized embeddings for scientific text. *arXiv preprint arXiv:1903.10676*.
- Gully A Burns, Pradeep Dasigi, and Eduard H Hovy. 2017. Extracting evidence fragments for distant supervision of molecular interactions. *BioRxiv*, page 192856.
- Gully A Burns, Xiangci Li, and Nanyun Peng. 2019. Building deep learning models for evidence classification from the open access biomedical literature. *Database*, 2019.
- Gully APC Burns, Pradeep Dasigi, Anita de Waard, and Eduard H Hovy. 2016. Automated detection of discourse segment and experimental types from the text of cancer pathway results sections. *Database*, 2016.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional lstm-cnns. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- François Chollet et al. 2015. Keras. <https://github.com/fchollet/keras>.
- Tim Clark, Paolo N Ciccarese, and Carole A Goble. 2014. Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications. *Journal of biomedical semantics*, 5(1):1–33.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning*, 20(3):273–297.
- Jessica Cox, Corey A Harper, and Anita de Waard. 2017. Optimized machine learning methods predict discourse segment type in biological research articles. In *Semantics, Analytics, Visualization*, pages 95–109. Springer.
- Pradeep Dasigi, Gully APC Burns, Eduard Hovy, and Anita de Waard. 2017. Experiment segmentation in scientific discourse as clause-level structured prediction using recurrent neural networks. *arXiv preprint arXiv:1702.05398*.
- Anita De Waard and Henk Pander Maat. 2012. Epistemic modality and knowledge attribution in scientific discourse: A taxonomy of types and overview of features. In *Proceedings of the Workshop on Detecting Structure in Scholarly Discourse*, pages 47–55. Association for Computational Linguistics.
- Franck Deroncourt and Ji Young Lee. 2017. **Pubmed 200k RCT: a dataset for sequential sentence classification in medical abstracts**. *CoRR*, abs/1710.06071.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Mihai Dusmanu, Elena Cabrio, and Serena Villata. 2017. Argument mining on twitter: Arguments, facts and sources. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2317–2322.
- S. Eker, M. Knapp, K. Laderoute, P. Lincoln, and C. Talcott. 2002. **Pathway Logic: Executable Models of Biological Networks**. In *Fourth International Workshop on Rewriting Logic and Its Applications (WRLA'2002*, volume 71 of *Electronic Notes in Theoretical Computer Science*. Elsevier.
- Dominique Fréard, Alexandre Denis, Françoise Détienne, Michael Baker, Matthieu Quignard, and Flore Barcellini. 2010. The role of argumentation in online epistemic communities: the anatomy of a conflict in wikipedia. In *ECCE 2010. European Conference on Cognitive Ergonomics*, pages pp–91.
- Yufan Guo, Anna Korhonen, Maria Liakata, Ilona Silins Karolinska, Lin Sun, and Ulla Steinius. 2010. Identifying the information structure of scientific abstracts: an investigation of three different schemes. In *Proceedings of the 2010 Workshop on Biomedical Natural Language Processing*, pages 99–107. Association for Computational Linguistics.
- Ivan Habernal, Judith Eckle-Kohler, and Iryna Gurevych. 2014. Argumentation mining on the web from information seeking perspective. In *ArgNLP*.
- Geoffrey E Hinton, Alex Krizhevsky, and Sida D Wang. 2011. Transforming auto-encoders. In *International Conference on Artificial Neural Networks*, pages 44–51. Springer.

- Kenji Hirohata, Naoaki Okazaki, Sophia Ananiadou, and Mitsuru Ishizuka. 2008. Identifying sections in scientific abstracts using conditional random fields. In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-I*.
- Jerry R. Hobbs. 2002. Information extraction from biomedical text. *Journal of biomedical informatics*, 35(4):260–264.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ting-Hao'Kenneth' Huang, Chieh-Yang Huang, Chien-Kuang Cornelia Ding, Yen-Chia Hsu, and C Lee Giles. 2020. Coda-19: Reliably annotating research aspects on 10,000+ covid-19 abstracts using non-expert crowd. *arXiv preprint arXiv:2005.02367*.
- Xiao Huang, Li Dong, Elizabeth Boschee, and Nanyun Peng. 2019. Learning a unified named entity tagger from multiple partially annotated corpora for efficient adaptation. In *The 2019 SIGNLL Conference on Computational Natural Language Learning (CoNLL)*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Di Jin and Peter Szolovits. 2018. [Hierarchical neural networks for sequential sentence classification in medical scientific abstracts](#). *CoRR*, abs/1808.06161.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature*, 521(7553):436–444.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*.
- Maria Liakata. 2010. Zones of conceptualisation in scientific papers: a window to negative and speculative statements. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 1–4. Association for Computational Linguistics.
- Maria Liakata, Shyamasree Saha, Simon Dobnik, Colin Batchelor, and Dietrich Rebholz-Schuhmann. 2012. Automatic recognition of conceptualization zones in scientific articles and two life science applications. *Bioinformatics*, 28(7):991–1000.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- SPFGH Moen and Tapio Salakoski2 Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. *Proceedings of LBM*, pages 39–44.
- Sandra Orchard, Mais Ammari, Bruno Aranda, Lionel Breuza, Leonardo Briganti, Fiona Broackes-Carter, Nancy H Campbell, Gayatri Chavali, Carol Chen, Noemi Del-Toro, et al. 2013. The mintact project—intact as a common curation platform for 11 molecular interaction databases. *Nucleic acids research*, 42(D1):D358–D363.
- Nanyun Peng and Mark Dredze. 2015. Named entity recognition for chinese social media with jointly trained embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 548–554.
- Nanyun Peng and Mark Dredze. 2016. Improving named entity recognition for chinese social media with word segmentation representation learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Nanyun Peng and Mark Dredze. 2017. Multi-task multi-domain representation learning for sequence tagging. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Erik F Sang and Jorn Veenstra. 1999. Representing text chunks. In *Proceedings of the ninth conference on European chapter of the Association for Computational Linguistics*, pages 173–179. Association for Computational Linguistics.
- Christos Sardianos, Ioannis Manousos Katakis, Georgios Petasis, and Vangelis Karkaletsis. 2015. Argument extraction from news. In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 56–66.
- Bahar Sateli and René Witte. 2015. Semantic representation of scientific literature: bringing claims, contributions and named entities onto the linked open data cloud. *PeerJ Computer Science*, 1:e37.
- Richard Socher, John Bauer, Christopher D Manning, et al. 2013. Parsing with compositional vector grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 455–465.

- Saurabh Srivastava, Puneet Agarwal, Gautam Shroff, and Lovekesh Vig. 2019. Hierarchical capsule based neural network architecture for sequence labeling. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Christian Stab, Christian Kirschner, Judith Eckle-Kohler, and Iryna Gurevych. 2014. Argumentation mining in persuasive essays and scientific articles from the discourse structure perspective. In *ArgNLP*, pages 21–25.
- Simone Teufel and Marc Moens. 1999. Discourse-level argumentation in scientific articles: human and automatic annotation. *Towards Standards and Tools for Discourse Tagging*.
- Simone Teufel and Marc Moens. 2002. Summarizing scientific articles: experiments with relevance and rhetorical status. *Computational linguistics*, 28(4):409–445.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018. The fact extraction and verification (fever) shared task. *arXiv preprint arXiv:1811.10971*.
- Dingquan Wang, Nanyun Peng, and Kevin Duh. 2017. A multi-task learning approach to adapting bilingual word embeddings for cross-lingual named entity recognition. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 383–388.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, Kathryn Funk, Rodney Kinney, Ziyang Liu, William Merrill, et al. 2020. Cord-19: The covid-19 open research dataset. *arXiv preprint arXiv:2004.10706*.

## A Appendices

### A.1 Sentence Representations via Attention

Each of the word representations in the input tensor  $D$  is first reduced from  $d$  to  $d_2$  dimensions. Then the word representations are projected from  $d_2$  dimension to  $p$  dimension. For attention without context,  $p$  dimensional reduced word representations directly perform dot product with a  $p$  dimensional vector to obtain attention scores without using RNN. For attention with context, a simple RNN or LSTM with unit size  $h$  is used to compute attention scores. After obtaining the summarized matrix  $D_{summ}$ , we use bidirectional LSTM with hidden state of size  $H$  to tag the clauses.

**LSTM-Attention** We take the input tensors  $D$  of shape  $c \times w \times d$  and output a matrix  $A$  of shape  $c \times w$  which contains the attention weights of all the words in each clause or sentence. We first project each input word into a lower dimensional space using a projection matrix  $P$  of shape  $d \times p$ .

$$D_l = \tanh(D \cdot P) \in \mathbb{R}^{c \times w \times d}$$

We score  $D_l$  with context that is summarized by an LSTM. Specifically, we score each word in the  $i^{th}$  clause in the context of other words in the same clause or sentence using an LSTM. The score for each word is a function of its  $p$  dimensional representation  $W_j$  and the previous words in the clause represented by the hidden states ( $h_{j-1}^i$ ) in the LSTM cell. The equations are the following:

$$\begin{aligned} D_l^i &= D_l[i, :, :] \in \mathbb{R}^{w \times p} \\ W_j &= D_l^i[j, :] \in \mathbb{R}^p \\ h_j^i &= LSTM(W_j, h_{j-1}^i) \in \mathbb{R}^h \\ h^i &= [h_1^i \ h_2^i \ \dots \ h_w^i] \in \mathbb{R}^{w \times h} \\ a^i &= softmax(h^i \cdot s) \in \mathbb{R}^w \\ A &= [a^1 \ a^2 \ \dots \ a^i \ \dots \ a^c] \in \mathbb{R}^{c \times w} \end{aligned}$$

where  $LSTM$  is an LSTM cell with the unit size of  $h$ .  $s$  is a vector of length  $h$ .

Finally like Dasigi et al. (2017), a  $c \times d$  shaped weighted sum  $D_{summ}$  of the input tensor  $D$  is computed, with the weights computed by the attention mechanism, then it is fed to a clause/sentence-level sequence tagger to tag discourse labels.

$$D_{summ}[i, :] = A[i, :] \cdot D[i, :, :] \in \mathbb{R}^d$$

Hyper-Parameter	Used
$d_{BERT}$	768
$c$	40
$w$	60
$d$	768
$d_2$	300
$p$	200
$h$	75
$H$	350
$lr$	$10^{-3}$
Validation Set Ratio	0.1
Embedding dropout	0.4
Dense dropout	0.4
Attention dropout	0.6
LSTM dropout	0.5
Batch size	10

Table 6: Optimal hyper-parameters of scientific discourse tagger model

### A.2 Implementation and Training Details

The scientific discourse tagging model is implemented using Keras (Chollet et al., 2015) with Tensorflow (Abadi et al., 2015) backend. We use early stopping mechanism with toleration of 2 epochs. We schedule the training by training the model with a learning rate of  $lr$  for 20 epochs. We use Adam (Kingma and Ba, 2014) as our optimizer. The optimal hyper-parameters and the attempted range if applicable are listed in Table 6.