Examining Macro-level Argumentative Structure Features for Argumentative Relation Identification

Tatsuki Kuribayashi[1] Paul Reisert[2] Naoya Inoue[1] Kentaro Inui[1,2]

Abstract: Argumentative relation identification (i.e. detecting *attack*, *support* and *no relation*) in argumentative texts is an important task that has received much attention over the years. In this paper, we examine the effectiveness of macro-level argumentative structure features which have yet to be explored in previous work. For example, we consider features such as the following: when a claim is supported by a premise, it is additionally supported by another premise. To examine the effectiveness, we enhance an existing argumentative relation identification model with macro-level features. Our evaluation demonstrates the potential effectiveness of macro-level argumentative structure features for the task of argumentative relation identification.

1. Introduction

Argument mining is the task of identifying argument structures in argumentative texts. In the literature, several subtasks for argument mining have been extensively studied, such as *argument component type classification* and *argumentative relation identification* [3], [4], [5], [6], [7], [9].

Suppose the argumentative text[*1] in Fig. 1, where argument components (ACs), basic units of arguments, are already identified. Argument component type classification aims at classifying ACs into a premise or claim (e.g. classifying AC_1 into a claim and AC_2 into a premise). Argumentative relation identification aims to identify an argumentative link between two ACs, and if the link exists, classify the relation into two classes: *attack* or *support* (e.g. identifying a support relation from AC_2 to AC_1). Argument mining is useful for many applications such as document summarization, opinion mining, and automated essay scoring [8], [11].

This paper addresses the task of argumentative relation identification. The conventional approaches to relation identification are to train a three-class (i.e. *attack*, *support* or no relation (henceforth, *neither*)) classifier that encodes two input ACs as a feature vector by using local information such as bag-of-words [5], [9] or word embeddings [1] of the input ACs rather than using macro-level information such as the overall structure of an argument. Our hypothesis is that macro-level argumentative structure constructed with argumentative relations will be a hint for predicting a relation. For example, in Fig. (1), AC_1 is supported by AC_2, which will be helpful for predicting the support relation from AC_3 to AC_1, because if an AC is supported by another AC once, the AC is likely to be supported by another. This tendency is observed often when an AC is the main claim of the argument.

Should Germany introduce the death penalty?
AC_1 : *The death penalty is a legal means that as such is not practicable in Germany.*
AC_2 : *For one thing, inviolable human dignity is anchored in our constitution,*
AC_3 : *and furthermore no one may have the right to adjudicate upon the death of another human being.*

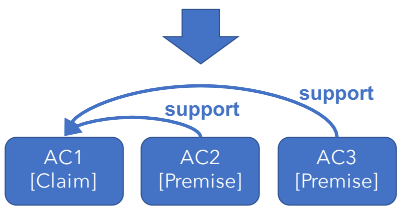

Fig. 1 An example of argument structure with three ACs

In this paper, we investigate the potential effectiveness of macro-level argumentative structures for argumentative relation identification. This work is motivated by the following research questions: (i) what kind of macro-level argumentative structure do we frequently observe in a corpus?; (ii) does exploiting macro-level argumentative structure improve the performance of an argumentative relation classifier? To answer these questions, we analyze macro-level structures in two reliable argumentative corpora and discuss the frequently observed argumentative macro-level structures. We then preliminarily test the effectiveness of macro-level structure features for relation identification by using a simple logistic regression classifier. This study has two main contributions:

- it reveals macro-level argumentative structures frequently observed in two reliable argumentative corpora;
- it demonstrates the potential effectiveness of macro-level features for the task of argumentative relation identification.

1 Tohoku University
2 RIKEN Center for Advanced Intelligence Project
*1 Slightly modified version of the text (micro_b006) taken from Peldszus and Stede (2015) [5]

2. Related work

Previous work has focused on two types of approaches for solving the task of argumentative relation identification. The first approach is to formalize argumentative relation identification as a structured prediction problem (i.e. predicting a graph consisting of argumentative relations from an argumentative text) [4], [5], [7], [9]. To predict a graph, Peldszus et al. [5] use the Maximum Spanning Tree (MST) algorithm, Stab et al. [9] use Integer Linear Programming (ILP), and Potash et al. [7] use Pointer Networks [10]. However, these works do not exploit information from other predicted argumentative relations when predicting a relation. Niculae et al. [4] use factor graphs for structured prediction. They report that higher-order features (e.g. combination features of argumentative links) increase the precision of AC type classification and link identification (i.e. whether argumentative relation exists or not between two ACs). However, it still remains an open question whether such higher-order features are useful for argumentative relation identification.

The second approach is to formalize argumentative relation identification as a pairwise multi-class classification problem [1], [3]. Cocarascu et al. [1] use a Siamese Neural Network-based classifier with a Long Short-Term Memory [2], where the input feature vector is constructed from the information from input ACs only. Nguyen et al. [3] exploit discourse structure features for argumentative relation identification. Their work is closest to our work in the sense of using macro-level information. However, we focus on macro-level information constructed using argumentative relations in a document.

To the best of our knowledge, this is the first work to empirically investigate macro-level argumentative structures frequently used in argumentative texts and to demonstrate the potential effectiveness of argumentative structures constructed using argumentative relations.

3. Data

This study uses two representative corpora used in the argument mining literature. First, we use the arg-micro text (henceforth, *MT*) corpus [5], which contains 112 argumentative short texts (one paragraph each, 5.1 ACs on average). Each text consists of an argumentative topic (e.g. *Should Germany introduce the death penalty?*) and a monologue text discussing this topic (e.g. Fig. 1). Each text is annotated with ACs, AC type (*Claim* or *Premise*), stance (*proponent* or *opponent*) and argumentative relations between ACs(*support, attack, rebut, undercut, normal, example*, or *add*). Following [5], we reduced *rebut* and *undercut* to *attack*, and *normal* and *example* to *support*.[*2] In our experiment, we use 174 attack relations and 290 support relations obtained by this conversion process.

Second, we use the persuasive essay (henceforth, *PE*) corpus [9], which consists of 402 essays (5 paragraphs, 15 ACs on average) posted in online forums. Similar to MT, ACs, AC type (*Premise*, *Claim* or *MajorClaim*), stance (*For* or *Against*) and argumentative relations (*support* or *attack*) are also annotated in PE. There are 219 attack relations and 3,613 support relations.

*2 For add relations from AC_i to AC_j, we first create a link from AC_i to AC_k, the grandparent of AC_i, with the same relation from AC_j to AC_k.

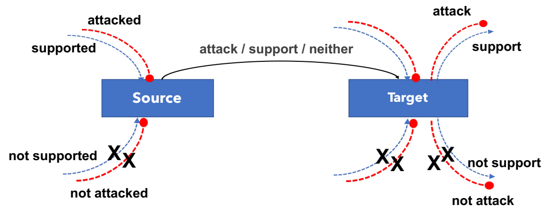

Fig. 2 Macro-level structures examined in this study, shown as dotted lines. Red, oval arrow depicts an *attack* relation and blue, open arrow depicts a *support* relation.

In PE, ACs across a paragraph are *not* annotated with argumentative relations. Therefore, we create *PE-IMP*, a new variant of PE, where (i) claim ACs with *for* stance and a major claim AC are annotated with a *support* relation, and (ii) claim ACs with *against* stance and a major claim AC are annotated with an *attack* relation. This implicit relation is also discussed in Stab and Gurevych (2016) [9]. Thus, analyzing PE-IMP requires analyzing the whole structure of the essay, and analyzing PE simply requires analyzing structures within paragraphs only.

4. Macro-level structure analysis

We analyze the corpora introduced in Section 3 to answer the following questions:

- do we find typical macro-level argumentative structures frequently observed in a corpus?;
- if so, what kind of macro-level argumentative structures are typically observed?

4.1 Methodology

We analyze MT, PE, and PE-IMP introduced in Section 3. We extract all *support*, *attack*, and *neither* relations from each corpora. For notational convenience, we call the starting point of each relation a *source AC* (i.e. an AC which supports/attacks something) and the end point of each relation a *target AC* (i.e. an AC which is supported/attacked by the source AC).

There are a wide variety of ways to define a *macro-level structure*. In this study, we define a macro-level structure as *the combination of a relation type and the state of the source AC and target AC*, as illustrated in Fig. 2. For the state of an AC, we consider the following properties:

- Whether the source AC is attacked/supported by another AC.
- Whether the target AC is attacked/supported by another AC.
- Whether the target AC attacks/supports another AC.

For example, we consider an attack relation where the source AC is attacked by another AC as a macro-level structure.

4.2 Results and discussion

The results are shown in Table 1. The skewed distribution indicates typical macro-level argumentative structures frequently used in each corpus. We also discovered that both corpora have similar tendencies on macro-level structures.

We found that supporting or attacking ACs are supported by another AC in PE-IMP more frequently than in MT and PE. For example, the percentage of an attacking AC which is supported

Table 1 Macro-level structures found in each corpus. The highest percentage among three relation types is in bold. In PE and PE-IMP, we randomly sample 9,000 *neither* relation pairs (about 10% of all *neither* pairs in the corpora).

MT

relation	source is attacked		source is supported		target is attacked (by another AC)		target is supported (by another AC)		target is attacking		target is supporting	
	True	False	True	False	True	False	True	False	True	False	True	False
attack	**62/174 (36%)**	112/174 (64%)	16/174 (9%)	**158/174 (91%)**	27/174 (16%)	**147/174 (84%)**	90/174 (52%)	84/174 (48%)	**64/174 (37%)**	110/174 (63%)	23/174 (13%)	**151/174 (87%)**
support	23/290 (8%)	**267/290 (92%)**	51/290 (18%)	239/290 (82%)	**139/290 (48%)**	151/290 (52%)	**202/290 (70%)**	88/290 (30%)	21/290 (7%)	**269/290 (93%)**	57/290 (20%)	233/290 (80%)
neither	582/2000 (29%)	1418/2000 (71%)	**686/2000 (34%)**	1314/2000 (66%)	354/2000 (18%)	1646/2000 (82%)	373/2000 (19%)	**1627/2000 (81%)**	638/2000 (32%)	1362/2000 (68%)	**1197/2000 (60%)**	803/2000 (40%)

PE

relation	source is attacked		source is supported		target is attacked (by another AC)		target is supported (by another AC)		target is attacking		target is supporting	
	True	False	True	False	True	False	True	False	True	False	True	False
attack	**36/219 (36%)**	183/219 (64%)	32/219 (15%)	187/219 (85%)	**68/219 (31%)**	151/219 (69%)	106/219 (48%)	113/219 (52%)	**38/219 (17%)**	181/219 (83%)	19/219 (9%)	**200/219 (91%)**
support	18/3613 (0%)	**3595/3613 (100%)**	499/3613 (14%)	3114/3613 (86%)	175/3613 (5%)	3438/3613 (95%)	**2978/3613 (82%)**	635/3613 (18%)	43/3613 (1%)	**3570/3613 (99%)**	624/3613 (17%)	2989/3613 (83%)
neither	224/9000 (2%)	2368/9000 (98%)	**2368/9000 (26%)**	6632/9000 (74%)	195/9000 (2%)	**8805/9000 (98%)**	1993/9000 (22%)	**7007/9000 (78%)**	243/9000 (3%)	8757/9000 (97%)	**5616/9000 (62%)**	3384/9000 (38%)

PE-IMP

relation	source is attacked		source is supported		target is attacked (by another AC)		target is supported (by another AC)		target is attacking		target is supporting	
	True	False	True	False	True	False	True	False	True	False	True	False
attack	**184/715 (26%)**	531/715 (74%)	**277/715 (39%)**	438/715 (61%)	**284/715 (40%)**	431/715 (60%)	599/715 (84%)	116/715 (16%)	**148/715 (21%)**	567/715 (79%)	71/715 (10%)	**644/715 (90%)**
support	110/5958 (2%)	**5848/5958 (98%)**	**2294/5958 (39%)**	3664/5958 (61%)	1186/5958 (20%)	4772/5958 (80%)	**5270/5958 (88%)**	688/5958 (12%)	334/5958 (6%)	**5624/5958 (94%)**	3279/5958 (55%)	2679/5958 (45%)
neither	859/9000 (29%)	8148/9000 (71%)	3379/9000 (37%)	**5627/9000 (63%)**	670/9000 (7%)	**8330/9000 (93%)**	2911/9000 (32%)	**6089/9000 (68%)**	729/9000 (8%)	8271/9000 (92%)	**7431/9000 (83%)**	1569/9000 (17%)

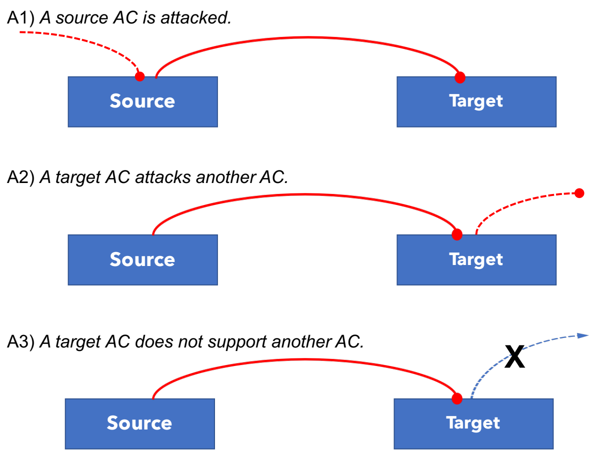

Fig. 3 Common macro-level structures for attack relations.

by another AC is 9% in MT and 15% in PE, but 39% in PE-IMP. This is because the tree structure becomes deeper when considering the whole (including the links between a paragraph's *Claim* and *MajorClaim*) argument structure of an essay. We also observed that one AC supporting another AC is rarely attacked, even in PE-IMP.

The target of an attack relation tends to be less frequently attacked by another AC in MT only (16% in MT, 31% in PE and 40% in PE-IMP). We believe that this is because the argumentative texts in MT are relatively short, so the writer does not write two or more opposing arguments for a specific AC.

In the rest of this section, we discuss common macro-level structures found for each relation.

4.3 Deeper analysis of common macro-level structures

Attack relations. As illustrated in Fig. 3, we observed the following common macro-level structures for an *attack* relation:

A1) *A source AC of an attack relation is attacked.*

A2) *A target AC of an attack relation attacks another AC.*

Fig. 4 shows an example of such common macro-level structures. We found that these structures are typically used when a writer attempts a "preemptive strike", namely gives a possible counter-argument to their claim and then attacks it immediately. In the example text, the writer starts with giving a positive aspect of "*living and studying overseas*" in AC_1. The writer then gives a possible counter-argument to this claim, namely "*living oversees will struggle with loneliness*" (AC_2), which has an attack relation to AC_1. This counter-argument AC_2 is then immediately attacked by AC_3 "*those difficulties will turn into valuable experience*". We observed such a chain of attack relations frequently in the corpora, where the chain typically consists of two attack relations. This pattern could be regarded as macro-level pattern of *support*, as shown in Fig. 5. AC_3 implicitly supports the claim AC_1 by refuting a possible negative opinion (AC_2) against the claim.

In addition, we found the following common macro-level structure (see Fig. 3 for the visualization):

A3) *A target AC of an attack relation does not support another AC.*

This is intuitive because writers do not attack their own support in general. Regardless, we observed that some target ACs of attack relations support another AC (13% of attack relations). We found that such a structure is observed when a writer attempts a preemptive strike to strengthen a supporting AC. For example, consider the text in Fig. 6. The writer begins by stating their main claim, "*Intelligent services must be regulated*" and then gives a support to this claim, i.e. the "*Edward Snowden*" incident. Then he gives a preemptive strike (see A1 and A2) to strengthen the support,

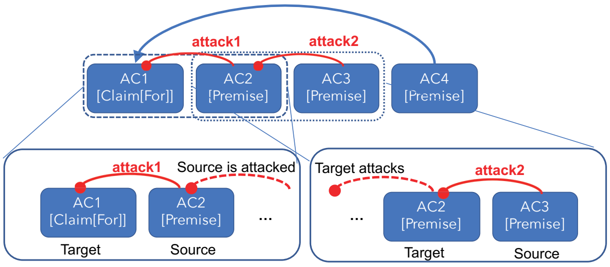

[Second, living and studying overseas is an irreplaceable experience when it comes to learn standing on your own feet.]$_{AC_1[Claim]}$ *[One who is living overseas will of course struggle with loneliness, living away from family and friends]*$_{AC_2[Premise]}$ *[but those difficulties will turn into valuable experiences in the following steps of life.]*$_{AC_3[Premise]}$ *[Moreover, the one will learn living without depending on anyone else.]*$_{AC_4[Premise]}$

Fig. 4 Example of common macro-level structure for attack relations in PE (essay004). The macro-level structure here is "*A source AC tends to be attacked.*" (in the relation *attack1*) and "*A target AC tends to attack.*"(in the relation *attack2*)

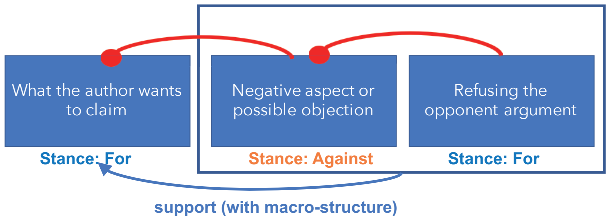

Fig. 5 Macro-level structure where a writer describes a negative aspect of something as a possible counter-argument against a claim (or premise) and refuses it later.

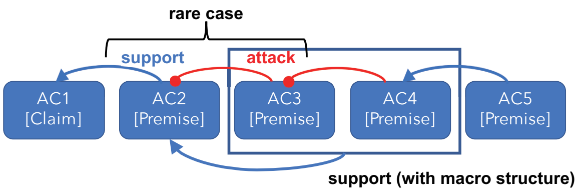

[Intelligence services must urgently be regulated more tightly by parliament;]$_{AC_1[Claim]}$ *[this should be clear to everyone after the disclosures of Edward Snowden.]*$_{AC_2[Premise]}$ *[Granted, those concern primarily the British and American intelligence services,]*$_{AC_3[Premise]}$ *[but the German services evidently do collaborate with them closely.]*$_{AC_4[Premise]}$ *[Their tools, data and expertise have been used to keep us under surveillance for a long time.]*$_{AC_5[Premise]}$

Fig. 6 Example of usage of macro-level structure "*a target AC of an attack relation supports another AC*" in MT (micro_b005).

namely, the writer mentions a possible counter-argument to the support (i.e. "*Snowden's case is about British and American*") and refutes it in the next AC (i.e. "*German services collaborate with them*").

Support relations. As shown in Fig. 7, we observed the following common macro-level structures in *support* relations:

S1) *A source AC of a support relation is not attacked.*

S2) *A target AC of a support relation is supported by another AC.*

S3) *A target AC of a support relation does not attack another AC.*

S1 is intuitive because writers do not use an AC attacked by another AC as a support in general. We found that S2 is frequently observed when the target AC is a main claim of an argument, as

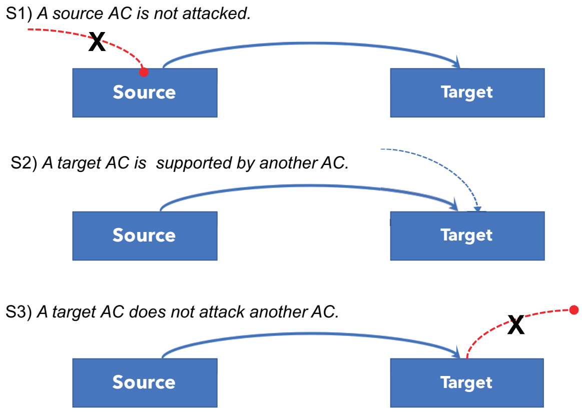

Fig. 7 Common macro-level structures for *support* relations.

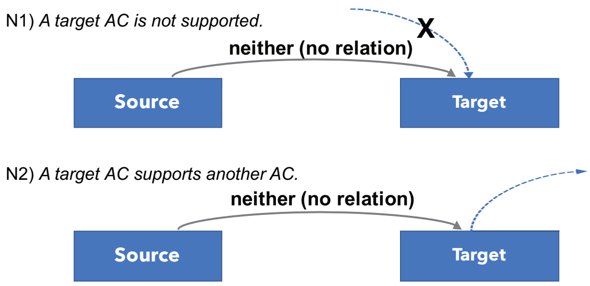

Fig. 8 Common macro-level structures for *neither* relations.

exemplified in Fig. 1. A claim tends to be supported by many premises. Concerning S3, we observe that writers do not provide support for an AC which attacks another AC. We assume this is because writers focus more on supporting their main claim rather than attacking it.

Neither relations. For *neither* relations, we observed the following common macro-level structures, which are illustrated in Fig. 8.

N1) *A target AC of a neither relation is not supported.*

N2) *A target AC of a neither relation supports another AC.*

N1 indicates that a target AC is unlikely to be supported when two ACs have no relation. This implies that two ACs may have a relation when a target AC is supported by another AC. This is consistent with S2, where we found that if a target AC is supported, the target AC tends to be supported additionally (or attacked sometimes). N2 indicates that if a target AC supports another AC, the target AC is not likely to be attacked nor supported. As shown in S1, we observed that a support is not likely to be attacked by other ACs (see Fig. 7). This also suggests that *support* relations are unlikely to form a chain.

5. Preliminary experiment

Given our observation in Section 4, we evaluate the effectiveness of macro-level structure features for the task of argumentative relation identification.

5.1 Setup

As a baseline model, we use a simple logistic regression classifier. The classifier takes two ACs as an input and outputs one of the following three classes: *support*, *attack* or *neither*. We represent a pair of ACs as a binary-valued feature vector, follow-

Table 2 Performance of argumentative relation identification (Binary-classification: *attack* or *support*, 3-class classification: *attack*, *support* or *neither*).

model	Binary classification			Three-class classification			
	macro	attack	support	macro	attack	support	neither
Baseline (MT)	0.670	0.688	0.652	0.492	0.542	0.394	0.539
Baseline (PE)	0.646	0.648	0.645	0.629	0.624	0.535	0.728
Macro (MT)	**0.791**	**0.783**	**0.798**	**0.690**	**0.707**	**0.643**	**0.720**
Macro (PE)	**0.722**	**0.708**	**0.737**	**0.731**	**0.697**	**0.671**	**0.827**

Table 3 Feature weights learned by the classifier of *attack* relation in MT when we use gold macro-level structure features. We list them in descending order of the feature weights and show macro-level structure features only.

rank	feature	weight
1	target does not attack	0.3897
2	source is not attacked	0.3614
3	target is attacked	0.3274
⋮		
29	target is not supported	0.0267
⋮		
8650	target is not supported	-0.2084
⋮		
8653 (worst 3)	target is not attacked	-0.5110
8654 (worst 2)	source is attacked	-0.5151
8655 (worst 1)	target attacks	-0.5635

ing Peldszus and Stede (2015) [5]. We extract surface features such as lemma, part-of-speech tags, and segment length from the source, target, and their adjacent ACs. See the original paper [5] for further details.

On top of the baseline model, we build a *macro-level* model. The macro-level model uses the feature set from the baseline model and additionally encodes a macro-level argumentative structure as the following binary features:

- 1 if the source AC is supported; 0 otherwise;
- 1 if the source AC is not supported; 0 otherwise;
- 1 if the target AC is supported by another AC; 0 otherwise;
- 1 if the target AC is not supported by another AC; 0 otherwise;
- 1 if the target AC supports another AC; 0 otherwise;
- 1 if the target AC doesn't support another AC; 0 otherwise;

The same feature set is introduced for an attack relation (i.e. replacing "support" with "attack"). Henceforth, we refer to these features *second-order features*. To obtain these features, we used the gold-standard information from the corpus. This experiment examines the potential effectiveness of macro-level structure for relation identification.

We ran our experiment on both MT and PE, and we balanced the data for each class with random sampling. The models are evaluated on 5×3-fold nested cross validation. The reported results are averaged over (the outer) 5-fold cross validation. We tuned all hyperparameters using one inner 3-fold CV from the training data. We evaluate the models using two configurations: binary-classification (*attack* or *support*) and 3-class classification (*attack*, *support* and *neither*). We use macro F1 and F1 for each class as an evaluation metric.

5.2 Results

The results are shown in Table 2. We found that the macro-level model significantly outperformed the baseline model on both corpora. This result indicates the potential effectiveness of macro-level features for the argumentative relation identification task.

We examine the feature weights learned by the classifier. Table 3 shows an example of the feature weights of the *attack* relation in MT. Many second-order features are in high rank and tendencies of effective feature are similar to the result of our analysis (see Section 4). For example, we found that macro-level features such as *A target AC does not attack* are ranked higher. This supports the importance of macro-level features for our task.

6. Conclusions and future work

In this paper, we discovered that common macro-level argumentative structure features exist in several argumentative corpora. We discussed why these features occur and found that they can be used as a strategy when a writer claims their opinion towards an argumentative topic. Our preliminary experiment demonstrated the potential effectiveness of macro-level features for the task of argumentative relation identification.

In our evaluation, we adopted a simple encoding of macro-level structure features, which are obtained from nearby argumentative relations. However, there are a wide variety of approaches for exploiting macro-level structures (e.g. sub-tree of the structure) for argumentative relation identification. In our future work, we will explore a model which can capture the overall macro-level structure of a given document using already predicted argumentative relations.

Acknowledgement

This work was supported by JST CREST Grant Number JPMJCR1513 and JSPS KAKENHI Grant Number 16H06614 and 15H01702.

References

[1] Oana Cocarascu and Francesca Toni. Identifying attack and support argumentative relations using deep learning. *ACL*, pages 1385–1390, 2017.

[2] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[3] Huy V Nguyen and Diane J Litman. Context-aware Argumentative Relation Mining. *ACL*, 2016.

[4] Vlad Niculae, Joonsuk Park, and Claire Cardie. Argument Mining with Structured SVMs and RNNs. *ACL*, 2017.

[5] Andreas Peldszus and Manfred Stede. Joint prediction in MST-style discourse parsing for argumentation mining. *EMNLP*, pages 938–948, 2015.

[6] Isaac Persing and Vincent Ng. End-to-End Argumentation Mining in Student Essays. *Naacl*, pages 1384–1394, 2016.

[7] Peter Potash, Alexey Romanov, and Anna Rumshisky. Here's My Point: Joint Pointer Architecture for Argument Mining. *EMNLP*, pages 1375–1384, 2017.

[8] Yi Song, Michael Heilman, Beata Beigman Klebanov, and Paul Deane. Applying Argumentation Schemes for Essay Scoring. *Proceedings of the First Workshop on Argumentation Mining*, (2011):69–78, 2014.

[9] Christian Stab and Iryna Gurevych. Parsing Argumentation Structures in Persuasive Essays. *arXiv preprint, under review*, pages 1–40, 2016.

[10] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700. 2015.

[11] Henning Wachsmuth, Khalid Al-Khatib, and Benno Stein. Using argument mining to assess the argumentation quality of essays. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics*, pages 1680–1691, 2016.