

Please quote as: [Wambsganss, T.; Niklaus, C.; Söllner, M.; Handschuh, S. & Leimeister, J. M. (2020): A Corpus for Argumentative Writing Support in German. In: 28th International Conference on Computational Linguistics (COLING)].

A Corpus for Argumentative Writing Support in German

Thiemo Wambsganss¹, Christina Niklaus^{1,3}, Matthias Söllner²,
Siegfried Handschuh^{1,3} and Jan Marco Leimeister^{1,2}

¹ University of St.Gallen

{thiemo.wambsganss, christina.niklaus, siegfried.handschuh, janmarco.leimeister}@unisg.ch

² University of Kassel

{soellner, leimeister}@uni-kassel.de

³ University of Passau

{christina.niklaus, siegfried.handschuh}@uni-passau.de

Abstract

In this paper, we present a novel annotation approach to capture claims and premises of arguments and their relations in student-written persuasive peer reviews on business models in German language. We propose an annotation scheme based on annotation guidelines that allows to model claims and premises as well as support and attack relations for capturing the structure of argumentative discourse in student-written peer reviews. We conduct an annotation study with three annotators on 50 persuasive essays to evaluate our annotation scheme. The obtained inter-rater agreement of $\alpha = 0.57$ for argument components and $\alpha = 0.49$ for argumentative relations indicates that the proposed annotation scheme successfully guides annotators to moderate agreement. Finally, we present our freely available corpus of 1,000 persuasive student-written peer reviews on business models and our annotation guidelines to encourage future research on the design and development of argumentative writing support systems for students.

1 Introduction

In today's world most information is readily available. Consequently, the sole reproduction of information is losing attention. This results in a shift of job profiles towards interdisciplinary, ambiguous and creative tasks (vom Brocke et al., 2018). Therefore, educational institutions need to evolve in their curricula when it comes to the compositions of skills and knowledge conveyed. In particular, teaching higher order thinking skills to students, such as critical thinking, collaboration or problem-solving, has become more important (Fadel et al., 2015). This has already been recognized by the Organization for Economic Co-operation and Development (OECD), which included these skills as a major element of their Learning Framework 2030 (OECD, 2018). One subclass represents the skill of arguing in a structured, reflective and well-formed way. Argumentation is not only an essential part of our daily communication and thinking, but it also contributes significantly to the competencies of communication, collaboration and problem-solving (Kuhn, 1992). Building on studies by Aristotle, the ability to form convincing arguments is recognized as the foundation for persuading an audience of novel ideas, and it plays a major role in strategic decision-making and analyzing different standpoints. To develop skills such as argumentation, it is of great importance for the individual student to receive continuous feedback throughout their learning journey, also called formative feedback (Black and Wiliam, 2009; Hattie and Timperley, 2007). However, this is naturally hindered due to traditional large-scale lectures and due to the growing field of distance learning scenarios such as massive open online courses (MOOCs) (Seaman et al., 2018). In fact, educational institutions, such as universities, face the challenge of providing individual formative feedback effectively (Fortes and Tchantchane, 2010), since every student would need a personal tutor to have an optimal learning environment for learning how to argue (Vygotsky, 1980).

One possible path for providing individual feedback is to leverage recent developments in Computational Linguistics in the form of computer-assisted writing which enables the development of writing support systems that provide tailored feedback about textual documents (Song et al., 2014; Stab and Gurevych, 2014a). Argumentation Mining (AM), a research field in Computational Linguistics, aims at automatically identifying arguments in unstructured texts (Lippi and Torroni, 2015). An argument is a set of statements made up of three parts: a claim, a set of evidence or premises (e.g., facts) and an

inference from the evidence to the claim (Toulmin, 1984). Claim and premise represent the argument components. The claim is the central component of an argument, representing an arguable text unit, while the premises are propositions that either support or attack the claim, underpinning its plausibility. Support and attack are argumentative relations that model the discourse structure of arguments. In the identification of these argumentation structures, two main tasks can be distinguished:

- **Argument component classification:** classification of argumentative text into claims and premises
- **Argument relation classification:** identification of support and attack relationships between pairs of argument components

To design, train and test AM algorithms, the availability of high-quality labeled corpora is crucial. Therefore, numerous prior works have dealt with creating annotated data sets. However, they are all limited to a particular genre, ranging from well-structured legal (Palau and Moens, 2009; Ashley and Walker, 2013) and scientific documents (Kirschner et al., 2015; Hougbo and Mercer, 2014), to rather ambiguous, vague and less formal social web content (Wachsmuth et al., 2014; Aharoni et al., 2014; Cabrio and Villata, 2014; Habernal and Gurevych, 2015a). Corpora that are applicable for the design and development of argumentative writing support systems are scarce (Stab and Gurevych, 2017a; Wambsganss and Rietsche, 2020; Wambsganss et al., 2020b). The only collection from the education domain that is annotated for argumentative discourse structures was presented in Stab and Gurevych (2014a). It is composed of 90 persuasive essays written by students in English language and later extended to include 402 essays (Stab and Gurevych, 2017a). However, these corpora are 1) annotated in English language only and 2) not derived from a specific learning scenario which would leverage the effective use for an argumentative writing support system. Consequently, there is a *lack of linguistic corpora* for training models that provide students with adaptive feedback about the quality of their argumentation in common scenarios in large-scale lectures or the growing field of MOOCs (Seaman et al., 2018), such as peer reviews where students provide each other argumentative feedback on a specific task, e.g., on a previously developed business model (Rietsche and Söllner, 2019).

Creating gold standards and test collections requires a formal representation model as well as corresponding annotation guidelines. In this paper, we introduce an argumentation annotation scheme for persuasive student-generated peer reviews extracted from a common learning scenario. Moreover, we present a corpus of 1,000 student-written peer reviews that are annotated for argumentation components and their relations. Our contribution is threefold: 1) we derive an annotation scheme for a new data domain for AM based on argumentation theory and previous work on annotation schemes for persuasive student essays (Stab and Gurevych, 2017a; Stab and Gurevych, 2014a), 2) we present an annotation study based on 50 persuasive peer reviews and three annotators to show that the annotation of student peer reviews is reliably possible, 3) we present our final and freely available corpus of 1,000 student peer reviews collected in our lecture about business innovation in German language. We therefore hope to encourage future research on student-generated argumentative texts and on writing support systems to train argumentation skills of students based on AM.

2 Related Work

2.1 Argumentation Mining

AM is a research field in Computational Linguistics, gaining momentum in a lot of areas, including the legal domain (Mochales Palau and Ieven, 2009), newswire articles (Reed et al., 2008a; Deng and Wiebe, 2015; Sardianos et al., 2015), user-generated web content (Wachsmuth et al., 2014; Habernal and Gurevych, 2015b; Abbott et al., 2016), or online debates (Cabrio and Villata, 2014; Eckle-Kohler et al., 2015). AM aims at automatically identifying arguments in unstructured textual documents based on the classification of argumentative and non-argumentative text units and the extraction of argument components and their relations. Recently, researchers have built increasing interest in intelligent writing assistance based on AM (Song et al., 2014; Stab and Gurevych, 2014a; Stab and Gurevych, 2014b), since it enables argumentative writing support systems that provide tailored feedback about arguments in

student-generated texts. However, the effectiveness of using this technology in a certain learning scenario for educational purposes has rarely been assessed (Stab and Gurevych, 2017b; Lippi and Torroni, 2015), as argumentation corpora from student-generated texts in the field of education are rather uncommon (Lawrence and Reed, 2019).

2.2 Argument Annotated Corpora and Annotation Schemes

As Lawrence and Reed (2019) state, “*one of the challenges faced by current approaches to argument mining is the lack of large quantities of appropriately annotated arguments to serve as training and test data.*” Since the availability of labeled corpora is crucial for designing, training and evaluating AM algorithms, numerous prior works have dealt with creating annotated data sets, such as the Araucaria corpus (Reed et al., 2008b), the European Court of Human Rights (ECHR) corpus of Mochales and Moens (2008), the Debatepedia corpus (Cabrio and Villata, 2012), the ChangeMyView corpus (Egawa et al., 2019) or the persuasive essays corpus of Stab and Gurevych (2014a) with 90 essays and the corpus of Stab and Gurevych (2017a) with 402 persuasive student essays. These corpora have been widely used for various AM tasks, such as the identification of argument components (Rooney et al., 2012), corpus wide AM (Ein-Dor et al., 2019) or end-to-end AM (Persing and Ng, 2016).

Creating gold standards and test collections requires a formal representation model as well as corresponding annotation guidelines. While a number of well-defined models exist in the field of AM (Freeman, 2001; Perelman, 1971; Pollock, 1995; Walton and Macagno, 2015; Walton, 1996; Wambsganss et al., 2020a), there is no general argumentation annotation scheme across all domains and genres of texts. Instead, the proposed representations differ in granularity, expression power and categorization (Lawrence and Reed, 2019). Therefore, conducting annotation studies with several annotators when introducing new annotation schemes is crucial for the quality of argumentation corpora.

2.3 Argument Annotated Corpora in Education

With the exception of the corpora proposed in Stab and Gurevych (2014a) and Stab and Gurevych (2017a), previously presented argument annotated datasets are not easily applicable for the development of argumentative writing support systems for students in a real-world case, since they are 1) not extracted from an educational learning scenario in which the annotation could be used for training a model that provides students feedback on the texts, and 2) often not annotated at the level of discourse (Stab and Gurevych, 2017a; Lawrence and Reed, 2019), which is necessary for example to give students feedback on insufficiently supported claims. Stab and Gurevych (2014a) identified the lack of linguistic corpora in the domain of student-written texts for designing and developing argumentative writing support systems for adaptive feedback by leveraging AM techniques (Stab and Gurevych, 2014a). Therefore, they introduced an annotation scheme for annotating argument components and their relationships in persuasive student essays. Afterwards, several studies built on their corpus, including e.g., Carlile et al. (2018) who take a subset of the essays and annotate their persuasiveness, or Ke et al. (2018) who train a persuasiveness scoring model on them. However, the transfer of argumentation corpora to other educational domains (e.g., common learning scenarios such as peer reviews) and other languages falls short in current literature.

2.4 Essay Scoring Corpora

Besides annotated datasets of argumentative student-written texts, several corpora exist in the field of automatic essay scoring. The goal of this task is to automatically rate textual documents in the form of holistic scores based on their content and form (Horbach et al., 2017). Most corpora are built on student-written content, e.g. the Cambridge Learner Corpus (CLC) (Yannakoudakis et al., 2011) with 1,244 English essays, the Swedish high school corpus with 1,702 essays (Östling et al., 2013), or the TOEFL11 corpus with 1,100 English essays written by language students (Blanchard et al., 2013). However, the corpora are usually annotated with a holistic score on the level of full documents only, e.g., *low*, *medium*, *high* in the TOEFL11 corpus, while specific argumentation structures are commonly ignored. In fact, the persuasiveness of essays is - if at all considered - usually only regarded as one sub variable of the overall document (e.g., in the annotations of Östling et al. (2013) and Yannakoudakis et al. (2011)).

Consequently, these corpora lack applicability for the development of argumentative writing support systems. The argumentation quality is often only annotated as a qualitative score on a 1 to 3 range (e.g., Horbach et al. (2017)) and therefore not sufficient enough to train a sophisticated supervised machine learning model for argumentative writing assistants. Corpora from the field of automatic essay scoring usually neither distinguish between different types of argument components (e.g., claims or premises) nor are they built on a rich argumentation annotation scheme. Nguyen and Litman (2018) demonstrated the value of AM for automated persuasive essay scoring by evaluating different AM features for improving essay scores. However, essay scoring corpora mostly do not focus on the annotation of argumentation relations, and therefore, disqualify for a foundation for sophisticated models for feedback on argumentative discourse through writing support systems (e.g., feedback on unsupported claims). We aim to address this literature gap by presenting and evaluating an annotation scheme as well as an annotated argumentation corpus built on student-written texts with the objective of enabling researchers to develop novel argumentative writing support systems for students.

3 Corpus Construction

We compiled a corpus of 1,000 student-generated peer reviews in which students provide each other feedback on previously developed business models. Peer reviews are a modern learning scenario in large-scale lectures, enabling students to reflect on their content, receive individual feedback from peers and thus deepen their understanding of the content (Rietsche and Söllner, 2019). Moreover, they are easy to set up in traditional large-scale learning scenarios or in the growing field of distance-learning scenarios such as MOOCs (Seaman et al., 2018). This can be applied to train skills such as argumentation. However, since not many suitable corpora are available to provide argumentation feedback that A) contain annotated persuasive student peer reviews, B) consist of a sufficient size to be able to use trained models in a real-world scenario and C) follow an annotation guideline for guiding the annotators towards an adequate agreement, we propose an new annotation scheme to model argument components as well as argumentation relations that reflects the argumentative discourse structure in persuasive peer reviews. We based our annotation scheme on the model of Toulmin (1984) and the studies of Stab and Gurevych (2014a; Stab and Gurevych (2017a)). To build a reliable corpus, we followed a four step methodology, illustrated in Figure 1: 1) We examined scientific literature and theory on how to model argumentation discourse structures in texts from different domains, 2) we randomly sampled 50 student-generated peer reviews and, on the basis of our findings from literature and theory, developed a set of annotation guidelines consisting of rules and limitations on how to annotate argumentation discourse structures, 3) we applied, evaluated and improved our guidelines with three native-speakers in three consecutive workshops to resolve annotation ambiguities, 4) we applied the final annotation scheme based on our 15 pages guidelines to a corpus of 1,000 student-generated peer reviews.¹




Figure 1: Process of compiling a corpus of 1,000 annotated student-generated peer reviews.

3.1 Data Source

We gathered a corpus of 1,000 student-generated peer reviews written in German. The data was collected in one of our mandatory business innovation lectures in a master program at our university. In this lecture, around 200 students develop and present a new business model for which they receive three peer reviews each. Here, a fellow student from the same course elaborates on the strengths and weaknesses

¹The annotation guidelines as well as the entire corpus can be accessed at https://github.com/thiemowa/argumentative_student_peer_reviews.

of the business model and gives persuasive recommendations on what could be improved. We sampled a random subset of 1,000 of these reviews out of around 7,000 documents collected between 2014 and 2018.

3.2 Annotation Scheme

Our objective is to model the argumentation discourse structures of student-generated peer reviews by annotating argumentation components and their relations. Most of the peer reviews in our corpus follow a similar structure. They describe several strengths or weaknesses of the business model under consideration, backing them up by examples, statistics, intuitions or citations. These strengths and weaknesses can also be formulated to make a certain recommendation for improvements of the business model. An argumentation component, e.g., a strength, weakness or suggestion is only regarded as a “*claim*” if it contains a certain standpoint, which can also represent a complete sentence. Our basic annotation scheme is illustrated in Figure 2.




Figure 2: Argumentation annotation scheme including argument components (claim, premise) and argumentative relations (support or attack).

3.2.1 Argument Components

For argumentation components, we follow established models in argumentation theory which provide detailed definitions of argument components (e.g., (Toulmin, 1984; Walton et al., 2008; Weinberger and Fischer, 2006; Perelman, 1971; Pollock, 1995; Walton, 1996; Freeman, 2001; Walton and Macagno, 2015; Van Eemeren and Grootendorst, 2016). These theories generally agree that a basic argument consists of multiple components and that it includes a *claim* that is supported or attacked by at least one *premise*. Also in the domain of student-written peer reviews, we found that a *claim* is the central component of an argument. It is a controversial statement (e.g., claiming a strength or weakness of a business model - see examples below) that is either true or false and should not be accepted by the receiver of the feedback without additional support or backing. The *premise* supports the validity of the *claim* (e.g., by providing a statistic, quote or a value-based intuition). It is a reason given by the author to persuade the reader of her *claim*. Below are two examples of *claims* from our corpus (1. being a strength, 2. being a weakness) and their supporting *premises*.²

1. “**The value proposition is very well done.**_{claim (strength)} *It is short and concise and the advantages or benefits are well highlighted.*_{premise (example)} *As a customer I would like to try the product after reading it.*”_{premise (intuition)}
2. “**Unfortunately, the value proposition canvas is very poorly filled in.**_{claim (weakness)} *The points are far too little elaborated and far too little described or explained. The customer jobs should be described much more precisely.*”_{premise (example)}

3.2.2 Argumentative Relations

The basic argumentation structure in our corpus of student-generated peer reviews consists of several claims, each independently supported by one or more premises. Since in our data domain weaknesses and strengths of a business model are discussed, the texts do generally not present a major claim as is

²Since the original texts are written in German, we translated the examples to English for the sake of this paper.

the case for example in English essays annotated by Stab and Gurevych (2014a) and Stab and Gurevych (2017a). Instead, the documents we deal with consist of a set of claims supported by one or more premises. However, premises not only *support* a statement, but may also *attack* a claim, e.g., when used as a stylistic device or to illustrate uncertainty in the argumentation. Hence, more complicated constellations of claims and premises are possible, in which a claim is supported by several different premises or by a chain of premises, in which each premise is in turn supported by another premise. In the same way, a claim can be supported by one premise and attacked by another, or supported by a premise which is attacked by another premise. Nevertheless, the simplest form consists of a claim supported by a single premise. To provide an overview, we illustrated three basic examples of annotated relations in our corpus (see Figure 3). Some statistics on the occurrences of different patterns in our dataset can be found in Table 6.




Figure 3: Examples of possible argumentation relations. The arrow signifies a relation, which can either be *support* or *attack*. The rectangle denotes an argument component in the form of a *claim* or *premise*.

3.3 Annotation Process

Three native German speakers annotated the peer reviews independently from each other for *claims* and *premises* as well as their *argumentative relationships* in terms of support and attack, according to the annotation guidelines we specified. Inspired by Stab and Gurevych (2017a), our guidelines consisted of 15 pages, including definitions and rules for what is an argument, which annotation scheme is to be used and how argument components and argumentative relations are to be judged. Several private training sessions and three team workshops were performed to resolve disagreements among the annotators and to reach a common understanding of the annotation guidelines. We used the brat rapid annotation tool, since it provides a graphical interface for marking up text units and linking their relations (Stenetorp et al., 2012). After the first 50 reviews were annotated by all three annotators, we calculated the inter-annotator agreement (IAA) scores. As we obtained satisfying results, we proceeded with a single annotator who marked up the remaining 950 documents. Following Stab and Gurevych (2014a), we annotated argument components in terms of claims and premises. The specific component types, e.g., *weakness*, *strength* for claims or *intuition*, *statistic* for premise were not annotated in this study. However, we believe this would be a useful addition in future work. In Figure 4 we display an example of an entire peer review with the corresponding annotations.

4 Corpus Analysis

We analyzed the results of the annotation process in order to examine (1) the *reliability of the corpus* and (2) the *major disagreements* in argument component and relation annotations between the annotators. In addition, we calculated some statistics of the final corpus.

4.1 Inter-Annotator Agreement

To evaluate the reliability of the argument component and argumentative relation annotations, we followed the approach of Stab and Gurevych (2014a).

4.1.1 Argument Components

With regard to the argument components, two strategies were used. Since there were no predefined markables, annotators not only had to identify the *type of argument component*, but also its *boundaries*. In order to assess the latter, we use Krippendorff’s α_U (Krippendorff, 2004), which allows for assessing the reliability of an annotated corpus, considering the differences in the markable boundaries. To evaluate the annotators’ agreement in terms of the selected category of an argument component for a given sentence, we calculate percentage agreement and two chance-corrected measures, multi- π (Fleiss, 1971)




Figure 4: Fully annotated example of a peer review according to our annotation scheme and guidelines. The left column represents the claims (C1 - C6), while the premises are listed on the right (P1 - P7). N signifies a non-argumentative text unit.

and Krippendorff's α (Krippendorff, 1980). We decided to operate at sentence level, since only 20.56% of the sentences in the corpus contain annotations of different argument components. Thus, evaluating the reliability at sentence level served as a good approximation of the IAA. At the level of individual sentences, 33.45% contain a claim, 35.64% a premise and 55.52% a non-annotation. 20.56% of the sentences contain several annotations, with a combination of a premise and a non-argumentative span (32.21%) and a combination of a claim and a non-argumentative span (35.78%) representing the majority of those cases. Only 4.08% contain both a premise and a claim. At the token level, the following class distribution is achieved: 43.54% claim, 45.42% premise and 11.04% are not annotated.

Table 1 displays the resulting IAA scores. We obtained an IAA of 78.65% for the claims and 76.63% for the premises. The corresponding multi- π scores are 55.47% and 51.06%. Regarding Krippendorff's α , a score of 55.49% and 51.08% is obtained, indicating a moderate agreement for both categories. With a score of 44.04% and 47.76%, the unitized α of both the claim and premise annotations is slightly smaller compared to the sentence-level agreement. Thus, the boundaries of argument components are

	%	Multi- π	Krippendorff's α	Krippendorff's α_U
Claim	0.7865	0.5547	0.5549	0.4404
Premise	0.7663	0.5106	0.5108	0.4776
None	0.8572	0.6410	0.6412	0.3590

Table 1: Inter-rater agreement of argument component annotations.

less precisely identified in comparison to the classification into argument types. Yet the scores still suggest that there is a moderate level of agreement between the annotators. With a score of $\alpha_U=35.90\%$, the boundaries of non-argumentative units are less reliably detected. In contrast, the agreement scores targeting the component types are considerably higher for the non-argumentative spans as compared to the claims and premises (85.72%, 64.10%, 64.12%), indicating a substantial agreement between the annotators. Hence, we conclude that the annotation of the argument components in student-generated peer reviews is reliably possible.

4.1.2 Argumentative Relations

	%	Multi- π	Krippendorff's α
Support	0.9413	0.4903	0.4903
Non-support	0.9413	0.4903	0.4903

Table 2: Inter-rater agreement of argumentative relation annotations.

To evaluate the reliability of the argumentative relations, we used the set of all pairs between argument components that were possible during the annotation task according to our annotation scheme, i.e. all pairs between a claim and a premise and between two premises. In total, the markables include 4,792 pairs of which 7.41% are annotated as support relations and 0.35% as attack relations. 92.24% of the possible pairs were not identified by an annotator. Since the number of attack relations is so small, we decided to focus on the support relations, distinguishing only between the two types *support* and *non-support*. We obtained an IAA of 94.13% for both support and non-support relations. The corresponding multi- π and Krippendorff's α scores both amount to 49.03% (see Table 2). Therefore, we conclude that argumentative relations, too, can be reliably annotated in student-generated argumentative peer reviews.

4.2 Disagreement Analysis

4.2.1 Confusion Probability Matrices

To analyze the disagreement between the three annotators, we created confusion probability matrices (CPM) (Cinková et al., 2012) for argument components and argumentative relations. A CPM contains the conditional probabilities that an annotator assigns a certain category (column) given that another annotator has chosen the category in the row for a specific item. In contrast to traditional confusion matrices, a CPM also allows for the evaluation of confusions if more than two annotators are involved in an annotation study (Stab and Gurevych, 2014a).

	Claim	Premise	None
Claim	0.5500	0.3296	0.1205
Premise	0.3741	0.4570	0.1689
None	0.1850	0.2285	0.5865

Figure 5: CPM for argument component annotations.

	Supports	None
Support	0.3528	0.6472
Non-support	0.0606	0.9394

Figure 6: CPM for argumentative relation annotations.

While there is a broad agreement between the annotators in distinguishing non-argumentative discourse units from argument components, the major disagreement is between claims and premises (Table 5). This result is in accordance with the findings in Stab and Gurevych (2014a). This could be expected since a claim can also serve as premise for another claim, and it is difficult to distinguish these two concepts in the presence of reasoning chains. For instance, examples (1–3) establish a reasoning chain in which (1) is supported by (2) and (2) is supported by (3):

- (1) “The client would have to wait twice as long as usual for his parcel.”
- (2) “LivePrice orders the products from the provider that offers them at the best price and packs them.”
- (3) “In the Business Model Canvas it is described that LivePrice packs the products itself and ships them.”

Considering the structure of this argumentation, (1) can be classified as a claim. However, if (1) is omitted, (2) becomes a claim that is supported by (3). Thus, the distinction between claims and premises depends not only on the context and the intention of the author but also on the structure of a specific argument (Stab and Gurevych, 2014a).

The CPM for argumentative relations (see Table 6) reveals that there is a rather high confusion between support relations and none classified relations. This result is again in line with the findings of Stab and Gurevych (2014a). In sum, according to our error analysis, the annotation of argumentative relations yields less reliable results than that of argument components.

4.2.2 Manual Analysis

We manually analyzed the most common differences in the annotated text spans between the annotators on a random sample of 10 documents that are composed of 292 sentences. Our findings are as follows: while some annotators included *clausal conjunctions*, such as “because”, “since” or “as”, in their annotations of premises, others did not mark them up. Moreover, in some cases *phrases with verbs of reported speech or cognition*, such as “I think that” or “I believe that”, were included in the annotations of argument components, whereas in other cases this phrasal type was ignored in the annotations. The same holds for *introductory phrases* at the beginning of sentences or clauses, e.g. “however”, “moreover”, “for example”. In addition, some annotators included *punctuation marks* denoting the end of sentences or clauses in their annotated spans, while others did not. In addition, we encountered a larger number of cases where annotators *mistakenly missed the first or last character(s) of a word* starting or ending, respectively, the annotated span. Often, the annotations were not even consistent within a single annotator.

4.3 Corpus Statistics

	total	mean	standard deviation	min	max	median
Sentences	20,217	20.217	9.4299	2	69	19.0
Tokens	246,980	246.98	101.4358	40	720	236.0

Table 3: Distribution of *sentences* and *tokens* in the created corpus. Mean, standard deviation, minimum, maximum and median refer to the number of sentences and tokens, respectively, per document.

Tables 3, 4, 5 and 6 present some statistics of the final corpus. It consists of 1,000 student-written peer reviews in German that are composed of 20,125 sentences with 246,980 tokens in total. Hence, on average, each document has 20 sentences and 272 tokens. A total of 7,996 claims (31.64%) and 8,479 premises (33.55%) were annotated. 8,796 textual spans (34.81%) were not identified as an argument component (“None”). On top of that, 8,291 support (97.03%) and 254 attack (2.97%) relationships were marked up by the annotators. The little percentage of attack relations are explainable due to the domain of our annotated texts. The nature of peer reviews is to provide feedback about a certain topic by highlighting strengths, weaknesses and suggesting improvements. Usually students used premises to support their claims, in fact, only about 3% of the claims were further elaborated by discussing them more controversially through attacking premises.

	total	mean	standard deviation	min	max	median	percent
Claim	7,996	7.996	3.6487	1	25	7.0	0.3164
Premise	8,479	8.479	4.8346	0	26	8.0	0.3355
None	8,796	8.796	3.9545	2	26	8.0	0.3481
All	25,271	8.4237	4.1878	0	26	8.0	1.0

Table 4: Distribution of types of *argument components* in the created corpus.

	total	mean	standard deviation	min	max	median	percent
Support	8,291	8.291	4.7771	0	26	7.0	0.9703
Attack	254	0.254	0.6196	0	5	0.0	0.0297
All	8545	4.2725	5.2681	0	26	2.0	1.0

Table 5: Distribution of types of *argumentative relations* in the created corpus.

Table 6 presents the distribution of support relationships in our corpus. With 46.30%, the majority of claims is not supported by any premise. 39.07% of the claims are backed up by exactly one premise, while 10.93% are supported by two premises. Patterns with more than two supporting premises rarely occur in our dataset. The majority of premises is not backed up by another premise (75.07%). However, in 21.08% of the cases, there is one supporting premise. In that way, more complex reasoning chains are established.

	#0 support	#1 support	#2 support	#3 support	#4 support	>#4 support
Claim	0.4630	0.3907	0.1093	0.0230	0.0074	0.0066
Premise	0.7507	0.2108	0.0289	0.0073	0.0019	0.0005

Table 6: Distribution of patterns of *support relationships* in the created corpus. The columns denote the percentage of claims and premises, respectively, that are supported by the specified number of premises.

5 Application of the Corpus

After constructing and analysing our corpus, we leveraged the novel data to train a machine learning model. Our objective was to embed a classification algorithm in the backend of a user-centered adaptive writing support system to provide students with formative argumentation feedback in the writing process. Therefore, we trained and tuned a model with different text features (Stab and Gurevych, 2014a; Fromm et al., 2019). We performed a multiclass classification on the sentence level to detect the argument components and their relations. For argument component classification, we found that a Support Vector Machine (SVM) achieved the best results, with an accuracy of 65.4% on the test set. Regarding the argumentative relation classification, a binary classification task, an SVM again achieved the best results on our corpus, obtaining an accuracy of 72.1% on the test set. A detailed description of the features and the embedding of our corpus in a user-centered adaptive writing support tool, as well as its effect on students’ argumentation writing skills can be found in Wambsganss et al. (2020b).

6 Conclusion

We introduce an argumentation annotation scheme as well as an annotated corpus of persuasive student-written peer reviews extracted from a real-world learning scenario which can be leveraged to provide students feedback on the quality of their argumentation (e.g., by highlighting insufficiently backed claims) through a writing support system. Regarding the educational domain, previously presented argument annotation schemes and argumentation corpora fall short on several aspects: they are not derived from a modern learning scenario, they do not follow a systematic methodology based on detailed inter-rater agreement studies or they do not include annotations of argumentative relations. To overcome these limitations, we present a corpus of 1,000 student-written peer reviews that are annotated for argument components and their relations. Our contribution is threefold: 1) we derive an annotation scheme for a new data domain for AM based on argumentation theory and previous work on annotation schemes for persuasive student essays (Stab and Gurevych, 2014b; Stab and Gurevych, 2017a), 2) we present an annotation study based on 50 persuasive peer reviews that were annotated by three native German speakers, demonstrating that the annotation of student-generated peer reviews is reliably possible, and 3) we present our final and freely available corpus of 1,000 student peer reviews collected in our lecture about business innovation in German language. We hope to encourage fellow researchers to leverage our annotation scheme and argumentation corpus to design and develop writing support systems for students in large-scale learning scenarios.

References

- Rob Abbott, Brian Ecker, Pranav Anand, and Marilyn Walker. 2016. Internet argument corpus 2.0: An sql schema for dialogic social media and the corpora to go with it. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Helene Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- Ehud Aharoni, Anatoly Polnarov, Tamar Lavee, Daniel Hershcovich, Ran Levy, Ruty Rinott, Dan Gutfreund, and Noam Slonim. 2014. A Benchmark Dataset for Automatic Detection of Claims and Evidence in the Context of Controversial Topics. In *Proceedings of the First Workshop on Argumentation Mining*, pages 64–68, Baltimore, Maryland. Association for Computational Linguistics.
- Kevin D. Ashley and Vern R. Walker. 2013. Toward constructing evidence-based legal arguments using legal decision documents and machine learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, ICAIL '13*, pages 176–180, New York, NY, USA. ACM.
- Paul Black and Dylan Wiliam. 2009. Developing the theory of formative assessment. *Educational Assessment, Evaluation and Accountability*, 21(1):5–31.
- Daniel Blanchard, Joel Tetreault, Derrick Higgins, Aoife Cahill, and Martin Chodorow. 2013. Toefl11: a Corpus of Non-Native English. *ETS Research Report Series*, 2013(2):i–15.
- Elena Cabrio and Serena Villata. 2012. Natural language arguments: A combined approach. *Frontiers in Artificial Intelligence and Applications*, 242:205–210.
- Elena Cabrio and Serena Villata. 2014. Towards a Benchmark of Natural Language Arguments. *CoRR*, abs/1405.0.
- Winston Carlile, Nishant Gurrupadi, Zixuan Ke, and Vincent Ng. 2018. Give me more feedback: Annotating argument persuasiveness and related attributes in student essays. *ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, 1:621–631.
- Silvie Cinková, Martin Holub, and Vincent Križ. 2012. Managing uncertainty in semantic tagging. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 840–850, Avignon, France, April. Association for Computational Linguistics.
- Lingjia Deng and Janyce Wiebe. 2015. MPQA 3.0: An Entity/Event-Level Sentiment Corpus. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1323–1328, Denver, Colorado. Association for Computational Linguistics.
- Judith Eckle-Kohler, Roland Kluge, and Iryna Gurevych. 2015. On the role of discourse markers for discriminating claims and premises in argumentative discourse. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2236–2242, Lisbon, Portugal. 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, July. Association for Computational Linguistics.
- Liat Ein-Dor, Eyal Shnarch, Lena Dankin, Alon Halfon, Benjamin Sznajder, Ariel Gera, Carlos Alzate, Martin Gleize, Leshem Choshen, Yufang Hou, Yonatan Bilu, Ranit Aharonov, and Noam Slonim. 2019. Corpus Wide Argument Mining – a Working Solution. (1).
- Charles Fadel, Bernie Trilling, and Maya Bialik. 2015. *Four-Dimensional Education: The Competencies Learners Need to Succeed*.
- J.L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382.
- Pauline Carolyne Fortes and Abdellatif Tchantchane. 2010. Dealing with large classes: A real challenge. *Procedia - Social and Behavioral Sciences*, 8:272–280.
- James B Freeman. 2001. Argument structure and disciplinary perspective. *Argumentation*, 15(4):397–423.

- Hansjörg Fromm, Thiemo Wambsganss, and Matthias Söllner. 2019. Towards a Taxonomy of Text Mining Features. In *European Conference of Information Systems (ECIS)*, pages 1–12.
- Ivan Habernal and Iryna Gurevych. 2015a. Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2127–2137, Lisbon, Portugal. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2015b. Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2127–2137, Lisbon, Portugal, September. Association for Computational Linguistics.
- John Hattie and Helen Timperley. 2007. The Power of Feedback. *Review of Educational Research*, 77(1):81–112.
- Andrea Horbach, Dirk Scholten-Akoun, Yuning Ding, and Torsten Zesch. 2017. Fine-grained essay scoring of a complex writing task for native speakers. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 357–366.
- Hospice Hougbo and Robert Mercer. 2014. An automated method to build a corpus of rhetorically-classified sentences in biomedical texts. In *Proceedings of the First Workshop on Argumentation Mining*, pages 19–23, Baltimore, Maryland. Association for Computational Linguistics.
- Zixuan Ke, Winston Carlile, Nishant Gurrupadi, and Vincent Ng. 2018. Learning to give feedback: Modeling attributes affecting Argument persuasiveness in student essays. *IJCAI International Joint Conference on Artificial Intelligence*, 2018-July:4130–4136.
- Christian Kirschner, Judith Eckle-Kohler, and Iryna Gurevych. 2015. Linking the thoughts: Analysis of argumentation structures in scientific publications. In *Proceedings of the 2nd Workshop on Argumentation Mining held in conjunction with the 2015 Conference of the North American Chapter of the Association for Computational Linguistics – Human Language Technologies (NAACL HLT 2015)*, pages 1–11.
- Klaus Krippendorff. 1980. *Content Analysis: An Introduction to Methodology*. Sage Publications, Inc., Beverly Hills, CA.
- Klaus Krippendorff. 2004. Measuring the reliability of qualitative text analysis data. *Quality and Quantity*, 38(6):787–800, Dec.
- Deanna Kuhn. 1992. Thinking as Argument. *Harvard Educational Review*, 62(2):155–179, 7.
- John Lawrence and Chris Reed. 2019. Argument mining: A survey. *Computational Linguistics*, 45(4):765–818.
- Marco Lippi and Paolo Torroni. 2015. Argumentation Mining: State of the Art and Emerging Trends. *IJCAI International Joint Conference on Artificial Intelligence*, 2015-Janua(2):4207–4211.
- Raquel Mochales and Marie Francine Moens. 2008. Study on the structure of argumentation in case law. *Frontiers in Artificial Intelligence and Applications*, 189(1):11–20.
- Raquel Mochales Palau and Aagje Ieven. 2009. Creating an argumentation corpus: do theories apply to real arguments? A case study on the legal argumentation of the ECHR. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Law (ICAIL 2009), Twelfth international conference on artificial intelligence and law (ICAIL 2009), Barcelona, Spain, 8-12 June 2009*, pages 21–30. ACM.
- Huy V. Nguyen and Diane J. Litman. 2018. Argument mining for improving the automated scoring of persuasive essays. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 5892–5899.
- OECD. 2018. *The Future of Education and Skills - Education 2030*.
- Robert Östling, André Smolentzov, Björn Tyrefors Hinnerich, and Erik Höglin. 2013. Automated Essay Scoring for Swedish. *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 42–47.
- Raquel Mochales Palau and Marie-Francine Moens. 2009. Argumentation mining: The detection, classification and structure of arguments in text. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law, ICAIL '09*, pages 98–107, New York, NY, USA. ACM.
- Chaim Perelman. 1971. The New Rhetoric. In Yehoshua Bar-Hillel, editor, *Pragmatics of Natural Languages*, pages 145–149. Springer Netherlands, Dordrecht.

- Isaac Persing and Vincent Ng. 2016. End-to-End Argumentation Mining in Student Essays. In *NAACL HLT*, pages 1384–1394.
- John L Pollock. 1995. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press, Cambridge, MA, USA.
- Chris Reed, Raquel Mochales Palau, Glenn Rowe, and Marie-Francine Moens. 2008a. Language Resources for Studying Argument. In Bente Maegaard Joseph Mariani Jan Odijk Stelios Piperidis Daniel Tapias Nicoletta Calzolari (Conference Chair) Khalid Choukri, editor, *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Chris Reed, Raquel Mochales Palau, Glenn Rowe, and Marie Francine Moens. 2008b. Language resources for studying argument. *Proceedings of the 6th International Conference on Language Resources and Evaluation, LREC 2008*, pages 2613–2618.
- Roman Rietsche and Matthias Söllner. 2019. Insights into Using IT-Based Peer Feedback to Practice the Students Providing Feedback Skill. *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- Niall Rooney, Hui Wang, and Fiona Browne. 2012. Applying kernel methods to argumentation mining. *Proceedings of the 25th International Florida Artificial Intelligence Research Society Conference, FLAIRS-25*, pages 272–275.
- 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, Denver, Colorado. Association for Computational Linguistics.
- Julia E. Seaman, I. E. Allen, and Jeff Seaman. 2018. Higher Education Reports - Babson Survey Research Group. Technical report.
- Yi Song, Michael Heilman, Beata Beigman Klebanov, and Paul Deane. 2014. Applying argumentation schemes for essay scoring. In *Proceedings of the First Workshop on Argumentation Mining*, pages 69–78, Baltimore, Maryland, June. Association for Computational Linguistics.
- Christian Stab and Iryna Gurevych. 2014a. Annotating Argument Components and Relations in Persuasive Essays. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers* ,, pages 1501–1510.
- Christian Stab and Iryna Gurevych. 2014b. Identifying Argumentative Discourse Structures in Persuasive Essays. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)(Oct. 2014)*, Association for Computational Linguistics, p.(to appear), pages 46–56.
- Christian Stab and Iryna Gurevych. 2017a. Parsing Argumentation Structures in Persuasive Essays. *Computational Linguistics*, 43(3):619–659, 9.
- Christian Stab and Iryna Gurevych. 2017b. Recognizing insufficiently supported arguments in argumentative essays. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 980–990, Valencia, Spain, April. Association for Computational Linguistics.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. brat: a web-based tool for NLP-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France, April. Association for Computational Linguistics.
- Stephen E. Toulmin. 1984. *Introduction to Reasoning*.
- Frans H Van Eemeren and Rob Grootendorst. 2016. *Argumentation, communication, and fallacies: A pragma-dialectical perspective*. Routledge.
- Jan vom Brocke, Wolfgang Maaß, Peter Buxmann, Alexander Maedche, Jan Marco Leimeister, and Günter Pecht. 2018. Future Work and Enterprise Systems. *Business and Information Systems Engineering*, 60(4):357–366.
- Lev Semenovich Vygotsky. 1980. *Mind in society: The development of higher psychological processes*. Harvard university press.

- Henning Wachsmuth, Martin Trenkmann, Benno Stein, Gregor Engels, and Tsvetomira Palakarska. 2014. A review corpus for argumentation analysis. In *Proceedings of the 15th International Conference on Computational Linguistics and Intelligent Text Processing - Volume 8404, CICLing 2014*, pages 115–127, New York, NY, USA. Springer-Verlag New York, Inc.
- Douglas Walton and Fabrizio Macagno. 2015. A Classification System for Argumentation Schemes. *New Publisher: Ios Press*, 6(3):219–245.
- Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press, Cambridge.
- D N Walton. 1996. *Argumentation Schemes for Presumptive Reasoning*. Argumentation Schemes for Presumptive Reasoning. L. Erlbaum Associates.
- Thiemo Wambsganss and Roman Rietsche. 2020. Towards designing an adaptive argumentation learning tool. In *40th International Conference on Information Systems, ICIS 2019*, pages 1–9, Munich, Germany.
- Thiemo Wambsganss, Nikolaos Molyndris, and Matthias Söllner. 2020a. Unlocking Transfer Learning in Argumentation Mining : A Domain-Independent Modelling Approach. In *15th International Conference on Wirtschaftsinformatik*,, number MI, Potsdam, Germany.
- Thiemo Wambsganss, Christina Niklaus, Matthias Cetto, Matthias Söllner, Jan Marco Leimeister, and Siegfried Handschuh. 2020b. AL : An Adaptive Learning Support System for Argumentation Skills. In *ACM CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Armin Weinberger and Frank Fischer. 2006. A framework to analyze argumentative knowledge construction in computer-supported collaborative learning. *Computers and Education*, 46(1):71–95.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading ESOL texts. *ACL-HLT 2011 - Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 1:180–189.