

Multilingual Argument Mining: Datasets and Analysis

Orith Toledo-Ronen, Matan Orbach, Yonatan Bilu, Artem Spector and Noam Slonim

IBM Research

{oritht, matano, yonatanb, artems, noams}@il.ibm.com

Abstract

The growing interest in argument mining and computational argumentation brings with it a plethora of Natural Language Understanding (NLU) tasks and corresponding datasets. However, as with many other NLU tasks, the dominant language is English, with resources in other languages being few and far between. In this work, we explore the potential of transfer learning using the multilingual BERT model to address argument mining tasks in non-English languages, based on English datasets and the use of machine translation. We show that such methods are well suited for classifying the stance of arguments and detecting evidence, but less so for assessing the quality of arguments, presumably because quality is harder to preserve under translation. In addition, focusing on the translate-train approach, we show how the choice of languages for translation, and the relations among them, affect the accuracy of the resultant model. Finally, to facilitate evaluation of transfer learning on argument mining tasks, we provide a human-generated dataset with more than 10k arguments in multiple languages, as well as machine translation of the English datasets.

1 Introduction

Argument mining has received much attention in recent years, with research mainly focused on English and, to some extent, German texts. Recent advancements in Natural Language Understanding suggest that in order to train appropriate models for argument mining tasks in other languages, we do not need to manually label text in these languages, but rather employ transfer learning from the English-based models (Eger et al., 2018a).

In this work we examine three argument mining tasks: (1) *stance classification*: given a topic and an argument that supports or contests the topic, determine the argument’s stance towards the topic;

(2) *evidence detection*: given a topic and a sentence, determine if the sentence is an evidence relevant to the topic; (3) *argument quality*: given a topic and a relevant argument, rate the argument so that higher-quality arguments are assigned a higher score.

To facilitate transfer learning from English datasets for these tasks, we employ Multilingual BERT (mBERT) released by Devlin et al. (2019), a pre-trained language model that supports 104 languages, and use it mainly in a translate-train approach. Namely, the English dataset is automatically translated into the desired language(s) using machine translation (MT); an augmented dataset composed of the original English text and all the translated copies is created; the mBERT model is fine-tuned on a subset of the dataset; and the resultant model is then used to solve the relevant downstream task in the desired language. Previous works have suggested that translating the original dataset to as large a number of languages as possible is beneficial (Liang et al., 2020). In this work, we show a more nuanced picture, where often selecting a subset of related languages is preferable.

In addition, we also examine the translate-test approach, in which one creates a classification model only with English data. At prediction time, the non-English text is automatically translated into English, and then analyzed by the model. This approach is less appealing since the initial translation step increases prediction run-time, and on our data also tends to perform worse.

We examine two text sources for performance evaluation on non-English texts. The first one is a “pseudo test” set – an automatic translation of an English evaluation set for a task. While such texts can be easily generated, it is not clear how well they represent “real” texts, authored by humans. Hence, we also examine human-authored texts, in several non-English languages, collected via crowdsourcing specifically for this work. Both

datasets are released as part of this work.¹

When translating the evaluation set, either automatically or by a human translator, one would like to assume that the initial label of the English text is maintained after translation. While this is often the case, we show that this assumption becomes more dubious as the argument mining task becomes more complex and subjective, as well as when the original labels are not clearly agreed upon.

In summary, the main contributions of this paper are: (1) a comparative analysis of the translate-train approach on three central argument mining tasks using different subsets of languages, showing that training on more data helps, but that, in some cases, training on related languages is sufficient; (2) multilingual benchmark datasets for the three tasks; (3) an analysis of the three tasks, showing how well labels are preserved across translation, and the impact that has on the success of the translate-train approach.

2 Related Work

Argument mining has been expanding from identifying argumentative passages to a variety of Natural Language Understanding (NLU) tasks (Stede and Schneider (2018); Lawrence and Reed (2020)). In this work, we explore three argumentation tasks in multilingual settings.

Stance detection (or classification) is often contrasted with sentiment analysis, in that the task is not simply to classify the sentiment of a text, but rather its stance w.r.t. some given target. Early work on this task includes Thomas et al. (2006); Lin et al. (2006), while in the context of argument mining, the task has probably been introduced by Sobhani et al. (2015). As with many other NLU tasks, earlier works developed classifiers based on various features (e.g., Bar-Haim et al. (2017)), while more modern approaches rely on deep learning. See Schiller et al. (2020) for a recent benchmarking report on such methods.

Research on stance detection in a multilingual setting is rather recent. Zotova et al. (2020) explore stance detection in Twitter for Catalan and Spanish; Lai et al. (2020) do this for political debates in social media in these two languages as well as French and Italian; Vamvas and Sennrich (2020) analyze the stance of comments in the context of

the Switzerland election in German, French and Italian. Stance detection is reminiscent of the Natural Language Inference (NLI) problem, where one is given two sentences, and the objective is to determine whether one entails the other, contradicts it or is neutral. This task had been researched extensively, with Conneau et al. (2018) providing a 15-language benchmark for the multilingual setting. Another earlier work on a related task is addressing the support/attack relation prediction of two arguments in Italian (Basile et al., 2016).

Evidence detection is the task of determining, given some text and a topic, whether the text can serve as evidence in the context of the topic (Rinott et al., 2015). We follow Ein-Dor et al. (2020), in defining *evidence* as *a single sentence that clearly supports or contests the topic, yet is not merely a belief or a claim. Rather, it provides an indication for whether a relevant belief or a claim is true*, and, since we use their datasets, we restrict our analysis to evidence of type Study and Expert. In a multilingual setting, a similar task, that of premise detection, was considered in Eger et al. (2018a) for German, French, Spanish, and Chinese; in Fishcheva and Kotelnikov (2019) for Russian; in Eger et al. (2018b) for French and German; and in Aker and Zhang (2017) for Chinese.

Argument quality prediction is the task of evaluating the quality of an argument, either on an objective scale - the input is an argument and the output is a quality score; or in a relative manner - the input is a pair of arguments and the output is which of them is of higher quality. While there are many, arguably independent, dimensions for quality (Wachsmuth et al., 2017b), it seems that people - and, consequently, algorithms - can usually perform this task in a consistent manner (Habernal and Gurevych, 2016; Wachsmuth et al., 2017a; Toledo et al., 2019; Gretz et al., 2020). To the best of our knowledge, this task was not previously considered in a multilingual setting.

In contrast with previous multilingual research on argument mining, in this work we address three different problems, of varying complexity, over a relatively large number of languages. This allows us to draw more holistic conclusions on the efficacy - and pitfalls - of transfer learning in the argument mining domain. It is interesting to compare these conclusions with other wide-scope multilingual NLU research, such as the XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020).

¹https://www.research.ibm.com/haifa/dept/vst/debating_data.shtml#MultilingualArgumentMining

3 Data Sets

3.1 Translated Data

English Datasets The sources for our translated training data and the "pseudo test" sets are two existing argument mining datasets in English, collected by our colleagues as part of our work on Project Debater.² One is a corpus of 30,497 arguments on 71 controversial topics, annotated for their stance towards the topic and for their quality (Gretz et al., 2020). This dataset (referred herein as ArgsEN) is used for the stance classification and argument quality tasks.

The second dataset is a corpus of 35,211 sentences from Wikipedia on 321 controversial topics, annotated for their stance towards the topic and the extent to which they can serve as an evidence for the topic (Ein-Dor et al., 2020). This dataset (referred as EviEN) is used for the stance classification and evidence detection tasks. Example 1 shows an argument and an evidence for one topic.

Example 1 (Argument and evidence)

Topic: *We should legalize cannabis*

Argument: *Cannabis can provide relief for a number of ailments without side effects.*

Evidence: *In 1999, a study by the Division of Neuroscience and Behavioral Health found no evidence of a link between cannabis use and the subsequent abuse of other illicit drugs.*

A third dataset was used to augment the training data for evidence detection and stance classification – the so-called VLD dataset of Ein-Dor et al. (2020), which includes around 200k sentences from newspaper articles, pertaining to 337 topics.³

Data Selection The ArgsEN dataset was filtered for stance classification by selecting arguments with a clear stance (confidence > 0.75) for training and evaluation. For argument quality, arguments with a clear agreement on their quality were selected – quality score above 0.9 or below 0.4.

A positive label for evidence detection was assigned for evidence from EviEN with a score above 0.7, and a negative label to those with score below 0.3 (those with in-between scores were not used). For stance classification on the evidence data, all

²Project Debater is the first AI system that can debate humans on complex topics: <https://www.research.ibm.com/artificial-intelligence/project-debater/>

³The stance labels of the EviEN and VLD datasets are not part of their official releases. They are included in our release.

sentences with a non-neutral stance were selected, since the EviEN dataset does not provide a confidence score for the stance label.

The VLD corpus was also filtered, taking sentences with evidence score above 0.95 or below 0.05. This yielded a total of 52,037 sentences for training, of which 19,406 have a positive label.

Translation We used the Watson Language Translator⁴ to translate the selected English data into 5 languages: German (DE), Dutch (NL), Spanish (ES), French (FR), and Italian (IT). The translation is a one-time process, which can be applied to any target language (TL) that the MT engine in use supports. The labels for the MT data were projected from the English data.

Following the data splits provided in the official release of each English corpus (into a training, development, and test sets), the translations of the training data were used for fine-tuning mBERT, and the translations of the test data were used for evaluation. The translations of the test data of ArgsEN and EviEN into the 5 non-English languages – the pseudo test sets – are herein referred to as ArgsMT and EviMT. The statistics of the translated English data are summarized in Table 1.

3.2 Human-Generated Data

Arguments written in the TL provide a more realistic evaluation set than translated texts, specifically for tasks where labels are not well-preserved across automatic translation. Therefore, we created a new multilingual evaluation set by collecting arguments in all 5 languages (ES, FR, IT, DE, and NL) for all the 15 topics of the ArgsEN test set, using the Appen⁵ crowdsourcing platform. The human-authored evaluation dataset is herein referred to as ArgsHG.

Annotation Setup Initially, crowd contributors wrote up to two pairs of arguments per topic, with one argument supporting the topic and another contesting it in each pair. Next, the arguments were assessed for their stance and quality by 10 annotators (per-language). Given an argument, they were asked to determine the stance of the argument towards the topic and to assess whether it is of high quality. The full argument annotation guidelines are included in the Appendix, and Table 2 details the number of arguments collected and labeled per

⁴www.ibm.com/watson/services/language-translator/

⁵www.appen.com

Set	ArgsEN				EviEN					
	#T	Stance		Quality	#T	Stance			Evidence	
		Pro	Con	#Args		Pro	Con	#T	Ev	Non-Ev
Train	49	10,162	9,766	8,373	171	5,592	3,622	174	3,522	14,275
Dev	7	1,564	1,497	1,329	46	1,726	1,202	47	1,145	3,967
Test	15	3,024	2,952	2,449	100	2,614	1,209	100	2,068	1,937
Total	71	14,750	14,215	12,151	317	9,932	6,033	321	6,735	20,179

Table 1: Statistics of the data selected from the **ArgsEN** and **EviEN** datasets and translated into 5 non-EN languages. For the tasks of *stance* classification, argument *quality* prediction and *evidence* detection, the table shows: the number of topics (**#T**) discussed by the arguments (from ArgsEN) or sentences (from EviEN) for each task; the number of **Pro** and **Con** arguments or sentences for stance classification; the number of arguments (**#Args**) for argument quality; the number of evidence (**Ev**) and non-evidence (**Non-Ev**) sentences for evidence detection.

language. To set a common standard, annotators were instructed to mark about half of the arguments they labeled as high quality. Annotation quality was controlled by integrating test questions (TQs) with an a-priori known answer in between the regular questions, measuring the per-annotator accuracy on these questions, and excluding underperformers.

A per-annotator average agreement score was computed by considering all peers sharing at least 50 common answers, calculating pairwise Cohen’s Kappa (Cohen, 1960) with each of them, and averaging. Those not having at least 5 peers meeting this criterion were excluded and their answers were discarded. Averaging the annotator agreements yields the average inter-annotator agreement (agreement- κ) of each question.

To derive a label (or score) for each question we use the WA-score of Gretz et al. (2020). Roughly, answers are aggregated with a weight proportional to the agreement score for the annotators who chose them. At least 5 answers were required for a question to be considered as labeled.

Scaling the annotation from English to new languages required some adjustments, such as restricting participation to countries in which the TL is commonly spoken, and the use of TQs for the argument quality question. Further details are provided in the Appendix.

Results Table 2 presents the agreement- κ for all TLs and each task for the human-generated dataset. For stance, the agreement is comparable to previously reported values for English (0.69 by Toledo et al. (2019) and 0.83 for ArgsEN). For quality, the agreement is significantly better than previously reported on ArgsEN (0.12 by Gretz et al. (2020)), presumably due to the use of TQs in this task, which

were not included before. The annotation in each of the non-EN languages involved a distinct group of annotators, producing varying annotation quality among languages which is reflected in their agreement- κ values.

The results also include the percentage of arguments labeled as supporting arguments, computed separately for each annotator and averaged over all annotators. All values are close to 0.5, confirming that the collected arguments are balanced for stance, as instructed. Similarly, the results show the percentage of arguments labeled as high quality, averaged over all annotators, confirming that annotators mostly followed the instruction to label about half of the arguments as high quality.

The same confidence filtering thresholds described in §3.1 were applied to the data of ArgsHG. The statistics of the arguments selected for evaluation are shown in Table 2 (right).

4 Experimental Setup

Our experiments are aimed at providing a comparative analysis of the translate-train approach when trained on different subsets of languages, and identifying when that approach is beneficial on the three argumentation tasks. We begin by describing the setup used in all experiments.

Training Configuration We used the BERT-Base multilingual cased model configuration (12-layer, 768-hidden, 12-heads, total of 110M parameters) with a sentence-topic pair input. Training was performed on one GPU.

The parameters configuration of the binary classification tasks, namely, stance classification and evidence detection, was: maximum sequence length of 128, batch size of 32, dropout rate of

Language		Collection							Evaluation			
		#C	Stance			Quality			Stance			Quality
			#L	κ	Sup.	#L	κ	HQ	#Args	Pro	Con	#Args
Spanish	ES	2995	2995	0.73	0.51	828	0.29	0.62	2541	1337	1204	440
French	FR	2201	1109	0.66	0.49	903	0.41	0.53	957	500	457	556
Italian	IT	3018	987	0.82	0.50	969	0.24	0.67	923	465	458	586
German	DE	1962	801	0.60	0.50	801	0.39	0.56	628	347	281	467
Dutch	NL	925	599	0.72	0.47	382	0.40	0.49	478	237	241	264

Table 2: Statistics of the ArgsHG multilingual arguments dataset, collected in five languages (See §3.2). On the left are statistics pertaining to its *collection*: the number of unique arguments collected (**#C**); the number of arguments labeled (**#L**) for their *stance* and *quality*; the agreement- κ obtained for each task; the average percentage of arguments labeled by each annotator as supporting the topic (**Sup.**) and as high-quality (**HQ**). On the right are statistics describing the *evaluation* data selected from ArgsHG for the stance and quality tasks: the number of arguments (**#Args**) selected for the evaluation of each task; for stance classification, the number of **Pro** and **Con** arguments within that selection.

0.1 and learning rate of $5e-5$. Each model was fine-tuned over 10 epochs, using a cross-entropy loss function. The regression model for argument quality prediction, similar to the one used by Gretz et al. (2020), used a maximum sequence length of 100, a batch size of 32, a dropout rate of 0.1 and a learning rate of $2e-5$. Each model was fine-tuned over 3 epochs, using a mean-squared-error loss function. In all cases, the model from the last epoch was selected for evaluation.

Translate-Train Models For each task, mBERT was trained using data translated into one of the target languages (ES, FR, IT, DE and NL). These per-language models, denoted herein as *TL*, are the simplest application of the translate-train approach. Two more models were trained for the language families that are represented in the above languages together with English: *RM* – for the Romance languages (ES, FR, IT), and *WG* – for the West-Germanic languages (EN, DE, NL). Each language family model was trained using the data of the languages in that family. Lastly, a model was trained on data from all 6 languages (denoted *6L*). To summarize, our evaluation includes 4 models based on the translate-train approach (*TL*, *6L*, *RM* and *WG*) for each task and TL.

Baselines Another mBERT model, denoted *EN*, was trained on the source English data. Using this model, the results of two baselines are reported: (i) zero-shot (denoted *ZS*) – which passes an input text in a non-English language to the *EN* model, utilizing the cross-lingual transfer capabilities of mBERT; (ii) translate-test (denoted *TT*) – in which

an input text is machine translated into English, and that translation is provided as input to the *EN* model. The results on English data are also reported as a performance benchmark for other languages. Obviously, for English, the results of the *TL* model and the *ZS* and *TT* baselines are identical.

Related and Distant Languages For each TL, we define two types of models. The *RL* model is the one trained on *related* languages from the same family as the TL, and *DL* is the model that is trained on languages that are more *distant* from the TL. In other words, given a TL, the *RL* model refers to the language family that includes the TL, and the *DL* model refers to the other family that does not include the TL. For example, in the case of the TL being German, *TL* denotes the model trained only on translated data in German; *RL* – the *WG* model trained on the 3 West-Germanic languages (including German); *DL* – the *RM* model trained on the 3 Romance languages.

Evaluation Metrics The reported metrics are macro-F1 for the classification tasks (stance classification and evidence detection), and Pearson correlation for the regression setting of argument quality.

5 Results

The results below for arguments (for stance classification on that data and argument quality) are averages over 5 evaluation runs of randomly initialized models that were trained in the same manner. For evidence sentences (stance classification on that data and evidence detection), the results are from a single evaluation run.

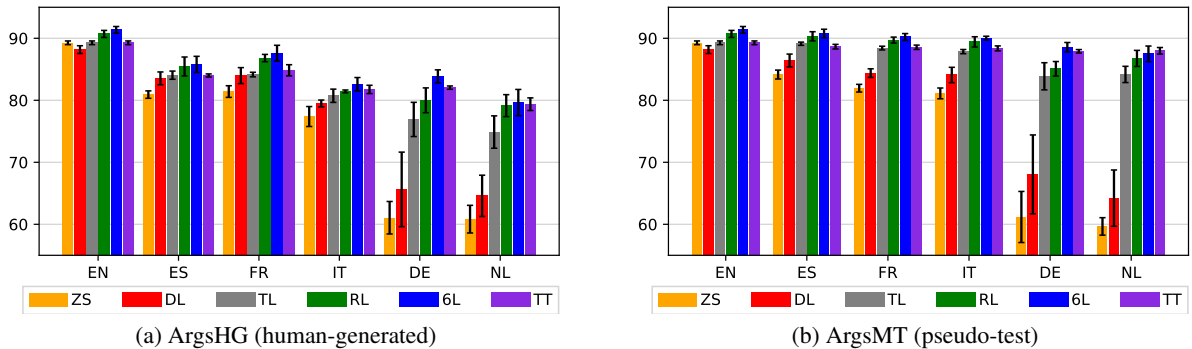


Figure 1: Stance classification results on arguments data, showing macro-F1 averaged over 5 evaluation runs, and its standard deviation. The results compare four translate-train methods (*DL*, *TL*, *RL* and *6L*), and two baselines – zero short (*ZS*) and translate-test (*TT*). See §5.1 for details.

5.1 Stance Classification

Arguments Figure 1a shows the evaluation results on the human-generated arguments of ArgsHG. For the non-English languages, the performance over the *ZS* baseline improves when adding translated data, even when that data is from distant languages (*DL*). The other baseline *TT* is better, yet the best performance is attained by the *6L* models – significantly so for 3 of the 5 languages (ES, FR and DE). Notably, ordering the translate-train models by their performance yields the same order for all languages: *DL* is always the worst, followed by *TL*, *RL* and the best performing model *6L*.

Repeating the same experiments on the pseudo-test data of ArgsMT resulted in similar trends, depicted in Figure 1b. Further augmentation of the training data with translations to more languages beyond the languages included in the training of the *6L* models (e.g. with 9 or 17 languages) did not significantly improve performance on these languages. These results are detailed in Table 5 within the Appendix.

Evidence To explore whether the observed trends are data-specific, we repeated the evaluation of the stance classification task with the EviMT dataset of evidence sentences from Wikipedia. Training was performed on the training set of that corpus (called Wikipedia models). The results on its pseudo-test set are depicted in Figure 2a. For the non-English languages, the best performing models are *6L* and *RL*, consistent with our findings for arguments (Figure 1).

The VLD evidence corpus allows further exploration of the stance classification task within the evidence domain. We trained models on a

larger dataset of translated evidence combining the Wikipedia data and selected data from the VLD corpus (called Extended models). Figure 2b shows the results obtained using these models. Overall, the performance of the Extended models is significantly better than the performance of the Wikipedia models, in almost all cases, and the *TL* models become competitive even with the *6L* models.

Performance on English In comparison with the *ZS* baseline (trained only on English), adding translated training data improves performance on English (leftmost bars in Figures 1a and 2a), for both domains. For the evidence data, even translations to distant languages (*DL*) help the Wikipedia model, yet when a lot of training data is available in English (leftmost bars in 2b), there is no significant gain from adding translations to the training set.

Summary Overall, the best performing models for the stance classification task are the *RL* and *6L* models, in both domains. Our finding that the *6L* models outperform the *TL* models is consistent with previous results on the XNLI task (Hu et al., 2020). Interestingly, translated data can be used to improve performance on English as well.

The *ZS* and *TT* baselines are almost always outperformed by the best translate-train model. However, when a large-scale English corpus is available (Figure 2b), the *TT* baseline becomes comparable to the best translate-train models.

5.2 Evidence Detection

To examine whether the above observations are task-specific, we move on to the task of evidence detection. The results for that task on the EviMT

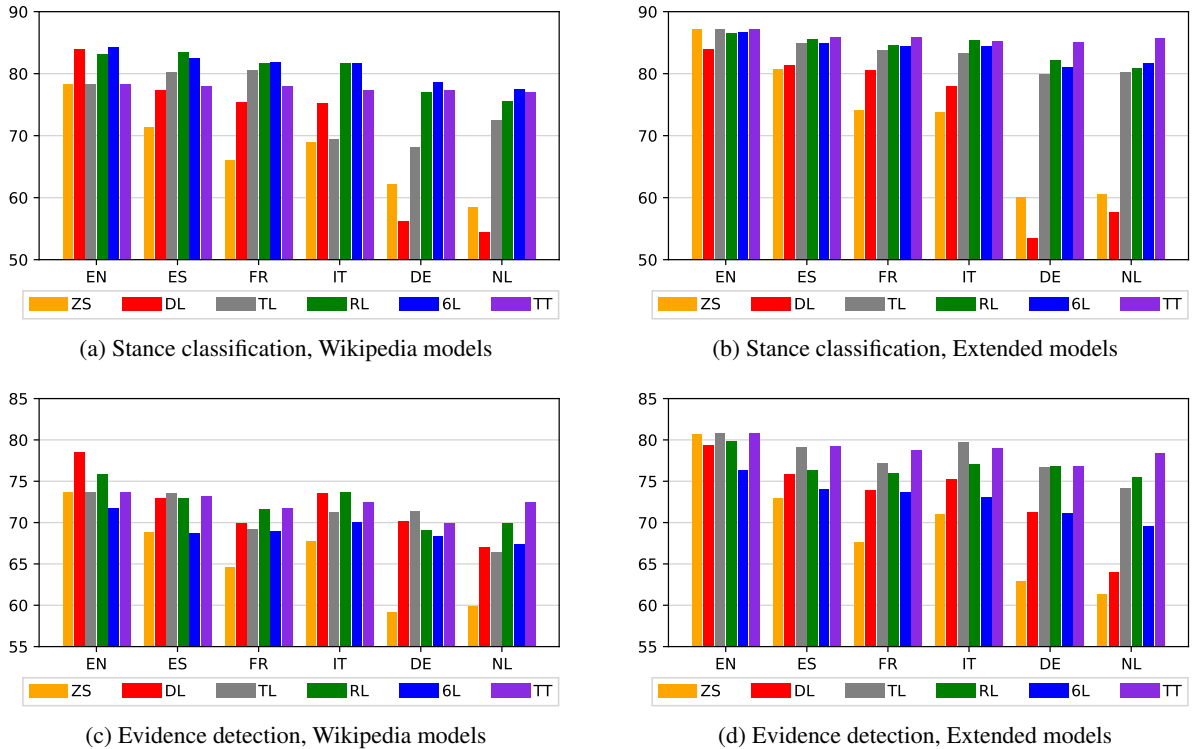


Figure 2: Stance classification (top) and evidence detection (bottom) macro-F1 results on the EviMT pseudo-test set with Wikipedia models (left) and Extended models (right). The results compare four translate-train methods (*DL*, *TL*, *RL* and *6L*), and two baselines – zero short (*ZS*) and translate-test (*TT*). See §5.2 for details.

pseudo-test are depicted in Figure 2c (Wikipedia models) and in Figure 2d (Extended models).

In contrast with the stance results, where in most cases the *6L* models were best, for evidence detection performance may degrade when adding languages. The best performing translate-train models are either the *TL* or *RL* models, in all cases.

As in the stance classification results for this corpus, the additional training data used in the Extended models improves performance. In addition, the English benchmark results for the Wikipedia models (leftmost bars in Figure 2c) can improve by adding languages, or by adding English data (*ZS* bar for EN in Figure 2d), but there is no significant gain from doing both (leftmost bars in Figure 2d).

5.3 Argument Quality Prediction

Moving to our last task, Figure 3 shows the Pearson correlation results on the human-generated arguments, between the predicted quality score and the labeled argument quality score. In contrast with the stance results, adding data from related languages (the *RL* bars) does not help, and training on the English dataset (the *ZS* bars) is sufficient to obtain a

competitive model.⁶ We suspect that the reason for this is that this task is more complex and nuanced than the previous two.

6 Analysis

The performance of a translate-train model may be affected, among other factors, by the translation quality, the extent in which a task-specific label is preserved across that translation, and, for our data, the discussed topic. These are analyzed below.

6.1 Translation Quality Assessment

We assessed our machine translation quality by computing the BLEU score (Papineni et al., 2002) between the English arguments from the test set of ArgsEN and the same arguments after translation to a TL and back to English. For all languages, these scores are above 0.5 (see Table 3), suggesting the translations are of high quality.

6.2 Translated Label Assessment

An important prerequisite for training and evaluating models on automatically translated texts is

⁶The performance of the *ZS* model on English is 0.61 – comparable to its performance for FR and DE.

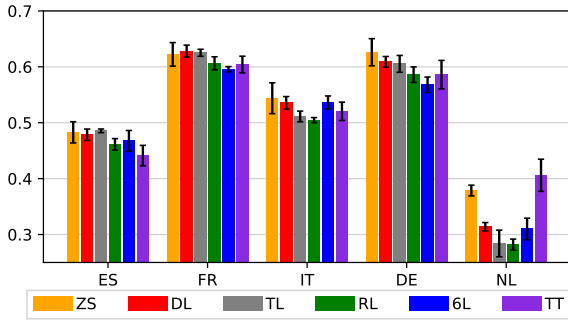


Figure 3: Argument quality prediction results (Pearson correlation) on the ArgsHG dataset (See §5.3).

ES	FR	IT	DE	NL
0.59	0.51	0.63	0.52	0.55

Table 3: Per-language BLEU scores between the ArgsEN test set arguments and their back-translation to English.

that the labels of the original texts are preserved under translation, which depends on the specific task at hand. Example 2 shows one argument and its translation to Spanish and back to English. The translation preserves the original stance, but the argument quality is degraded. Hence, we annotated a sample of the translated texts to assess how often this happens in each task. The annotation focuses on one Romance and one West-Germanic language – Italian and German.

Example 2 (Translation quality)

Topic: *We should ban algorithmic trading*

English argument: *Algorithmic trading results in unfair advantages for those able to access it to the detriment of ordinary investors.*

Back-translation: *The algorithmic trading of results in unjust advantages for those able to access it to the detriment of common investors.*

Annotation Setup 14 arguments were randomly sampled from each topic of the ArgsEN test set, yielding 210 arguments per language. Similarly, two sentences were sampled from each topic in the EviEN test set, producing 200 sentences per language. All texts were machine translated and human translated by native speakers of each TL. Both translations of each argument were labeled for their stance and quality, as in §3.2. Similarly, the potential evidence sentences were annotated for whether they are valid evidence, and those which are so were also annotated with their stance towards the topic, as in Ein-Dor et al. (2020). In this an-

	κ		HT	MT	HT	MT
	DE	IT	DE	DE	IT	IT
Stance-A	0.74	0.82	0.98	0.96	0.98	0.95
Stance-E	0.82	0.76	0.88	0.86	0.80	0.85
Qual	0.26	0.12	0.49	0.45	0.47	0.32
Det	0.38	0.33	0.74	0.75	0.81	0.84

Table 4: Translated labels assessment results for two languages and all tasks: stance classification on arguments (**Stance-A**) or evidence (**Stance-E**), argument quality (**Qual**) and evidence detection (**Det**). The results show the agreement- κ obtained in the annotations, and Pearson correlations between the original English labels and the labels of human (HT) and machine (MT) translated texts.

notation, TQs were formed from translated texts, with the correct answer taken from the English labels. The full evidence annotation guidelines are included in the Appendix.

Results Table 4 shows the assessment results for all tasks and the two languages. The obtained agreement- κ is on par with previously reported values for these tasks (as detailed in §3.2), though somewhat lower for evidence detection. The table further shows Pearson correlation between the original English WA-scores, and the WA-scores of the translated texts. For evidence detection and argument quality, this computation was performed on texts matching the criteria defined in §3.1. The correlation for evidence stance classification was computed on sentences with at least 6 stance labels on their translated version.

The results show that for both datasets, stance is well preserved after translation. For evidence detection, the correlation is lower, yet the difference between MT and HT is small, suggesting the change in the labels is not due to the automatic translation. Thus, the use of translated texts in these tasks is acceptable, for both training and evaluation.

For argument quality, the correlation is considerably lower, and there is a significant difference between MT and HT in IT, as may be expected for such a nuanced task. This could be the reason that the translate-train models do not improve performance for this task – since the quality label is not maintained when an argument is translated, projecting these labels into translated texts introduces significant noise into such training data.

6.3 Per-Topic Analysis

In both of our data domains, arguments and evidence, the texts are relevant to a specific topic, and the obtained performance may depend on that topic in various ways. Focusing on stance classification, we measured the per-topic performance on the human-generated arguments of ArgsHG. Figure 4 shows these results averaged over the 5 non-English languages, for the *TL* and *6L* translate-train models, and the *ZS* and *TT* baselines. The topics are ordered by their performance on English.⁷

The results demonstrate the performance variability among the different topics. For some, the average performance on the non-English languages is close to their performance on English (e.g. topics 9 or 10), yet for others it is far from it (e.g. topic 5). The performance of the *ZS* baseline is low for topics 2 through 8, from which 5 are discussing imposing a ban. This implies that the stance towards the discussed topic, or the "action" within the topic (e.g. ban, legalize, etc.) may be an important factor.

We further manually analyzed the results on Topic 10, a low-performing outlier in French for the *ZS* baseline (see Figure 5 in the Appendix). A native speaker examined 3 batches of 20 arguments each, containing: 1) prediction errors from that topic; 2) randomly sampled correct predictions from the same topic; 3) all 4 prediction errors and 16 randomly sampled correct predictions from the topic with the highest performance (Topic 1). Within the first batch, 40% of the samples were incoherent or syntactically wrong arguments, compared to only 20% in each of the other two batches.

7 Conclusions

We have examined the translate-train paradigm for three multilingual argument mining tasks: stance classification, evidence detection, and argument quality, evaluating a wide range of multilingual models on machine-translated and human-authored data. These tasks differ in their complexity, as reflected in the agreement of annotators on the correct label, the extent to which this label is preserved across translation, and, ultimately, in the accuracy of the models.

Accordingly, our results show that the translate-train approach is well suited for stance classification, as performance improves when augmenting the English training data with automatic translations from other languages. For evidence detection,

⁷Table 6 in the Appendix lists the topic for each topic ID.

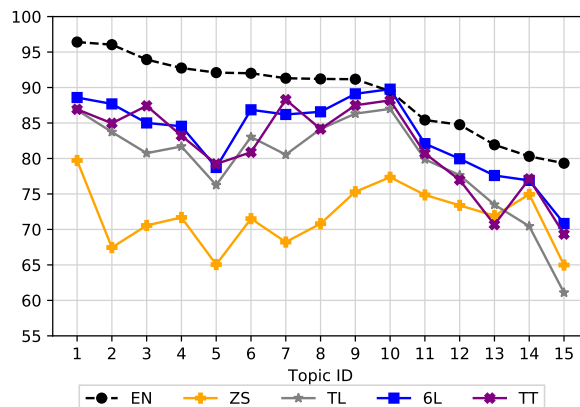


Figure 4: Per-topic stance classification average macro-F1 results on ArgsHG, averaged over the 5 non-English languages. The topics are ordered by their performance on English with the *EN* model (dashed black line).

adding data from the target language or related languages improves performance, yet adding more languages is not helpful.

For both tasks on the evidence data, adding more English training data improves performance. In these cases, augmenting the large English training set with data of other languages only leads to a marginal gain for stance classification, and even degrades performance for evidence detection.

In contrast with the above two tasks, the results on argument quality show that training only on English is at least as good, if not better, than any translate-train model. This is reflected by the clearly opposite trends observed in Figure 3 vs. those observed in Figure 1a.

Taken together, our results confirm the validity of the common translate-train paradigm for argument mining tasks such as stance and evidence detection, for which the label is relatively well preserved under translation. However, for the more subtle argument-quality task, where the label – as might be expected – is far less preserved, a new approach might be needed. Future work might wish to explore how translation can preserve not only the semantics of texts, but also finer aspects that contribute to its quality.

Acknowledgments

Special thanks go to Michal Jacovi, Christopher Sciacca, and Mohamed Nasr for coordinating the manual translation effort. We thank the teams who performed the translation, and Benjamin Sznajder for helping with the analysis. We also thank the anonymous reviewers for their valuable comments.

References

- Ahmet Aker and Huangpan Zhang. 2017. Projection of Argumentative Corpora from Source to Target Languages. In *Proceedings of the 4th Workshop on Argument Mining*, pages 67–72.
- Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. 2017. **Stance Classification of Context-Dependent Claims**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 1, pages 251–261.
- Pierpaolo Basile, Valerio Basile, Elena Cabrio, and Serena Villata. 2016. Argument Mining on Italian News Blogs. In *Proceedings of Third Italian Conference on Computational Linguistics (CLiC-it 2016) & Fifth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2016)*, Napoli, Italy.
- Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Alexis Conneau, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating Cross-lingual Sentence Representations. *arXiv preprint arXiv:1809.05053*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding**. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Steffen Eger, Johannes Daxenberger, Christian Stab, and Iryna Gurevych. 2018a. **Cross-lingual Argumentation Mining: Machine Translation (and a bit of Projection) is All You Need!** In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 831–844, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Steffen Eger, Andreas Rücklé, and Iryna Gurevych. 2018b. **PD3: Better Low-Resource Cross-Lingual Transfer By Combining Direct Transfer and Annotation Projection**. In *Proceedings of the 5th Workshop on Argument Mining*, pages 131–143.
- 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. 2020. **Corpus Wide Argument Mining – a Working Solution**. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Irina Fishcheva and Evgeny Kotelnikov. 2019. Cross-Lingual Argumentation Mining for Russian Texts. In *International Conference on Analysis of Images, Social Networks and Texts*, pages 134–144. Springer.
- Shai Gretz, Roni Friedman, Edo Cohen-Karlik, Assaf Toledo, Dan Lahav, Ranit Aharonov, and Noam Slonim. 2020. **A Large-scale Dataset for Argument Quality Ranking: Construction and Analysis**. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*.
- Ivan Habernal and Iryna Gurevych. 2016. **Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1589–1599.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. **XTREME: A Massively Multilingual Multi-task Benchmark for Evaluating Cross-lingual Generalization**. *arXiv preprint arXiv:2003.11080*.
- Mirko Lai, Alessandra Teresa Cignarella, Delia Irazú Hernández Farías, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. Multilingual Stance Detection in Social Media Political Debates. *Computer Speech & Language*, page 101075.
- John Lawrence and Chris Reed. 2020. **Argument mining: A survey**. *Computational Linguistics*, 45(4):765–818.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. **XGLUE: A New Benchmark Dataset for Cross-lingual Pre-training, Understanding and Generation**. *arXiv preprint arXiv:2004.01401*.
- Wei-Hao Lin, Theresa Wilson, Janyce Wiebe, and Alexander G Hauptmann. 2006. **Which Side Are You On? Identifying Perspectives at the Document and Sentence Levels**. In *Proceedings of the Tenth Conference on Computational Natural Language Learning (CoNLL-X)*, pages 109–116.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **BLEU: a Method for Automatic Evaluation of Machine Translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-2002)*, page 311–318.
- Ruty Rinott, Lena Dankin, Carlos Alzate Perez, Mitesh M. Khapra, Ehud Aharoni, and Noam Slonim. 2015. **Show Me Your Evidence - an Automatic Method for Context Dependent Evidence Detection**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 440–450, Lisbon, Portugal. Association for Computational Linguistics.
- Benjamin Schiller, Johannes Daxenberger, and Iryna Gurevych. 2020. **Stance Detection Benchmark: How Robust Is Your Stance Detection?** *arXiv preprint arXiv:2001.01565*.

- Parinaz Sobhani, Diana Inkpen, and Stan Matwin. 2015. [From Argumentation Mining to Stance Classification](#). In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 67–77.
- Manfred Stede and Jodi Schneider. 2018. *Argumentation Mining*. Morgan & Claypool Publishers.
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. *arXiv preprint cs/0607062*.
- Assaf Toledo, Shai Gretz, Edo Cohen-Karlik, Roni Friedman, Elad Venezian, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. 2019. [Automatic Argument Quality Assessment - New Datasets and Methods](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5625–5635, Hong Kong, China. Association for Computational Linguistics.
- Jannis Vamvas and Rico Sennrich. 2020. [X-Stance: A Multilingual Multi-Target Dataset for Stance Detection](#). *arXiv preprint arXiv:2003.08385*.
- Henning Wachsmuth, Nona Naderi, Ivan Habernal, Yufang Hou, Graeme Hirst, Iryna Gurevych, and Benno Stein. 2017a. [Argumentation Quality Assessment: Theory vs. Practice](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 250–255, Vancouver, Canada. Association for Computational Linguistics.
- Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017b. [Computational Argumentation Quality Assessment in Natural Language](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 176–187.
- Elena Zotova, Rodrigo Agerri, Manuel Nuñez, and German Rigau. 2020. [Multilingual Stance Detection: The Catalonia Independence Corpus](#). *arXiv preprint arXiv:2004.00050*.

A Appendices

A.1 Additional Results

A.1.1 Stance Classification on Machine-Translated Arguments

As described in §4, our evaluations were conducted on 6 European languages (EN, ES, FR, IT, DE, and NL). Models were trained for several language groups: *RM* – for the Romance languages, *WG* – the West-Germanic, and *6L* – a model that covers all the TLs in our evaluation.

We further explored the translate-train approach by augmenting the training data of our models with machine-translated data of other language families. First, we trained a model for the North-Germanic family (*NG*) with three languages – Danish (DA), Swedish (SV), and Norwegian (NB). Next, we combined the Romance languages with the two German families (*RM*, *WG*, and *NG*), and created the *9L* model with 9 languages. Finally, we trained a model with a relatively large number of languages (17) and a variety of language families. This model, denoted *17L*, consists of all the languages in *9L* and 8 additional languages: Slavic languages – Polish (PL), Slovak (SK), Russian (RU); Semitic – Arabic (AR), Hebrew (HE); and Chinese/Japonic – Simplified Chinese (ZH), Traditional Chinese (ZT), and Japanese (JA).

The stance classification results on the EviMT pseudo-test for all 17 languages using all the aforementioned models are presented in Table 5. We see that expanding the training set beyond the six languages in *6L* by adding more distant languages, as in the *9L* and *17L* models, does not significantly improve the performance on English. On average, training on the TL is better than training on the original English arguments (average performance over all 17 languages is 73.7% with *EN* and 86.5% with *TL*). Training on all 17 languages tends to yield the best performance (with an average of 88.9%), though training on a subset of them is often nearly as good - and sometimes even better, especially on the *6L* and *9L* groups.

A.1.2 Per-Topic Analysis Information

Table 6 contains a list of 15 topics that are included the test set of the ArgsEN dataset, along with their IDs used during error analysis.

Figure 5 shows per-language results of the zero-shot baseline on the 15 topics of the ArgsHG evaluation dataset. As indicated by the overall results on the same data in Figure 1a, we see high performance on the three Romance languages consistently across most topics, and low performance on DE and NL for about half of the topics.

A.2 Annotation Details

This section describes further annotation details, such as the adjustment of the argument assessment annotation task to multiple languages, and the guidelines used in each annotation task.

Model	EN	DE	NL	ES	FR	IT	DA	SV	NB	PL	SK	RU	AR	HE	ZH	ZT	JA
<i>EN</i>	89.3	61.2	59.7	84.2	82.0	81.1	69.0	72.8	66.7	74.8	74.3	73.6	62.7	71.9	81.4	80.4	67.7
<i>TL</i>	89.3	83.9	84.2	89.1	88.4	87.9	84.8	86.6	87.0	84.8	87.5	86.8	85.3	85.5	87.7	87.5	83.6
<i>ES, FR, IT</i>	88.2	68.0	64.2	90.3	89.7	89.4	65.2	74.3	74.4	73.7	72.8	75.0	71.1	77.0	83.1	84.0	70.5
<i>EN, DE, NL</i>	90.7	85.1	86.8	86.4	84.4	84.1	79.8	76.5	79.6	78.1	75.9	79.1	68.7	67.4	80.8	81.4	70.9
<i>SV, NB, DA</i>	83.7	64.0	68.3	82.8	80.1	79.0	88.8	88.9	88.6	72.8	73.8	76.7	64.1	71.1	79.1	79.6	67.3
<i>PL, SK, RU</i>	80.5	60.1	56.5	80.8	82.7	81.5	75.8	77.4	74.2	86.8	88.3	88.4	70.9	75.4	82.1	81.5	66.3
<i>AR, HE</i>	81.2	60.7	54.7	78.1	79.2	78.9	66.1	71.9	63.0	75.7	73.6	74.2	85.7	86.2	80.0	81.3	66.5
<i>ZH, ZT, JA</i>	81.3	63.8	55.6	81.3	69.3	78.7	61.6	65.6	61.9	72.4	70.9	75.8	62.9	71.5	88.6	88.3	85.8
<i>6L</i>	91.4	88.6	87.5	90.8	90.2	90.0	70.6	81.2	75.1	79.4	73.2	83.4	69.2	70.2	81.9	82.3	73.2
<i>9L</i>	91.3	86.6	88.8	90.9	90.3	90.0	89.7	89.0	89.7	76.7	72.3	81.7	70.6	71.7	83.0	84.1	73.2
<i>17L</i>	91.5	86.8	88.8	90.7	90.5	90.0	89.3	88.9	90.0	88.3	89.1	88.9	87.8	86.9	88.9	88.8	86.8

Table 5: Stance classification macro-F1 results on the ArgsEN test set in English (leftmost column) and on the ArgsMT evaluation set in 16 languages with models trained on various language groups (see §A.1.1). The results in italics are showing averages of 5 training runs. The other results are from a single training run of each model.

ID	Topic
1	We should abolish the Olympic Games
2	We should ban factory farming
3	We should ban algorithmic trading
4	We should ban targeted killing
5	We should prohibit school prayer
6	We should ban private military companies
7	We should adopt libertarianism
8	We should ban missionary work
9	Social media brings more harm than good
10	We should legalize cannabis
11	We should abolish the three-strikes laws
12	We should prohibit women in combat
13	Holocaust denial should be a criminal offence
14	The use of public defenders should be mandatory
15	We should adopt atheism

Table 6: Topics and IDs of the ArgsEN test set.

A.2.1 Multilingual Argument Assessment

While Gretz et al. (2020) mention using a group of annotators with whom they have worked before for the assessment of arguments written in English, no such group was available to us for the non-English languages. In addition, no test questions (TQs) were available, since they are typically formed from existing labeled data. Initially, the first issue was addressed by relying on workers from appropriate countries, and the second by using machine-translated arguments from ArgsEN, with a high-confidence label *in English*, as TQs. At first, since the quality label is sensitive to translation (as described in §6.2) such TQs were limited

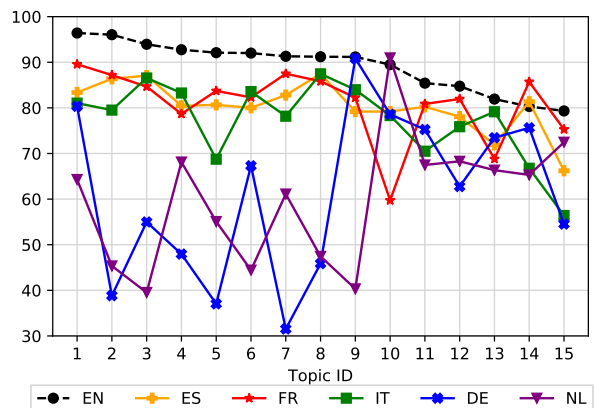


Figure 5: Per-topic stance classification average macro-F1 results per language on ArgsHG test set for the ZS baseline. The dashed black line is the performance on English.

to stance.

A pilot on Spanish arguments showed a good agreement- κ for stance (0.71), yet a low value for quality (0.04). The results showed that many of the annotators labeled a vast majority (>80%) of the arguments as high-quality, even though they were instructed to consider only half as such. Therefore, only those labeling $\leq 80\%$ of arguments as high-quality were allowed further work. Others were excluded and their argument quality answers were ignored.

A second pilot extended this procedure to other languages. However, the size of the workforce meeting the above criteria was small for DE, FR and NL, preventing progress altogether for the last two. This required integrating TQs for quality question despite the risk of the quality label changing due to the automatic translation. To mitigate that risk, one of the authors carefully monitored each annotation task, reviewed TQs which many anno-

tators answered incorrectly, and disabled those in which the translation introduced errors in the correct label or made the text unclear.

A.2.2 Argument Authoring Guidelines

Below is an example of the argument authoring guidelines for German. The guidelines for the other languages were similar.

PLEASE READ:

All your submitted arguments will be assessed for their quality. For each argument determined as a high-quality one, you will receive a bonus of up to 0.4\$.

Overview

In the following task you are presented with a debatable topic, to which you should suggest high quality supporting/contesting arguments in German.

A supporting/contesting argument will be considered as a high-quality one, if a person preparing a speech to support/contest the topic, respectively, will be likely to use this argument as is in her speech.

Note: Copying texts from the web or elsewhere is prohibited. The content you provide must be written by you in your own language.

Requirements

- The argument must be phrased in German.
- The argument must either clearly support or clearly contest the topic.
- You should write a single argument in each text box.

A.2.3 Argument Assessment Guidelines

In this annotation, the guidelines for all languages were the same.

In the following task you should answer two questions concerning an argument suggested in the context of a debatable topic.

1. What is the stance of the argument towards the topic? (supporting, contesting or neutral)
2. For someone with this stance towards the topic, is this a high-quality argument to use? (yes or no)

IMPORTANT! For the second question please answer "YES" only for high-quality arguments, and only for about half of the time.

Your answers will be monitored not only using test questions. If you are interested in participating in future similar tasks, please answer thoroughly.

A.2.4 Evidence Assessment Guidelines

Below is an example of the evidence assessment guidelines for German. The guidelines for the other languages were similar.

General instructions

In this task you are given a topic and evidence candidates for the topic. The candidates are in German. Consider each candidate independently. For each candidate please select Accept if and only if it satisfies ALL the following criteria:

- The candidate clearly supports or clearly contests the given topic. A candidate that is neutral towards the topic should not be accepted.
- The candidate represents a coherent, stand-alone statement, that one can articulate (nearly) "as is" while discussing the topic, with no need to change/remove/add more than two words.
- The candidate represents valuable evidence to convince one to support or contest the topic. Namely, it is not merely a belief or merely a claim, rather it provides an indication whether a belief or a claim is true. A candidate which presents detailed information (typically quantitative) that clearly support or clearly contest the topic, should be accepted.

If you select Accept, you should further indicate whether the evidence supports the topic (Pro) or contests it (Con).

Note: if you are unfamiliar with the topic, please briefly read about it in a relevant data source like Wikipedia.

Examples

The following examples outline several candidates along with their suggested annotations; please read all these examples before performing the task.

Topic: We should ban the sale of violent video games to minors.

Example 1

The research clearly suggests that, among other risk factors, exposure to violent video games can lead to aggression and other potentially harmful effects.

Annotation: Accept – Pro

Note: even though the text is not explicitly referring to the proposed ‘ban’ policy, it should still be accepted, since highlighting the negative aspects of violent video games can be used to support the suggested ban.

Example 2

A university of Oxford study negates the idea that violent video game content leads to violence.

Annotation: Accept - Con

Note: here as well, even though the proposed ‘ban’ policy is not explicitly mentioned, the text should be accepted since clearly it can be used to contest the suggested ban.

Example 3

There is no reason to suppose that violent video games cause harm to children.

Annotation: Reject

Reason: The candidate states a claim. It does not offer any additional information to convince the reader that this claim is true.

Example 4

The American Psychological Association argues that violent video-game play leads to increased moral sensitivity.

Annotation: Accept - Con

Reason: The candidate states a claim, but the fact that it is raised by an authority figure (organization or human) turns it into a valuable evidence.

Example 5

Kennelly said there is no scientific evidence that violent video games cause “serious harm” in kids such as heightened aggression that would require protection of the law.

Annotation: Accept - Con

Note: If you are not certain whether the speaker is an authority figure or not, you should typically give him/her the benefit of the doubt and consider them as such (in this case the speaker is Matthew F. Kennelly, a United States District Judge). However, if the candidate states a claim and the speaker is only mentioned by he/she/they you should reject it.

Example 6

The issue as “Psychological research confirms that violent video games can increase children’s aggression.”

Annotation: Reject

Reason: The candidate does not represent a coherent, stand-alone statement.

Example 7

Some studies have clearly demonstrated that video game violence is leading to serious aggressive behaviour in real life, although other studies have shown the opposite.

Annotation: Reject

Reason: The pro/con stance of the candidate towards the topic is unclear, since the end of the text contradicts its beginning.

Example 8

The Entertainment Software Association reports that 17% of violent video game players are boys under the age of eighteen.

Annotation: Reject

Reason: The candidate states a fact with no clear pro/con stance towards the topic.

Example 9

Studies show that watching violent movies increases aggression amongst youth.

Annotation: Reject

Reason: The candidate is not related to the topic as it discusses violent movies and not violent video games.

Example 10

Another 2001 meta-analyses and a more recent 2009 study focusing specifically on serious

aggressive behavior concluded that video game violence is not related to serious aggressive behavior in real life.

Annotation: Accept - Con

Note: Even though the candidate's first word better be omitted to make it a stand-alone statement, this is a minor change which is acceptable.

Example 11

Limiting the sale of violent video games will cause 15,000 people to lose their jobs.

Annotation: Accept - Con

Note: The candidate presents a specific numeric piece of information that clearly contest the topic. You are not expected to fact check the provided piece of information, don't reject such a candidate just because you are not sure that the provided piece of information is true.