

# A School Student Essay Corpus for Analyzing Interactions of Argumentative Structure and Quality

Maja Stahl<sup>1</sup>, Nadine Michel<sup>2</sup>, Sebastian Kilsbach<sup>2</sup>,  
Julian Schmidtke<sup>1</sup>, Sara Rezat<sup>2</sup>, and Henning Wachsmuth<sup>1</sup>

<sup>1</sup>Leibniz University Hannover, Institute of Artificial Intelligence

<sup>2</sup>Paderborn University, Institute for German Language and Comparative Literature

{m.stahl,h.wachsmuth}@ai.uni-hannover.de, julian.schmidtke@stud.uni-hannover.de

{nadine.michel, sebastian.kilsbach, sara.rezat}@uni-paderborn.de

## Abstract

Learning argumentative writing is challenging. Besides writing fundamentals such as syntax and grammar, learners must select and arrange argument components meaningfully to create high-quality essays. To support argumentative writing computationally, one step is to mine the argumentative structure. When combined with automatic essay scoring, interactions of the argumentative structure and quality scores can be exploited for comprehensive writing support. Although studies have shown the usefulness of using information about the argumentative structure for essay scoring, no argument mining corpus with ground-truth essay quality annotations has been published yet. Moreover, none of the existing corpora contain essays written by school students specifically. To fill this research gap, we present a German corpus of 1,320 essays from school students of two age groups. Each essay has been manually annotated for argumentative structure and quality on multiple levels of granularity. We propose baseline approaches to argument mining and essay scoring, and we analyze interactions between both tasks, thereby laying the ground for quality-oriented argumentative writing support.

## 1 Introduction

Writing argumentative texts, in particular argumentative essays, constitutes an essential part of school students' writing education. However, learning to write arguments of high quality can be challenging (Zhu, 2001; Ferretti et al., 2007; Peloghitis, 2017; Alexander et al., 2023). It requires various skills, from writing fundamentals, such as syntax and grammar, to argumentation-specific skills, such as meaningfully organizing and structuring arguments and counter-considerations (Rezat, 2011). This takes time and effort to master (Ka-kan-dee and Kaur, 2014; Dang et al., 2020). Given teachers' limited time to give students feedback on their writing, automatic argumentative writing support could



Figure 1: Exemplary annotated school student essay on the use of school funding, taken from our corpus. The text is from the FD-LEX corpus (Becker-Mrotzek and Grabowski, 2018), translated from German for display.

benefit students as it offers guidance at their own pace and convenience (Wambsganss et al., 2022a).

Argumentative writing support systems employ argument mining to analyze input texts (Stab, 2017; Wambsganss and Niklaus, 2022; Weber et al., 2023), that is, computational methods that identify argumentative components and their relations. Common components are *major claim* (main standpoint of the text, also known as *thesis*), *claim* (controversial statement), and *premise* (reason for justifying or refuting the claim) along with their argumentative relations *support* and *attack*. This knowledge enables the systems to give feedback on the structure of a text, e.g., by highlighting unwarranted claims (Stab and Gurevych, 2017a), or by analyzing the number of argumentative components quantitatively (Stab and Gurevych, 2017a; Wambsganss and Niklaus, 2022; Weber et al., 2023).

Unlike argument mining, automated essay scoring explicitly evaluates essay quality, either holis-

tically (Uto et al., 2020; Yang et al., 2020; Wang et al., 2023) or specific linguistic aspects, such as coherence (Li et al., 2018; Farag et al., 2018), grammar (Ajit Tambe and Kulkarni, 2022), and organization (Persing et al., 2010; Rahimi et al., 2015). Combining argument mining with essay scoring may enable support systems to give students comprehensive feedback on their writing. In addition, it helps identify how different argumentative structures influence the overall essay quality and which structures are common for different levels of quality (Wachsmuth et al., 2016). However, student essay corpora for argument mining are scarce and do not include ground-truth essay quality annotations (Stab and Gurevych, 2017a). Moreover, no corpus with structure annotations for essays written by school students has been published yet.

To fill this research gap, we present a German corpus of 1,320 school student essays with manual annotations for argumentative structure and essay quality. The essays have been systematically selected from an existing corpus (Becker-Mrotzek and Grabowski, 2018), equally distributed over two age groups (fifth-graders and ninth-graders) and binary genders, three per student. We present an extensive annotation scheme focused on school student essays that covers argumentative structure on four levels of granularity as well as five essay quality aspects, as shown in Figure 1. To achieve consistent annotations, we developed annotation guidelines in close dialogue with our expert annotators from the field of language education. This led to high agreement between the annotations.

Our analyses of the corpus provide various insights into the correlation between the different levels of argumentative structure and essay quality, as well as the interaction between these two types of annotation. We experiment with fine-tuned transformers and adapters as baseline approaches to mining argumentative structure and scoring essay quality. Moreover, we demonstrate that the information on argumentative structure helps predicting the essay quality, which is in line with what previous studies showed on other corpora (Wachsmuth et al., 2016; Beigman Klebanov et al., 2016; Nguyen and Litman, 2018). This result underlines the usefulness of our corpus annotations for quality-oriented argumentative writing support.

More explicitly, this work aims to answer (i) how the argumentative structure and essay quality of school student essays can be modeled, (ii)

how different levels of argumentative structure and essay quality correlate for school student essays, and (iii) how this correlation can be exploited to automatically score the essay quality.

Altogether, this paper’s main contributions are:

- A corpus for studying argumentative structure and essay quality on school student essays
- Empirical insights into the interactions of argumentative structure and essay quality
- Baseline approaches to argument mining and essay scoring<sup>1</sup>

## 2 Related Work

Argumentative writing is a key capability that is taught in school across age groups and disciplines (Becker-Mrotzek et al., 2010; Rezat, 2011). A common educational form of argumentative text is the essay, where school students should introduce a thesis, to which they provide pro and con arguments, and finally conclude (Townsend et al., 1993; Schröter, 2021). The components of an argumentative text take on different roles (e.g., *claim* or *premises*), and they may operationalize different actions (e.g., *conceding* or *reasoning*) (Feilke, 2017). Learning to write argumentative text is complex and requires continuous and detailed feedback (Kellogg et al., 2010; Wambsganss et al., 2022b).

Analyzing the argumentative structure of texts computationally, also known as argument(ation) mining, is a crucial and widely-studied step in providing automatic support for argumentative writing (Stede and Schneider, 2019). Student essays are a prominent domain for argument mining. A respective annotated corpus of 402 English student essays is available (Stab and Gurevych, 2017a), for which also quality issues such as insufficient claim support have been modeled (Stab and Gurevych, 2017b; Gurcke et al., 2021). Additionally, student corpora are available for more specific domains, such as argumentative legal texts (Weber et al., 2023), persuasive peer reviews on business models (Wambsganss et al., 2020), and business model pitches (Wambsganss and Niklaus, 2022).

Some research has used argumentative structure to assess essay quality. In consecutive works, Persing et al. (2010) and Persing and Ng (2013, 2014, 2015) graded different argumentation-related quality aspects for the well-known essay corpus ICLE

<sup>1</sup>Our corpus and experiment code can be found under: <https://github.com/webis-de/NAACL-24>.

(Granger et al., 2009), namely organization, thesis clarity, prompt adherence, and argument strength. In contrast, Horbach et al. (2017) targeted different quality aspects of argumentative writing at once. Some works further investigated the interaction between argumentative structure and essay quality. Wachsmuth et al. (2016) found that multiple argumentation-related essay scoring tasks benefit from argument mining, underlining the impact of argumentative structure on essay quality. The analyses by Beigman Klebanov et al. (2016) and Nguyen and Litman (2018) suggest that this finding also holds for predicting holistic essay scores, while Persing et al. (2010) observed similar for the quality aspect organization. However, these studies relied on automatically assigned quality scores only, due to the lack of ground-truth annotations.

In a related line of research, approaches have been proposed to suggest revisions for argumentative essays (Afrin and Litman, 2018), to assess the need for and the quality of revisions (Skitalinskaya and Wachsmuth, 2023; Liu et al., 2023), as well as to perform argument revisions computationally (Skitalinskaya et al., 2023). Other works towards writing support for argumentative essays presented a prototypical system that gives simple feedback in terms of missed criteria (Stab, 2017), design principles for an adaptive learning tool (Wambsganss and Rietsche, 2019), visual feedback to the learner to prompt them to repair broken argument structures (Wambsganss et al., 2022a), or point to enthymematic gaps in arguments and make suggestions on how to fill these gaps (Stahl et al., 2023). Most recently, Britner et al. (2023) proposed a tool that not only analyzes issues with argument quality but also generates an explanation for its prediction. Our corpus supports these steps towards support systems for argumentative writing by providing detailed ground-truth annotations for both the argumentative structure and the quality of essays.

However, all the works mentioned deal with texts written by university students, while our work targets argumentative essays written by school students, fifth-graders and ninth-graders specifically. To the best of our knowledge, the only other published corpora with school student essays is not openly available (Correnti et al., 2013) and has been analyzed for essay-level quality aspects only, such as the integration of evidence (Rahimi et al., 2014) and the essay’s organization (Rahimi et al., 2015). We recently came across another school

student essay corpus in English with annotations for argumentative structure and quality, which has yet to be published.<sup>2</sup> We go beyond in this work by assessing the quality of school student essays in terms of five aspects derived from language education literature while incorporating their interaction with annotated argumentative structures. This may foster the development of effective methods for helping school students improve their argumentative writing skills.

### 3 School Student Essay Corpus

This section presents the source data and annotation of our corpus for analyzing the argumentative structure and quality of school student essays.

#### 3.1 Source Data

As the basis, we systematically selected 1,320 German school student essays from the FD-LEX corpus (Becker-Mrotzek and Grabowski, 2018).<sup>3</sup> The authors instructed students to each write three argumentative essays on topics pertinent to school students: (a) a letter to school funding organization on possible use of funding, (b) a statement on how to deal with the misbehavior of a fellow student, (c) a statement on who is guilty in a bike accident.

We seek to enable analyses of differences across different groups of school student essays on the corpus. Therefore, we pseudo-randomly chose 440 school students equally distributed across genders (only *male* and *female* exist in the corpus) and age groups (*fifth-graders* and *ninth-graders*). Subsequently, we included all three essays written by each selected school student from the source data.

#### 3.2 Annotation Scheme

Our annotation scheme goes beyond existing corpora for argument mining, covering the macro and micro structure of argumentative essays on four levels in total. In addition, we evaluate the quality of the essays overall and in terms of four quality aspects. Figure 2 overviews our annotation scheme.

**Argumentative Structure** On the broadest level of granularity for argumentative structure, we annotate *discourse functions* (Persing et al., 2010):

<sup>2</sup>The unpublished pre-print is available at <https://zenodo.org/records/8221504>.

<sup>3</sup>Essays and metadata (grade, school form, age, gender, language background, and age group) are available here: <https://fd-lex.uni-koeln.de>

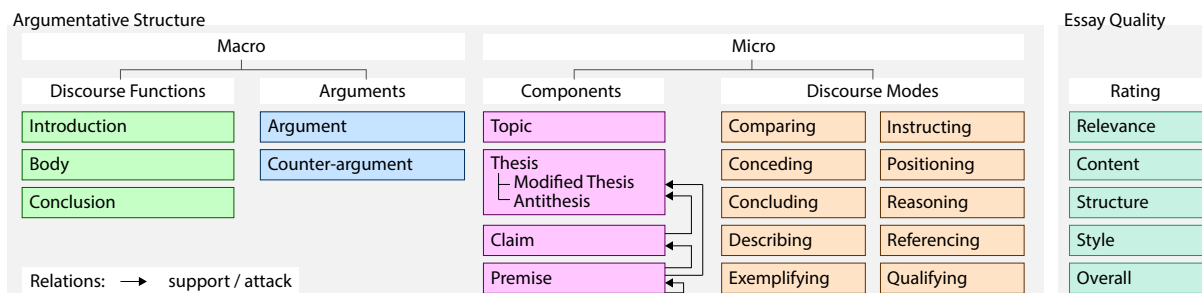


Figure 2: Proposed annotation scheme for argumentative school student essays: Four levels of argumentative macro and micro structure (discourse functions, arguments, components, discourse modes) and five essay quality aspects.

- *Introduction*. Initiates an essay by presenting the topic and possibly the context of an essay. This section is usually non-argumentative and placed at the beginning of an essay.
- *Body*. Core of the essay, containing the majority of argumentative components.
- *Conclusion*. Summary of main points, often with a final evaluation of the topic. This section is typically found at the end of the essay.

Next, we annotate *arguments* that comprise one point in an argumentative text, following [Walton et al. \(2008\)](#). We differentiate them by stance towards the main standpoint (thesis) of an essay:

- *Argument*. Ideally a claim (conclusion) and premises (reasons) supporting the claim.
- *Counter-argument*. An argument that attacks the thesis of an essay.

For analyzing the micro structure, we annotate argumentative and non-argumentative *components*. As [Stab and Gurevych \(2017a\)](#), we also mark support and attack relations between them (see Figure 2):

- *Topic*. Non-argumentative component that describes the subject or purpose of the essay.
- *Thesis*. Main standpoint of the whole argumentative text towards the topic. Repetitions of the thesis are also annotated as such.
- *Antithesis*. Thesis contrary to the actual thesis.
- *Modified Thesis*. Modified version of the actual thesis (e.g., more detailed or restricted).
- *Claim*. Statement that conveys a stance towards the topic.
- *Premise*. Reason that is given to support or attack a claim or another premise.<sup>4</sup>

<sup>4</sup>The components in our annotation scheme are similar to those of [Stab and Gurevych \(2017a\)](#). We added *antithesis* and *modified thesis* to reflect the changes in position in the essays.

On the finest level of granularity, we annotate *discourse modes* ([Smith, 2003](#)) specific to school student essays. They are derived from language education literature, where they are used for developing and analyzing argumentative writing skills ([Gätje et al., 2012](#); [Rezat, 2018](#); [Feilke and Rezat, 2021](#)):

- *Comparing*. Contrasting supporting and attacking points to a statement.
- *Conceding*. Addressing a counter-consideration and refuting it to support the own stance.
- *Concluding*. Drawing logical inferences using consecutive or final clauses (*so that, if... then*).
- *Describing*. Providing additional information, such as facts, statistics, and background data.
- *Exemplifying*. Providing examples or reporting on experiences.
- *Instructing*. Providing explicit instructions that recommend a specific course of action.
- *Positioning*. Expressing the own standpoint.
- *Reasoning*. Providing causal links to support a claim/thesis using markers (*because, then*).
- *Referencing*. Mentioning statements made by others, for example, by authorities.
- *Qualifying*. Presenting a variation of the all-or-nothing standpoints.

**Essay Quality** As [Persing et al. \(2010\)](#), we score essay quality on a 7-point scale from 1 (unsuccessful), 2 (rather unsuccessful), 3 (rather successful) to 4 (completely successful), with half points in between. We adapted the quality aspects of [Krusse et al. \(2012\)](#) for assessing school student essays in general to argumentative essays as follows:

- *Relevance*. The essay fits the prompt.
- *Content*. The selection of content helps to reach the essay's goal.

- *Structure*. The selected points are coherent and well-connected.
- *Style*. The use of language is adequate.
- *Overall*. The overall impression of the rater.

### 3.3 Annotation Process

We underwent the following process for both argumentative structure and essay quality:

To test and refine our annotation guidelines, we conducted pilot studies in which all annotators worked on the same 30 texts. We then discussed their understanding of the guidelines and annotation differences. We integrated their feedback into the guidelines to then test the reliability of the annotations in an inter-annotator agreement (IAA) study, where all annotators independently worked on the same 120 texts. Finally, each annotator annotated a set of 1,200 essays in the main annotation study.

As annotators, we employed experts in German language education from our lab and started the pilot and IAA studies on argumentative structure with three annotators. For the main part, only the two annotators with the most reliable annotations proceeded. The same annotators then annotated the essay quality, since they had already been trained in argumentative texts and our general procedure. However, we acknowledge that the annotators may have been predisposed to view essays in a certain way after the first annotation.

To assemble the final corpus, we combined the 1200 essays from the main study with the 120 essays from the IAA study after solving annotation conflicts. For conflicts between the three structure annotations per IAA essay, we kept the annotations that had the highest agreement across all levels with the other two. For conflicts between essay quality scores, we used their mean as the final score.<sup>5</sup>

### 3.4 Inter-Annotator Agreement

For the components, we follow [Stab and Gurevych \(2017a\)](#) in that we evaluate the agreement per essay at the token level, so the token labels are the unit of analysis. Thereby, overlaps of annotations are taken into account. For relations, we determined the component-level spans that at least two annotators agreed on with a relative overlap  $\geq 75\%$ . For all pairs of these, we then compared the relation labels (no relation, support, or attack) between the

<sup>5</sup>We rounded down to the next valid quality score for low values ( $< 2.5$ ) and rounded up for high values ( $> 2.5$ ) to prevent an upward distortion of the distribution.

Argumentative Structure		Essay Quality	
	$\alpha$		$\alpha$
Discourse Functions	0.89	Relevance	0.77
Arguments	0.86	Content	0.95
Components	0.81	Structure	0.84
Discourse Modes	0.74	Style	0.92
Relations	0.58	Overall	0.95

Table 1: Krippendorff’s  $\alpha$  agreement in the IAA study between three annotators for argumentative structure and two annotators for essay quality. The high values stress the high reliability of our annotations.

annotators. The mean Krippendorff’s  $\alpha$  scores over the 120 IAA essays are reported in Table 1. For essay quality, we computed the  $\alpha$ -value per quality aspect with essays as the unit of analysis.

The agreement is high for argumentative structure spans with values between 0.74 and 0.89. The agreement for relations is lower but reasonable, given that disagreement from the components annotations is propagated to the relations. The agreement for essay quality is also high, too, ranging from 0.77 to 0.95. Overall, we conclude that the annotations can mostly be seen as very reliable. Content and style quality annotations are very consistent between annotators, while assessing the relevance and structure seems slightly more subjective.

### 3.5 Corpus Statistics

Table 2 gives insights into the label distribution for argumentation structure. *Body* occurs most frequently among the discourse functions (1,335 times). With 56.75 tokens on average, bodies are also notably longer than *introductions* and *conclusions*. On the argument level, we note that *counter-arguments* are rather sparse in student essays. Among the components, *claims* are most frequent in total and per essay, followed by *theses*. Furthermore, we notice that *modified theses* are usually longer than theses, which matches our expectation that students add more details or restrictions to the thesis there. The most used discourse modes are *positioning*, *describing*, and *reasoning*, while *referencing*, *comparing*, and *exemplifying* occur rarely. Notable are also the differences in span length, e.g. *positioning* spans have on average only about half as many tokens as *comparing* spans.

Table 3 gives the frequency of annotated relations in our corpus. Most relations outgoing from *claims* are directed towards *theses* (92.6%), while most relations outgoing from *premises* are directed towards *claims* (85.8%). Overall, 96.2% of the re-

Label	# Spans	# Tokens	Tokens/Span	Spans/Essay
Introduction	114	2 329	20.43	0.09
Body	<b>1 335</b>	<b>75 766</b>	<b>56.75</b>	<b>1.01</b>
Conclusion	191	2 938	15.38	0.14
Argument	<b>2 692</b>	<b>51 560</b>	<b>19.15</b>	<b>2.04</b>
Counter-arg.	34	514	15.12	0.03
Topic	101	1 656	16.40	0.08
Thesis	1 687	19 581	11.61	1.28
Modified T.	267	4 490	<b>16.82</b>	0.20
Antithesis	14	174	12.43	0.01
Claim	<b>3 137</b>	<b>39 096</b>	12.46	<b>2.38</b>
Premise	1 020	12 533	12.29	0.77
Comparing	20	431	<b>21.55</b>	0.02
Conceding	142	2 874	20.24	0.11
Concluding	868	11 654	13.43	0.66
Describing	1 692	<b>22 258</b>	13.15	1.28
Exemplifying	63	926	14.70	0.05
Instructing	176	2 174	12.35	0.13
Positioning	<b>1 758</b>	19 178	10.91	<b>1.33</b>
Reasoning	1 553	17 204	11.08	1.18
Referencing	16	197	12.31	0.01
Qualifying	147	2 344	15.95	0.11

Table 2: Argumentative structure annotations in the corpus: Total number of spans and tokens per label, average span length in number of tokens (*Tokens/Span*) and average number of spans per essay (*Spans/Essay*). The highest value per column and level is marked bold.

From <i>Claim</i> to	#	%	From <i>Premise</i> to	#	%
Thesis	<b>2844</b>	<b>92.6</b>	Thesis	6	0.6
Modified Thesis	218	7.1	Claim	<b>872</b>	<b>85.8</b>
Antithesis	9	0.3	Premise	138	13.6

Table 3: Absolute and relative frequency of annotated relations outgoing from *claim* (left) or *premise* (right).

lations were labeled as *support* and 3.8% as *attack*.

The distribution of quality scores is shown in Table 4. We can see that *relevance*, *structure*, and *style* have a similar score distribution, while the distribution of *content* scores is shifted towards the higher scores, with the highest mean (2.80). Overall quality has the lowest mean score (2.20) and was most often scored with the lowest score of 1.0. These results suggest that overall quality is not just the average of the other annotated quality aspects but that it emerges from the annotators’ perception and possibly other aspects.

## 4 Analysis

This section reports on our corpus analysis of the interaction between the argumentative structure on micro vs. macro level and component vs. discourse mode level, and the different essay quality aspects.

Quality Aspect	1.0	1.5	2.0	2.5	3.0	3.5	4.0	Mean
Relevance	64	131	<b>548</b>	358	190	23	6	2.22
Content	62	33	123	312	<b>510</b>	<b>173</b>	<b>107</b>	<b>2.80</b>
Structure	59	91	423	414	292	24	17	2.35
Style	75	<b>159</b>	396	<b>436</b>	233	12	9	2.25
Overall	<b>90</b>	142	478	390	195	23	2	2.20

Table 4: Distribution and mean of scores per quality aspect. The highest value per column is marked bold.

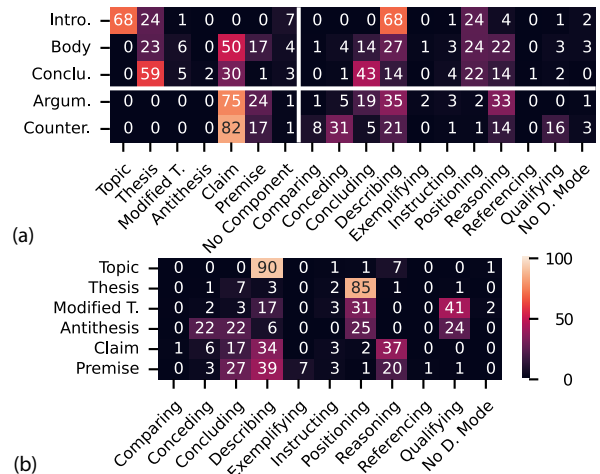


Figure 3: Cooccurrence matrices: Relative token-level overlap of (a) macro and micro structure and (b) component and discourse mode labels in percent. For example, 68% of all tokens labeled as *Introduction* on the macro level are also labeled as *Topic* on the micro level.

### 4.1 Macro vs. Micro structure

Figure 3(a) shows the overlap between macro structure (discourse functions and arguments) and micro structure (components and discourse mode) labels. The *introduction* mainly includes the *topic* (68%), while every second token in the *body* is a *claim* on average. The *thesis* cooccurs with all three discourse functions. We see that the proportion of *claims* and *premises* differs for *arguments* and *counter-arguments*. Counter-arguments contain more claim tokens than arguments, while fewer counter-argument tokens are part of a premise.

The usage of discourse modes differs between the macro-structure levels, too. In the introduction, students mostly *describe* (68%) and *position* (24%). The discourse modes are more diverse for the body, also including a notable portion of *reasoning*. As expected, in the *conclusion*, students focus more on *concluding*. While describing and reasoning are prevalent in arguments and counter-arguments, a notable portion of argument tokens is used for concluding. At the same time, more *conceding* and

	Relevance	Content	Structure	Style	Overall
Relevance		.53	.61	.47	<b>.75</b>
Content	.53		.48	.41	.60
Structure	.61	.48		.51	.71
Style	.47	.41	.51		.61
Overall	<b>.75</b>	<b>.60</b>	<b>.71</b>	<b>.61</b>	

Table 5: Kendall’s  $\tau$  correlation between the quality aspects. The highest value per column is marked bold.

*qualifying* tokens occur in counter-arguments. This is expected, since especially in counter-arguments other points of view should be varied or refuted.

## 4.2 Components vs. Discourse modes

The cooccurrences between components and actions can be seen in Figure 3(b). While the *topic* is mostly *described* (90%) and the *thesis* consists primarily of *positioning* (85%), the remaining components include more diverse discourse modes. In contrast to *theses*, *modified theses* also feature *describing* and *qualifying* tokens, while *antitheses* additionally cover *conceding* and *concluding*. *Claims* and *premises* mainly cooccur with *describing*, *reasoning*, and *concluding*. However, the proportions differ slightly. The cooccurrence matrix between all structure labels can be found in Appendix A.

## 4.3 Essay Quality

To further assess the interaction between the quality aspects, Table 5 shows all pairwise Kendall’s  $\tau$  correlations. All aspects correlate most with *overall quality*, most strongly *relevance* (.75). The correlation between *content* and *style* is lowest (.41), which underlines their distinctive nature.

## 5 Experiments

This section presents baseline approaches to the two main tasks our corpus enables: Predicting the argumentative structure (argument mining) and the essay quality (essay scoring). Additionally, we investigate whether information about the argumentative structure helps to predict the essay quality.

### 5.1 Argument Mining

We treat argument mining as a token classification task: Given a school student essay and a structure level, predict the label of each token on that structure level. The IOB2 format is used for the labels to separate adjacent spans of the same type. We performed 5-fold cross-validation for each structure level. For each folding, we used four folds

Approach	D. Func.		Argum.		Compon.		D. Mode	
	Acc.	F <sub>1</sub>	Acc.	F <sub>1</sub>	Acc.	F <sub>1</sub>	Acc.	F <sub>1</sub>
Random	.14	.00	.20	.00	.08	.00	.05	.00
Majority	.86	.52	.56	.00	.41	.00	.24	.00
mDeBERTaV3	.92	.46	.86	.29	.66	.21	.63	.21
-adapter	<b>.95</b>	<b>.68</b>	<b>.92</b>	<b>.52</b>	<b>.76</b>	<b>.49</b>	<b>.73</b>	<b>.46</b>
Human	.98	.94	.96	.85	.93	.89	.89	.84

Table 6: Argument mining results: Macro F<sub>1</sub>-score and accuracy of each approach in 5-fold cross-validation on all four argumentative structure dimensions. The best value per column is marked bold.

(80%) for training and divided the fifth fold in half: one half (10%) for selecting the best-performing checkpoint in terms of macro-averaged F<sub>1</sub>-score, and the remaining half (10%) for testing.

**Models** We used the multilingual model *mDeBERTaV3* (He et al., 2023) (microsoft/mdeberta-v3-base) from Huggingface (Wolf et al., 2020).<sup>6</sup> Besides, we tested the effect of training *adapters* (Houlsby et al., 2019), a set of task-specific parameters that are added to every transformer layer of mDeBERTaV3 and fine-tuned on the task while the model weights are fixed. To quantify the impact of learning, we compare against a random baseline that chooses a token label pseudo-randomly and a majority baseline that always predicts the majority token label from the training set. As upper bound, we report the human performance in terms of the average of each annotator in isolation on the 120 IAA texts annotated by all annotators.<sup>7</sup>

**Experimental Setup** We train mDeBERTaV3 for 30 epochs (1,980 steps) using the suggested hyperparameter values: a learning rate of  $3e - 5$ , batch size 16, and 500 warmup steps. For mDeBERTaV3-adapter, we follow Pfeiffer et al. (2020) who recommend to use a higher learning rate of  $1e - 4$  and train longer, here 50 epochs (3,300 steps).

**Results** Table 6 presents the token classification results for all levels of argumentative structure, averaged over all folds. Noteworthy, *mDeBERTaV3-adapter* outperforms training the whole model (*mDeBERTaV3*) in all cases. Given that the F<sub>1</sub>-scores improve more than the accuracy, the

<sup>6</sup>Exploratory experiments using instruction fine-tuned models such as Alpaca (Taori et al., 2023) did not lead to promising results for our token classification task.

<sup>7</sup>Note that the test set used for the model performance is a different subset of the dataset.

Approach	Relevance	Content	Structure	Style	Overall
Random	-0.013 ±0.084	-0.011 ±0.071	-0.014 ±0.073	0.017 ±0.084	-0.004 ±0.083
Majority	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000
mDeBERTaV3	0.530 ±0.069	0.295 ±0.109	0.513 ±0.044	0.492 ±0.059	0.616 ±0.040
-adapter	0.564 ±0.018	0.431 ±0.098	<b>0.575</b> ±0.038	0.579 ±0.077	0.648 ±0.054
-fusion-w/-discourse-functions	0.599 ±0.043	0.381 ±0.134	0.559 ±0.036	0.569 ±0.069	0.668 ±0.049
-fusion-w/-arguments	0.593 ±0.030	0.448 ±0.105	<b>0.575</b> ±0.019	0.581 ±0.054	0.668 ±0.036
-fusion-w/-components	† <b>0.600</b> ±0.025	0.437 ±0.137	0.543 ±0.044	0.585 ±0.053	0.663 ±0.046
-fusion-w/-discourse-modes	0.544 ±0.028	0.420 ±0.118	0.535 ±0.041	0.583 ±0.064	0.645 ±0.023
-fusion-w/-all	0.574 ±0.039	<b>0.454</b> ±0.142	0.546 ±0.013	<b>0.617</b> ±0.057	† <b>0.686</b> ±0.031
Human	0.636 ±0.055	0.632 ±0.003	0.734 ±0.007	0.766 ±0.005	0.746 ±0.003

Table 7: Essay scoring results: QWK of each approach in 5-fold cross-validation on all five quality dimensions. The best value per column is marked bold. We mark significant gains over *mDeBERTaV3-adapter* at  $p < .05$  with †.

adapters seem less prone to overfitting to the majority label. This learning success suggests the possibility of predicting all argumentative structure levels on our corpus. However, further improvements using more advanced approaches are expected.

## 5.2 Essay Scoring

We treat predicting the essay quality as a text classification task: Given a school student essay and a quality aspect, predict the corresponding quality score. As before, we performed 5-fold cross-validation for each quality aspect using the same folds. We selected the best-performing checkpoint on the validation set using quadratic weighted kappa (QWK), the most widely adopted metric for automatic essay scoring (Ke and Ng, 2019).

**Models** We adopted the previous approaches by changing the head to a text classification head. To analyze the interaction between argumentative structure and essay quality, we employed *Adapter-Fusion* (Pfeiffer et al., 2021), a multi-task learning framework that can be used to investigate relations between different dimensions by learning how to combine model weights with one or more adapters. We used the mDeBERTaV3-adapters trained on argumentative structure from the previous experiment. As the final adapter, we chose the one trained on the folding that performed most representative for all folds ( $F_1$ -score closest to the reported averaged  $F_1$ -score across folds). To measure the impact of each level of argumentative structure on the scoring performance, we used each adapter individually and a combination of all of them.

**Experimental Setup** The experimental setup for mDeBERTaV3 and mDeBERTaV3-adapter was adopted from before. For training the Adapter-

Fusion, we followed Pfeiffer et al. (2021) to use a learning rate of  $5e - 5$  and trained shorter than the adapters, in our case for 20 epochs (1,320 steps).

**Results** Table 7 shows the scoring results. All models outperform the lower-bound baselines (*random* and *majority*), suggesting that the quality scoring can be learned from our corpus. Furthermore, fusing all adapters trained on argumentative structure (*mDeBERTaV3-fusion-w/-all*) performs best for three out of five quality aspects, significantly beating mDeBERTaV3-adapter in predicting overall quality.<sup>8</sup> This underlines the need for all four levels of argumentative structure together in order to improve scoring overall quality (0.686 vs. 0.648). In addition, using only the adapter trained on the component level (*mDeBERTaV3-fusion-w/-components*) helps to significantly improve over mDeBERTaV3-adapter in predicting relevance (0.600 vs. 0.564), indicating an interaction between the structure on component level and this quality aspect. QWK scores greater or equal to 0.6 suggest substantial agreement between the predicted and ground-truth quality scoring of essays.

**AdapterFusion Activations** AdapterFusion extracts information from adapters only if they benefit the target task. Similar to Falk and Lapesa (2023), we visualize the average activations of our model *mDeBERTaV3-fusion-w/-all* over the layers in Figure 4 to investigate the influence of each level of argumentative structure on the quality scoring. All adapters are activated fairly evenly for all quality aspects, with slight deviations. This aligns with our previous results and underlines that all annotated

<sup>8</sup>We use Wilcoxon signed-rank tests at  $p < .05$  for testing the significance.






Figure 4: AdapterFusion activation on average over the layers for each *mDeBERTaV3-fusion-w/-all* model per quality aspect. We average the activation for each fused adapter (discourse functions, arguments, components, discourse modes) over all instances in the most representative test set folding.

structure levels are helpful for quality scoring. The activations per layer can be found in Appendix B.

## 6 Conclusion

Argumentative writing support of school students presupposes that the quality of their arguments can be assessed. Until now, no argument mining corpus with school student essays has been published, let alone any essay corpus with both argument and quality annotations. With this work, we fill both research gaps with a new corpus of 1,320 German school student essays, annotated by experts for argumentative structure and essay quality.

Our corpus analysis has provided various insights into the correlation between the different levels of argumentative structure and essay quality. In our experiments with fine-tuned transformers and adapters for mining argumentative structure and scoring essay quality we have demonstrated that combining information on all four argumentative structure levels helps the prediction of essay quality. This shows the usefulness of our corpus for research on quality-oriented argumentative writing support, which we seek to enable with this paper.

We point out that our corpus contains various information yet to be explored, such as argumentative relations and school student metadata. It thus lays the ground for further analyses—like identifying unwarranted claims and studying differences across age groups and genders.

## 7 Limitations

Aside from the still-improvable performance of the presented baseline models for argument mining and essay scoring, we see two notable limitations of our work: the restriction to German texts, and the pending utilization of the corpus for quality-oriented argumentative writing support.

First, we point out the specific language background of our work. The essays were written by German school students, and the annotations were developed in close communication with German

experts from the field of language education, while the discourse modes and essay quality aspects are, to a considerable extent, derived from work on German texts. This means that our findings may not perfectly align with argumentative writing in other countries or languages with different expectations for argumentative essays.

Second, while our analyses suggest that our corpus helps to enable quality-oriented argumentative writing support, the perceived usefulness of such a tool is still to be evaluated. We expect and encourage future work to utilize our corpus for such writing support tools, for example, by further analyzing which exact argumentative structures influence the essay quality and to what extent. Interpretable essay quality scoring based on the structure might generate helpful insights that can be used as writing feedback by school students.

## 8 Ethical Considerations

We see no apparent risk of the corpus or the methods presented in this paper being misused for ethically doubtful purposes. The authors of the FD-LEX corpus (Becker-Mrotzek and Grabowski, 2018) have already pseudonymized the author of each essay. Therefore, it is not possible to identify the individual school student from the provided data. However, we want to point out that one might find differences in the essays across gender or age groups that do not reflect reality but are rather due to an unintentional bias in the data selection.

## Acknowledgments

We would like to thank the participants of our study and the anonymous reviewers for the valuable feedback and their time. This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the project ArgSchool, project number 453073654.

## References

- Tazin Afrin and Diane Litman. 2018. [Annotation and classification of sentence-level revision improvement](#). In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 240–246, New Orleans, Louisiana. Association for Computational Linguistics.
- Aniket Ajit Tambe and Manasi Kulkarni. 2022. [Automated essay scoring system with grammar score analysis](#). In *2022 Smart Technologies, Communication and Robotics (STCR)*, pages 1–7.
- Patricia A. Alexander, Jannah Fusenig, Eric C. Schoute, Anisha Singh, Yuting Sun, and Julianne E. van Meerten. 2023. [Confronting the challenges of undergraduates’ argumentation writing in a “learning how to learn” course](#). *Written Communication*, 40(2):482–517.
- Michael Becker-Mrotzek and Joachim Grabowski. 2018. [Textkorpus Scriptoria](#). In Michael Becker-Mrotzek and Joachim Grabowski, editors, *FD-LEX (Forschungsdatenbank Lernertexte)*. Mercator-Institut für Sprachförderung und Deutsch als Zweitsprache, Köln. Available at: <https://fd-lex.uni-koeln.de>.
- Michael Becker-Mrotzek, Frank Schneider, and Klaus Tetling. 2010. [Argumentierendes Schreiben - lehren und lernen. Vorschläge für einen systematischen Kompetenzaufbau in den Stufen 5 bis 8](#).
- Beata Beigman Klebanov, Christian Stab, Jill Burstein, Yi Song, Binod Gyawali, and Iryna Gurevych. 2016. [Argumentation: Content, structure, and relationship with essay quality](#). In *Proceedings of the Third Workshop on Argument Mining (ArgMining2016)*, pages 70–75, Berlin, Germany. Association for Computational Linguistics.
- Sebastian Britner, Lorik Dumani, and Ralf Schenkel. 2023. [Aquadplane: The argument quality explainer app](#). In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 5015–5020, New York, NY, USA. Association for Computing Machinery.
- Richard Correnti, Lindsay Clare Matsumura, Laura Hamilton, and Elaine Wang. 2013. [Assessing students’ skills at writing analytically in response to texts](#). *The Elementary School Journal*, 114(2):142–177.
- Scott Crossley, Perpetual Baffour, Tian Yu, Alex Franklin, Meg Benner, and Ulrich Boser. 2023a. [A large-scale corpus for assessing written argumentation: PERSUADE 2.0](#).
- Scott Crossley, Yu Tian, Perpetual Baffour, Alex Franklin, Youngmeen Kim, Wesley Morris, Meg Benner, Aigner Picou, and Ulrich Boser. 2023b. [The english language learner insight, proficiency and skills evaluation \(ellipse\) corpus](#). *International Journal of Learner Corpus Research*, 9(2):248–269.
- Thi Hanh Dang, Thanh Hai Chau, and To Quyen Tra. 2020. [A study on the difficulties in writing argumentative essays of english-majored sophomores at tay do university, vietnam](#). *European Journal of English Language Teaching*, 6(1).
- Neele Falk and Gabriella Lapesa. 2023. [Bridging argument quality and deliberative quality annotations with adapters](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2469–2488, Dubrovnik, Croatia. Association for Computational Linguistics.
- Youmna Farag, Helen Yannakoudakis, and Ted Briscoe. 2018. [Neural automated essay scoring and coherence modeling for adversarially crafted input](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 263–271, New Orleans, Louisiana. Association for Computational Linguistics.
- Helmuth Feilke. 2017. [Schreib- und Textprozeduren](#). In Jürgen Baurmann, Clemens Kammler, and Astrid Müller, editors, *Handbuch Deutschunterricht. Theorie und Praxis des Lehrens und Lernens*, 1 edition, pages 51–57. Reihe Praxis Deutsch.
- Helmuth Feilke and Sara Rezat. 2021. [Textprozeduren und der Erwerb literaler Kompetenz](#). In Nikolaus Koch and Barbara Kozikowski, editors, *Sprach(en)erwerb*, pages 69–79. Der Deutschunterricht.
- Ralph P. Ferretti, Scott Andrews-Weckerly, and William E. Lewis. 2007. [Improving the argumentative writing of students with learning disabilities: Descriptive and normative considerations](#). *Reading & Writing Quarterly*, 23(3):267–285.
- Sylviane Granger, Estelle Dagneaux, Fanny Meunier, and Magali Paquot. 2009. *The International Corpus of Learner English*. Presses universitaires de Louvain, Louvain-la-Neuve.
- Timon Gurcke, Milad Alshomary, and Henning Wachsmuth. 2021. [Assessing the sufficiency of arguments through conclusion generation](#). In *Proceedings of the 8th Workshop on Argument Mining*, pages 67–77, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Olaf Gätje, Sara Rezat, and Torsten Steinhoff. 2012. [Positionierung. Zur Entwicklung des Gebrauchs modalisierender Prozeduren in argumentativen Texten von Schülern und Studenten](#). In Helmuth Feilke and Katrin Lehnen, editors, *Textroutinen. Theorie, Erwerb und didaktisch-mediale Modellierung*, pages 125–153. Lang, Frankfurt/Main.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTav3: Improving deBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations*.

- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Maleerat Ka-kan-dee and Sarjit Kaur. 2014. [Argumentative writing difficulties of thai english major students](#). In *The 2014 WEI International Academic Conference Proceedings*, pages 193–207.
- Zixuan Ke and Vincent Ng. 2019. [Automated essay scoring: A survey of the state of the art](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6300–6308. International Joint Conferences on Artificial Intelligence Organization.
- Ronald T. Kellogg, Alison P. Whiteford, and Thomas Quinlan. 2010. [Does automated feedback help students learn to write?](#) *Journal of Educational Computing Research*, 42(2):173–196.
- Norbert Kruse, Anke Reichardt, Maik Herrmann, Friederike Heinzl, and Frank Lipowsky. 2012. [Zur qualität von kindertexten. entwicklung eines bewertungsinstruments in der grundschule](#). *Didaktik Deutsch : Halbjahresschrift für die Didaktik der deutschen Sprache und Literatur*, 17(32):87–110.
- Xia Li, Minping Chen, Jianyun Nie, Zhenxing Liu, Ziheng Feng, and Yingdan Cai. 2018. [Coherence-based automated essay scoring using self-attention](#). In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, pages 386–397, Cham. Springer International Publishing.
- Zhexiong Liu, Diane Litman, Elaine Wang, Lindsay Matsumura, and Richard Correnti. 2023. [Predicting the quality of revisions in argumentative writing](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 275–287, Toronto, Canada. Association for Computational Linguistics.
- Huy Nguyen and Diane Litman. 2018. [Argument mining for improving the automated scoring of persuasive essays](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- John Peloghitis. 2017. [Difficulties and strategies in argumentative writing: A qualitative analysis](#). *Transformation in language education. JALT*.
- Isaac Persing, Alan Davis, and Vincent Ng. 2010. [Modeling organization in student essays](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 229–239, Cambridge, MA. Association for Computational Linguistics.
- Isaac Persing and Vincent Ng. 2013. [Modeling thesis clarity in student essays](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 260–269, Sofia, Bulgaria. Association for Computational Linguistics.
- Isaac Persing and Vincent Ng. 2014. [Modeling prompt adherence in student essays](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1534–1543, Baltimore, Maryland. Association for Computational Linguistics.
- Isaac Persing and Vincent Ng. 2015. [Modeling argument strength in student essays](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 543–552, Beijing, China. Association for Computational Linguistics.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Zahra Rahimi, Diane Litman, Elaine Wang, and Richard Correnti. 2015. [Incorporating coherence of topics as a criterion in automatic response-to-text assessment of the organization of writing](#). In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 20–30, Denver, Colorado. Association for Computational Linguistics.
- Zahra Rahimi, Diane J. Litman, Richard Correnti, Lindsay Clare Matsumura, Elaine Wang, and Zahid Kisa. 2014. [Automatic scoring of an analytical response-to-text assessment](#). In *Intelligent Tutoring Systems*, pages 601–610, Cham. Springer International Publishing.
- Sara Rezat. 2011. [Schriftliches Argumentieren. Zur Ontogenese konzessiver Argumentationskompetenz](#).

- Didaktik Deutsch: Halbjahresschrift für die Didaktik der deutschen Sprache und Literatur*, 16(31):50–67.
- Sara Rezat. 2018. Argumentative Textprozeduren als Instrumente zur Anbahnung wissenschaftlicher Textkompetenz. In Sabine Schmölzer-Eibinger, Bora Bushati, Christopher Ebner, and Lisa Niederdorfer, editors, *Wissenschaftliches Schreiben lehren und lernen. Diagnose und Förderung wissenschaftlicher Textkompetenz in Schule und Universität*, pages 125–146. Waxmann, Münster.
- Juliane Schröter. 2021. *Linguistische Argumentationsanalyse*. Kurze Einführungen in die germanistische Linguistik. Universitätsverlag Winter, Heidelberg.
- Gabriella Skitalinskaya, Maximilian Spliethöver, and Henning Wachsmuth. 2023. **Claim optimization in computational argumentation**. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 134–152, Prague, Czechia. Association for Computational Linguistics.
- Gabriella Skitalinskaya and Henning Wachsmuth. 2023. **To revise or not to revise: Learning to detect improvable claims for argumentative writing support**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15799–15816, Toronto, Canada. Association for Computational Linguistics.
- Carlota S. Smith. 2003. *Modes of Discourse: The Local Structure of Texts*. Cambridge Studies in Linguistics. Cambridge University Press.
- Christian Stab. 2017. *Argumentative Writing Support by means of Natural Language Processing*. Ph.D. thesis, Technische Universität Darmstadt, Darmstadt.
- Christian Stab and Iryna Gurevych. 2017a. **Parsing argumentation structures in persuasive essays**. *Computational Linguistics*, 43(3):619–659.
- 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. Association for Computational Linguistics.
- Maja Stahl, Nick Düsterhus, Mei-Hua Chen, and Henning Wachsmuth. 2023. **Mind the gap: Automated corpus creation for enthymeme detection and reconstruction in learner arguments**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4703–4717, Singapore. Association for Computational Linguistics.
- Manfred Stede and Jodi Schneider. 2019. *Argumentation Mining*. Springer International Publishing.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Michael A. R. Townsend, Lynley Hicks, Jacquilyn D. M. Thompson, Keri M. Wilton, Bryan F. Tuck, and Dennis W. Moore. 1993. **Effects of introductions and conclusions in assessment of student essays**. *Journal of Educational Psychology*, 85(4):670–678.
- Masaki Uto, Yikuan Xie, and Maomi Ueno. 2020. **Neural automated essay scoring incorporating hand-crafted features**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6077–6088, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Henning Wachsmuth, Khalid Al-Khatib, and Benno Stein. 2016. **Using argument mining to assess the argumentation quality of essays**. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1680–1691, Osaka, Japan. The COLING 2016 Organizing Committee.
- Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- Thiemo Wambsganss, Andrew Caines, and Paula Buttery. 2022a. **ALEN app: Argumentative writing support to foster English language learning**. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 134–140, Seattle, Washington. Association for Computational Linguistics.
- Thiemo Wambsganss, Andreas Janson, and Jan Marco Leimeister. 2022b. **Enhancing argumentative writing with automated feedback and social comparison nudging**. *Computers and Education*, 191:104644.
- Thiemo Wambsganss and Christina Niklaus. 2022. **Modeling persuasive discourse to adaptively support students' argumentative writing**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8748–8760, Dublin, Ireland. Association for Computational Linguistics.
- Thiemo Wambsganss, Christina Niklaus, Matthias Söllner, Siegfried Handschuh, and Jan Marco Leimeister. 2020. **A corpus for argumentative writing support in German**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 856–869, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Thiemo Wambsganss and Roman Rietsche. 2019. **Towards designing an adaptive argumentation learning tool**. In *International Conference on Interaction Sciences*.
- Cong Wang, Zhiwei Jiang, Yafeng Yin, Zifeng Cheng, Shiping Ge, and Qing Gu. 2023. **Aggregating multiple heuristic signals as supervision for unsupervised automated essay scoring**. In *Proceedings of the 61st Annual Meeting of the Association for Computational*

*Linguistics (Volume 1: Long Papers)*, pages 13999–14013, Toronto, Canada. Association for Computational Linguistics.

Florian Weber, Thimo Wambsganss, Seyed Parsa Neshaei, and Matthias Soellner. 2023. [Structured persuasive writing support in legal education: A model and tool for German legal case solutions](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2296–2313, Toronto, Canada. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Ruosong Yang, Jiannong Cao, Zhiyuan Wen, Youzheng Wu, and Xiaodong He. 2020. [Enhancing automated essay scoring performance via fine-tuning pre-trained language models with combination of regression and ranking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1560–1569, Online. Association for Computational Linguistics.

Wei Zhu. 2001. [Performing argumentative writing in english: Difficulties, processes, and strategies](#). *TESL Canada Journal*, 19(1):34–50.

## A Cooccurrence Matrix

The cooccurrence matrix between all argumentative structure levels (discourse functions, arguments, components, and discourse modes) is shown in Figure 5.

## B AdapterFusion Activations

Similar to Pfeiffer et al. (2021), we visualize the activations of our model *mDeBERTaV3-fusion-w/all* per layer in Figure 6 to further investigate the influence of each level of argumentative structure on the quality scoring. The first activation layers show for all five quality aspects that all structure adapters are activated quite diversely. In contrast, the later layers have a clear tendency towards activating only one or two adapters. Notable is the similar activation pattern between relevance and overall quality, which could come from their value correlation.

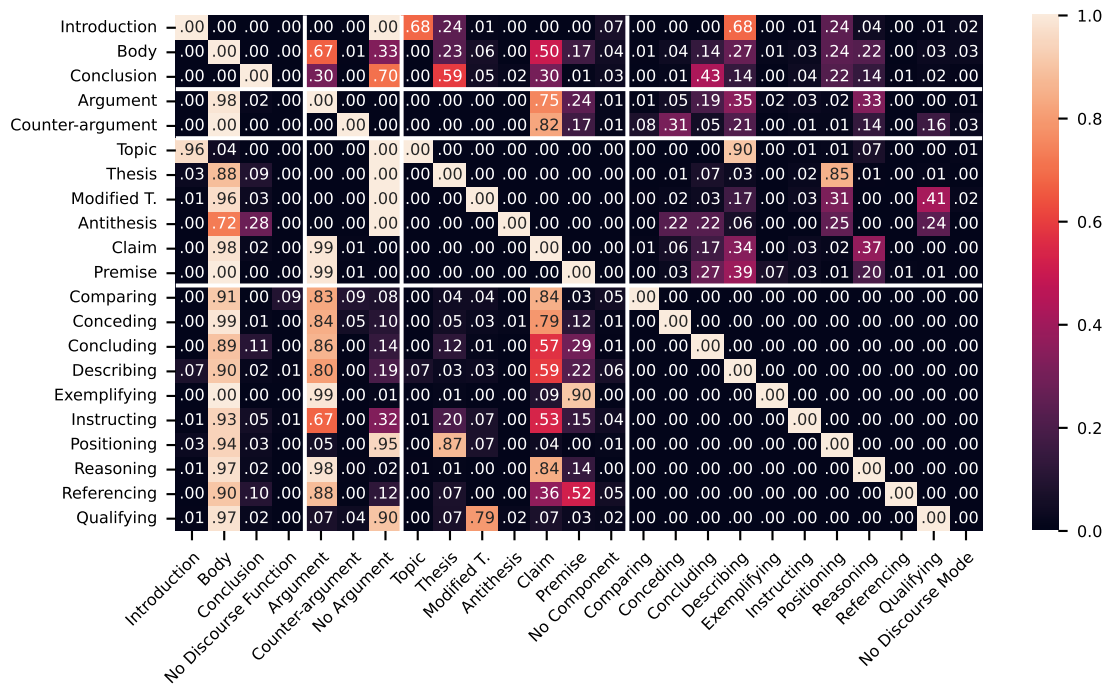


Figure 5: Relative token-level overlap of all argumentative structure labels, separated into the four levels of granularity. For example, 68% of all tokens labeled as *Introduction* are also labeled *Topic*.




Figure 6: AdapterFusion activation per layer (1–12) and on average over the layers (Avg) for each *mDeBERTaV3-fusion-w/-all* model per quality aspect. We average the activation for each fused adapter (for discourse functions, arguments, components, or discourse modes) over all instances in the test set of the most representative folding.