

# Revealing and Predicting Online Persuasion Strategy with Elementary Units

Gaku Morio and Ryo Egawa and Katsuhide Fujita

Tokyo University of Agriculture and Technology

Koganei, Tokyo, Japan

gakumorio@gmail.com, egawa@katfujii.lab.tuat.ac.jp,

katfujii@cc.tuat.ac.jp

## Abstract

In online arguments, identifying how users construct their arguments to persuade others is important in order to understand a persuasive strategy directly. However, existing research lacks empirical investigations on highly semantic aspects of elementary units (EUs), such as propositions for a persuasive online argument. Therefore, this paper focuses on a pilot study, revealing a persuasion strategy using EUs. Our contributions are as follows: (1) annotating five types of EUs in a persuasive forum, the so-called *ChangeMyView*, (2) revealing both intuitive and non-intuitive strategic insights for the persuasion by analyzing 4612 annotated EUs, and (3) proposing baseline neural models that identify the EU boundary and type. Our observations imply that EUs definitively characterize online persuasion strategies.

The annotated dataset, annotation guideline, and implementation of the neural model are available in public.<sup>1</sup>

## 1 Introduction

Persuasive discourses are common in social media. Recently, the automatic identification of such persuasive discourse is receiving attention in the field of computational linguistics (Habernal and Gurevych, 2016b; Tan et al., 2016; Habernal and Gurevych, 2016a; Persing and Ng, 2017; Musi et al., 2018; Hidey and McKeown, 2018; Ji et al., 2018; Durmus and Cardie, 2019; Gleize et al., 2019). Existing studies have examined persuasive arguments with regard to dynamics or lexical features (Tan et al., 2016), claim and premise analysis (Musi and Aakhus, 2018), and concessions (Musi

<sup>1</sup>Dataset and annotation guideline: [http://katfujii.lab.tuat.ac.jp/nlp\\_datasets/](http://katfujii.lab.tuat.ac.jp/nlp_datasets/).  
Implementation: [https://github.com/EdoFrank/emnlp2019\\_morio\\_egawa](https://github.com/EdoFrank/emnlp2019_morio_egawa)

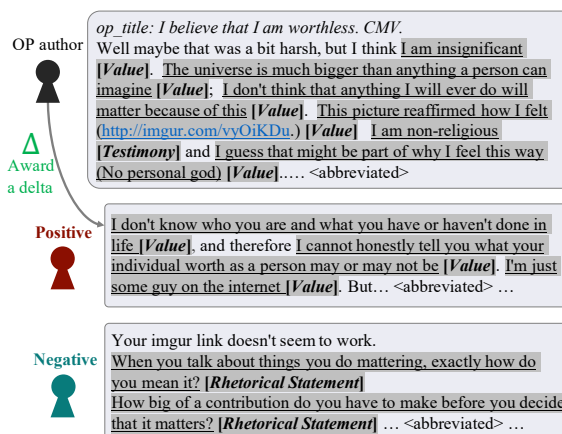


Figure 1: Overview of our EU annotation in *ChangeMyView*. We observed that Testimonies tend to appear in an original post (OP) (see Section 4.)

et al., 2018). However, to the best of our knowledge, fine-grained analyses with an **elementary unit (EU)** semantic, such as propositions, are still in their infancy, although we estimate that EUs are important in obtaining the persuasion strategy. Also, existing studies do not extract token-level EUs via machine learning, which would be a basic step for an automatic understanding of the strategy.

Therefore, this paper shows the fundamental role of EUs in a persuasive forum by annotating five types of token-level EUs (i.e., Testimony, Fact, Value, Policy, and Rhetorical Statement) in a persuasive forum and also provides a baseline neural model to identify the EUs automatically.<sup>2</sup>

In this study, we use the effective dataset of *ChangeMyView* (Tan et al., 2016), which is a subreddit where original post (OP) users provide a controversial view in the title to change their

<sup>2</sup>Although we use the partial data of the previous study (Egawa et al., 2019), this work is distinct from the previous one because we focus on the novel analysis of EUs and a novel task to predict EUs. See Section 4 and 5.

	source	#unit	level	IAA
<b>ours</b>	CMV	4612	token	.677
Al Khatib et al.	editorial	14313	token	.56
Hidey et al.	CMV	2615	token	.65
Park et al.	eRulemaking	4899	clause	.648

Table 1: Corpus comparison related to our study.

perspective through opponent users’ arguments. OP users award a delta ( $\Delta$ ) point when their view is changed. Figure 1 presents an overview of *ChangeMyView* where the positive post is an awarded post and the negative post is a non-awarded one.

Our analyses to find out the role of EUs in *ChangeMyView* show both intuitive and non-intuitive phenomena (Section 4.) We also inject a structural feature into the neural model to improve the performance; the results show that EUs have characteristic positional roles (Section 5.)

## 2 Related Work

EUs have been considered in investigating the characteristics of argumentation in an argument mining discipline. Park et al. (2015); Park and Cardie (2018) deemed Testimony to be a type of objective evidence using personal state or experience. Al Khatib et al. (2016) proposed the argument model for analyzing the strategy of a news editorial. The model considers token-level discourse units of six different types. The study that is most closely related to our research is (Hidey et al., 2017) because it annotated semantic types of claims and premises in *ChangeMyView*; however, the authors did not take objective propositions such as facts into consideration.

Table 1 summarizes the comparison in terms of size, annotation granularity level, and inter-annotator agreement (IAA) of the related corpora. The table shows that our corpus is middle sized but has sufficient granularity and the reasonable IAA reliability.

## 3 Annotating EUs

### Typology of EUs

This study considers five types of EUs. For the introduction of the scheme, the motivation is derived from our expectation that we can feature persuasive arguments by considering personal experience, facts (Park and Cardie, 2018) and rhetoric (Blankenship and Craig, 2006). The five types of EUs are defined as follows:

Ⓣ: **Testimony** is an objective proposition related to the author’s personal state or experience such as the following: *I do not have children.*

ⓕ: **Fact** is a proposition describing objective facts that can be verified using objective evidence and therefore captures the evidential facts in persuasions: *Empire Theatres in Canada has a “Reel Babies” showing for certain movies.*

Ⓥ: **Value** is a proposition that refers to subjective value judgments without providing a statement on what should be done: *it is absolutely terrifying.*

Ⓟ: **Policy** offers a specific course of action to be taken or what should be done: *intelligent students should be able to see that.*

Ⓡ: **Rhetorical Statement** implicitly states the subjective value judgment by expressing figurative phrases, emotions, or rhetorical questions: *does it physically hurt men to be raped by women (as in PIV sex)?*

### Annotation Process

We extracted 115 *ChangeMyView* threads from the train set of Tan et al. (Tan et al., 2016) through a simple random sampling.<sup>3</sup> Each thread contains a triple of OP, positive, and negative. Therefore, we use 345 posts for the annotation.

EUs are annotated by 19 non-native annotators who are proficient in English, excluding the authors. Three annotators independently annotated 87 threads; the remaining 28 threads were annotated by 8 experts who were selected from the 19 annotators. Each annotator was asked to read the annotation guideline prior to the actual annotation. We also held several meetings to train the annotators and re-annotated the erroneous annotations when required. Given that the annotators are non-native speakers, the posts are translated into their language, and the translated documents are only used as a reference for the annotators. For the 83 threads, a gold standard is established by merging three posts using a majority vote.

We consider the token-level annotation (Stab and Gurevych, 2017), rather than the clause level one, to extract accurate minimal EU boundary and to remove irrelevant boundaries such as *for example* and *therefore*.

<sup>3</sup>More precisely, 419 threads were randomly extracted from the train data range 2013/01/01–2015/05/07, and the older threads were actually used due to time constraints.




Figure 2: Distribution of the number of EUs. Most of the propositions in *ChangeMyView* are subjective.

	(F)	(T)	(P)	(R)	(V)
OP vs. reply		**		**	
positive vs. negative					

Table 2: Significance of the proportion of EU frequency of *OP* vs. *reply* and *positive* vs. *negative*. Significant at  $p < 0.01$  if \*\*.

### Annotation Results

As a result, 4612 EUs are annotated. Figure 2 shows the distribution of the number of EUs. We see that Value dominates in *ChangeMyView*, indicating most of the propositions are subjective.

At the end of the annotation, the IAA of Krippendorff’s  $\alpha_U$  (Krippendorff, 2004) was .677. Since sentence- or clause-level annotations (Park and Cardie, 2018) cannot accurately distinguish an inference step, we assume that the reasonable agreement is due to the token-level annotation.

## 4 Strategy Analysis of Persuasion: Revealing the Role of EUs

### Frequency of EUs Characterizes OP vs. Reply but Insignificant for Positive vs. Negative

Table 2 shows the significance of the proportion of each EU type for the post types, namely, *OP* vs. *reply* (*positive* and *negative*) and *positive* vs. *negative*. Given that the proportion is a real value, the Mann–Whitney U test is employed. A significant difference exists in Testimony in *OP* vs. *reply* and Rhetorical Statements in *OP* vs. *reply*. Therefore, *OP* authors are likely to assert their view with experiences, i.e., *OP*: *For many years, I’ve regularly skipped breakfast*. It is also intuitive that Rhetorical Statement exists to a high degree in persuasion (*reply*) posts because its strategy is for the persuasion.

In contrast, no significance is observed in a positive vs. negative case in Table 2. Although Habernal and Gurevych (2016b) stated that facts appear in persuasive comments, our result of Facts is insignificant as well as non-intuitive. We assume that this is because users are persuaded by effective indications of EUs, rather than simply present-




Figure 3: Kernel density estimate plots showing the distribution of EU positions in a post. Fact and Testimony tend to appear in the beginning of the post.

position	MannWhitneyU	KS
0.00 - 0.50	** (F)	* (F)
0.25 - 0.75		
0.50 - 1.00		

Table 3: Significance of the EU position between a positive and negative post. Significant at  $p < 0.05$  if \*.

ing objective statements.

### Positional Role of EUs is Characteristic

To examine the positional role of EUs as a persuasion strategy, we investigate the position with a one-dimensional axis, that is, a normalized position wherein the beginning of a post is 0.0 and the end is 1.0. Figure 3 shows the resulting histogram of the positions for each EU type. It presents the characteristic distributions: Fact and Testimony are possibly located at the beginning of the post, e.g., positive: *Eating first thing in the morning stabilizes blood sugar levels, which regulates appetite and energy*, indicating the persuasion strategy to state solid facts prior to claiming assertion. Moreover, a Policy appeared at the bottom, indicating that what should be done is regarded as a conclusion, e.g., negative: *Just keep in mind that eating breakfast may increase your ability to concentrate, even if you think you are doing ok right now*. Value is the type that dominates in most posts, and it appears in uniform distribution.

The statistical tests were performed to determine the existence of a difference between a positive and negative in terms of the position. In general, given that addressing the entire positions simultaneously could simplify the property, we set the window size to 0.5 and investigated three ranges as follows: beginning (0.00 – 0.50), middle (0.25 – 0.75), and end (0.50 – 1.00). For each window, we employed the Mann–Whitney U test and Kolmogorov–Smirnov (KS) significance for the average positions in threads which have at least five propositions, as described in Table 3.

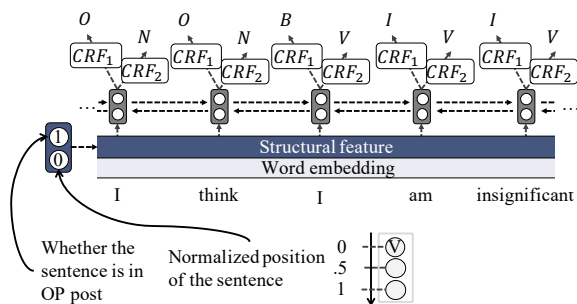


Figure 4: BiLSTM-CRF with structural features.

The table shows that the average position at the beginning of the post has significant differences in Fact.

## 5 Neural EU Extraction to Automatically Obtain the Strategy

In practice, we have to extract EUs computationally to obtain the strategy automatically. We consider the EU extraction problem as a sequence tagging problem in analogy using a related argument mining work (Eger et al., 2017; Habernal and Gurevych, 2017). The sequence tagging problem involves labeling each token in an input document. In each time step  $i$ , the label  $y_i \in \mathcal{Y}$ , where each token is associated with this label, is predicted as follows:  $\mathcal{Y} = \{(b, t) \mid b \in \{B, I, O\}, t \in \{F, T, P, R, V, N\}\}$ , where  $b$  denotes the unit identification whether the current token is non-EU ( $O$ ), begins ( $B$ ), or continues ( $I$ ) an EU. In addition,  $t$  is the unit classification, i.e., Non-EU ( $N$ ), Fact ( $F$ ), Testimony ( $T$ ), Policy ( $P$ ), Rhetorical Statement ( $R$ ), and Value ( $V$ ).

### The Multi-task Learning Approach

We employed neural multi-task learning (MTL) models for the sequence tagging problem because several existing argument mining studies achieved advanced performance when the MTL models were used (Schulz et al., 2018). Figure 4 shows the overview of the proposed model. This model comprises two modules. (i) The bi-directional LSTM (BiLSTM) module encodes the input tokens in a given document. (ii) The module for conditional random fields (CRFs) for each subtask (unit identification and classification) is applied to execute label classification, denoted by the layers toward the top of Figure 4. We call our model BiLSTM-CRF (BLC).

### (i) Module of the BiLSTM

In our model, LSTMs were used to encode the input tokens. We employed the BiLSTM that uses two LSTMs to represent the forward and backward contexts in a sentence. Given an input token representation  $x_i$  (i.e., a word embedding vector) at time  $i$ , we obtained the context-aware hidden state  $h_i = LSTM_F(x_{1:i}) \circ LSTM_B(x_{n:i})$ , where  $n$  denotes the length of the input tokens.

### (ii) Module of the CRFs

This output module discriminates the BIO tag and EU type using the hidden representations generated by the BiLSTM. When a CRF layer was utilized, the past and future tags to predict the present tag are used efficiently. Therefore, the CRF layer is effective for solving a non-independent sequence tagging problem, including BIO tagging, i.e., the label after “ $B$ ” is obviously “ $I$ .” Note that we provided an MTL model that shares all the trainable parameters of the BiLSTM module but does not share any parameter of the CRF module for each task.

### Experimental Setting

The train and test threads were randomly divided per a ratio of 8 : 2, where 30% of the train set is considered to be a development set. Herein, we set up two cases: *post-level*, a case in which an entire post is provided for the input and *sentence-level*, a case in which each sentence is given. The motivation is that the EUs cannot cross over a sentence. However, the sentence-level case would loss structural information. To remedy this problem, we focus on two structural features: whether or not the sentence is in OP post and the normalized position of the sentence in a post.

Each system was trained for 50 iterations with random seeds and the model that exhibited the optimal performance using the development set was selected (Schulz et al., 2018). The hyperparameters of our network are 100-dim GloVe embeddings (Pennington et al., 2014), the hidden layer size 200 for the BiLSTMs, the dropout rate of 0.5, and the optimizer is Adam (Kingma and Ba, 2014).

### Results: Overall Performance

Table 4 shows that our proposed sentence-level BLC with structural information performs best regarding macro F1 in either boundary identification or unit classification tasks. The presence of

BLC variant	unit identification				unit classification					
	B	I	O	macro	(V)	(R)	(P)	(T)	(F)	macro
sentence-level + structural	75.4	92.8	62.1	<b>*76.8</b>	<b>*80.7</b>	<b>61.5</b>	<b>*16.0</b>	<b>43.1</b>	<b>33.2</b>	<b>*49.2</b>
sentence-level	<b>75.6</b>	92.8	61.0	76.5	80.2	60.9	13.0	41.3	31.3	47.7
post-level	67.8	<b>92.9</b>	<b>*64.7</b>	75.1	79.9	49.4	2.1	35.3	32.2	43.8

Table 4: Performance of the average F1. Max score if **bold** and significant differences with  $p < 0.05$  if \*.




Figure 5: Performance improvement (F1 of a unit classification) by adding the structural features.

structural features provides an excellent boost in classifying unit types. For the post-level BLC, the model yields a better score in predicting non-EU parts because post-level discrimination can capture the entire post and thus identify irrelevant boundaries, such as supplementary notes, in posts.

### Results: Analysis of Structural Features

Figure 5 shows the performance improvement by structural features in a distinct experiment from Table 4. In general, we see our structural features perform better for EUs that have characteristic positional distributions such as Policy and Testimony as shown in Figure 3. Interestingly, the structural features did not perform well in Fact classification for OPs and positive posts. The result implies Facts in the OPs and positive posts would be solid and can be predicted without structural information.

## 6 Conclusion

This paper demonstrated empirical work on the five types of elementary units (EUs) in online arguments. To find out the role of EUs, we annotated 4612 EUs and showed both intuitive and non-intuitive results. We also proposed baseline neural models to discriminate EUs. Experimental results showed that positional roles of some EUs are essential in detecting EUs. The annotated corpus and neural models will be applicable for future persuasion reasoning or evaluation of persuasiveness.

### Acknowledgments

This work was supported by JST CREST (JPMJCR15E1) and JST AIP-PRISM (JPMJCR18ZL), Japan.

## References

- Khalid Al Khatib, Henning Wachsmuth, Johannes Kiesel, Matthias Hagen, and Benno Stein. 2016. A news editorial corpus for mining argumentation strategies. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3433–3443, Osaka, Japan. The COLING 2016 Organizing Committee.
- Kevin L. Blankenship and Traci Y. Craig. 2006. Rhetorical question use and resistance to persuasion: An attitude strength analysis. *Journal of Language and Social Psychology*, 25(2):111–128.
- Esin Durmus and Claire Cardie. 2019. A corpus for modeling user and language effects in argumentation on online debating. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 602–607, Florence, Italy. Association for Computational Linguistics.
- Ryo Egawa, Gaku Morio, and Katsuhide Fujita. 2019. Annotating and analyzing semantic role of elementary units and relations in online persuasive arguments. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 422–428, Florence, Italy. Association for Computational Linguistics.
- Steffen Eger, Johannes Daxenberger, and Iryna Gurevych. 2017. Neural end-to-end learning for computational argumentation mining. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11–22, Vancouver, Canada. Association for Computational Linguistics.
- Martin Gleize, Eyal Shnarch, Leshem Choshen, Lena Dankin, Guy Moshkovich, Ranit Aharonov, and Noam Slonim. 2019. Are you convinced? choosing the more convincing evidence with a Siamese network. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 967–976, Florence, Italy. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2016a. What makes a convincing argument? empirical analysis and detecting attributes of convincingness in web argumentation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1214–1223, Austin, Texas. Association for Computational Linguistics.

- Ivan Habernal and Iryna Gurevych. 2016b. Which argument is more convincing? analyzing and predicting convincingness of web arguments using bidirectional lstm. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1589–1599. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2017. Argumentation mining in user-generated web discourse. *Comput. Linguist.*, 43(1):125–179.
- Christopher Hidey and Kathleen McKeown. 2018. Persuasive influence detection: The role of argument sequencing.
- Christopher Hidey, Elena Musi, Alyssa Hwang, Smaranda Muresan, and Kathy McKeown. 2017. Analyzing the semantic types of claims and premises in an online persuasive forum. In *Proceedings of the 4th Workshop on Argument Mining*, pages 11–21, Copenhagen, Denmark. Association for Computational Linguistics.
- Lu Ji, Zhongyu Wei, Xiangkun Hu, Yang Liu, Qi Zhang, and Xuanjing Huang. 2018. Incorporating argument-level interactions for persuasion comments evaluation using co-attention model. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3703–3714. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Klaus Krippendorff. 2004. Measuring the reliability of qualitative text analysis data. *Quality and Quantity*, 38(6):787–800.
- Elena Musi and Mark Aakhus. 2018. Discovering argumentative patterns in energy polylogues: A macroscope for argument mining. *Argumentation*, 32(3):397–430.
- Elena Musi, Debanjan Ghosh, and Smaranda Muresan. 2018. Changemyview through concessions: Do concessions increase persuasion? *Dialogue and Discourse (D&D)*, Vol 9(1).
- Joonsuk Park, Cheryl Blake, and Claire Cardie. 2015. Toward machine-assisted participation in erulemaking: An argumentation model of evaluability. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law, ICAIL '15*, pages 206–210, New York, NY, USA. ACM.
- Joonsuk Park and Claire Cardie. 2018. A corpus of erulemaking user comments for measuring evaluability of arguments. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*.
- 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. Association for Computational Linguistics.
- Isaac Persing and Vincent Ng. 2017. Why can't you convince me? modeling weaknesses in unpersuasive arguments. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 4082–4088.
- Claudia Schulz, Steffen Eger, Johannes Daxenberger, Tobias Kahse, and Iryna Gurevych. 2018. Multitask learning for argumentation mining in low-resource settings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 35–41. Association for Computational Linguistics.
- Christian Stab and Iryna Gurevych. 2017. Parsing argumentation structures in persuasive essays. *Computational Linguistics*, 43(3):619–659.
- Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2016. Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *WWW*, pages 613–624. ACM.