

Crosslingual Argument Mining in the Medical Domain

Minería de Argumentos Crosslingüe en el Dominio Médico

Anar Yeginbergen, Rodrigo Agerri

HiTZ - Ixa, University of the Basque Country UPV/EHU

{anar.yeginbergen,rodrigo.agerri}@ehu.eus

Abstract: Nowadays the medical domain is receiving much more attention in applications involving Artificial Intelligence as clinicians decision-making is increasingly dependent on dealing with enormous amounts of unstructured textual data. In this context, Argument Mining (AM) helps to meaningfully structure textual data by identifying the argumentative components in the text and classifying the relations between them. However, as it is the case for many tasks in Natural Language Processing in general and in medical text processing in particular, the large majority of the work on computational argumentation has been focusing only on the English language. In this paper, we investigate several strategies to perform AM in medical texts for a language such as Spanish, for which no annotated data is available. Our work shows that automatically translating and projecting annotations (data-transfer) from English to a given target language is an effective way to generate annotated data without costly manual intervention. Furthermore, and contrary to conclusions from previous work for other sequence labelling tasks, our experiments demonstrate that data-transfer outperforms methods based on the the crosslingual transfer capabilities of multilingual pre-trained language models (model-transfer). Finally, we show how the automatically generated data in Spanish can also be used to improve results in the original English monolingual setting, providing thus a fully automatic data augmentation strategy. Data, code, and fine-tuned models are publicly available at <https://huggingface.co/datasets/HiTZ/AbstRCT-ES>.

Keywords: Argumentation, Information Extraction, AI in Healthcare, Multilinguality.

Resumen: La tecnología basada en Inteligencia Artificial tiene una gran potencialidad para desarrollar asistentes que ayuden a profesionales médicos en la toma de decisiones, las cuales en muchos casos están basadas en el procesamiento de una gran cantidad de textos no-estructurados. En este contexto, la minería de argumentos (AM) puede ayudar a estructurar los datos textuales en componentes argumentativos y las relaciones discursivas existentes entre ellos. Sin embargo, al igual que todavía ocurre en muchas tareas de Procesamiento del Lenguaje Natural, la gran mayoría del trabajo sobre argumentación computacional en el dominio médico se ha centrado únicamente en inglés. En este artículo investigamos varias estrategias para realizar AM en textos médicos para un idioma como el español, para el cual no existen datos manualmente etiquetados. Nuestro trabajo muestra que traducir y proyectar automáticamente anotaciones del inglés a un idioma de destino determinado como el español es una forma eficaz de generar datos anotados sin necesidad de realizar anotación manual. Por otra parte, se demuestra experimentalmente que traducir y proyectar obtiene mejores resultados que los métodos basados en las capacidades de transferencia crosslingüe de modelos de lenguaje multilingües. Finalmente, usamos los datos automáticamente generados para español para mejorar los resultados originales en inglés, proporcionando así una estrategia de aumento de datos totalmente automática.

Palabras clave: Argumentación, Extracción de Información, Inteligencia Artificial en Medicina, Multilingualismo, Etiquetado Secuencial, Procesamiento del Lenguaje Natural.

1 Introduction

Clinical decision-making is an essential process in which a medical doctor, for example, needs to identify and diagnose a disease and prescribe a treatment based on the patient’s health condition and the results of the clinical tests. However, it can also involve multiple challenges and add up stress, due to several reasons. First, there may be a large diversity of symptoms, one or more of which could be a sign of multiple diseases. Second, there exists an overwhelming and ever-increasing amount of data from previous patients with similar symptoms which is becoming increasingly difficult, if not impossible, to process manually. Lastly, the final decision should take into account the latest results published in Evidence-based Medicine reports (Sackett et al., 1996).

These challenges highlight the importance of medical professionals to be supported by AI-based technology able to extract, from the huge quantity of unstructured textual data available for the different diseases and treatments, relevant and timely information structured in a manner that is suitable to be easily analyzed.

Argument mining (AM) is a field of Natural Language Processing (NLP) that focuses on extracting argumentative structures from unstructured textual content. While there are a large number of theories on argumentation, in this work we follow the model proposed by Stab and Gurevych (2014) according to which the main objective of AM is to generate hierarchical argumentative tree structures by automatically identifying in running text the types of argument components (e.g., *evidences* or *claims*) and their boundaries (a sequence labelling task) and the relations held between them (e.g., *support* or *attack*), addressed as pair-wise classification. Given that such argumentative structures constitute the basis of evidence-based reasoning applications, research on AM for the medical domain has the potential of facilitating the development of enabling technology to help in clinical decision-making.

A number of works have been proposed to automatically detect, classify and assess the quality of argumentative structures for various specific domains such as law (Mochales and Ieven, 2009), biomedicine (Mayer et al., 2021; Accuosto, Neves, and Saggion, 2021), reviews (Li et al., 2017) and persuasive es-

says (Stab and Gurevych, 2014; Stab and Gurevych, 2017) but the large majority of them have been focused on English, especially for the medical domain.

Although some researchers have translated Persuasive Essays to study crosslinguality in AM for languages such as Chinese, German, French, Spanish, or Portuguese (Eger et al., 2018; Sousa et al., 2021), the generated data and method are not suitable for our purposes. First, translation and label projection for argument components have been done via manual translation or by applying fully automatic methods. However, while interesting, manually translating and projecting the labels is a route to avoid in the *data-transfer* approach since the objective is to reduce manual work to a minimum by applying Machine Translation (MT). Second, the automatic projection of labels should be manually checked *at least* for evaluation (García-Ferrero, Agerri, and Rigau, 2022) using both *model-transfer* and *data-transfer* methods. Finally, for the medical domain the only existing AM dataset is in English, which means that for a major language such as Spanish no annotated dataset for AM in the medical domain is currently available. In order to address these issues, this work empirically investigates optimal crosslingual transfer techniques to perform AM in medical texts for a target language such as Spanish for which no manually annotated data is available.

In this work we leverage AbstrCT, at the time of writing the only existing labeled dataset in English for AM in the medical domain (Mayer et al., 2021), to perform crosslingual transfer about argument components and relations from English to Spanish. More specifically, we experiment with: (i) *model-transfer* in which a multilingual Masked Language Model (MLM) is fine-tuned in English and the predictions are generated in Spanish; (ii) *data-transfer* techniques based on machine translating and projecting the argument component labels to generate annotated data for Spanish with a minimum effort and, (iii) *automatic data augmentation* by exploiting the multilingual capabilities of MLMs with the aim of improving results in the original source English gold-standard data. Summarizing, the main contributions of this paper are the following:

- We present the first Spanish dataset annotated with the argument components and re-

lations at the sequence level for the medical domain.

- We perform the first experiments comparing *model-transfer* and *data-transfer* methods for AM in order to establish which strategy works best when no annotated data is available for a target language. Contrary to previous work for other sequence labelling tasks (García-Ferrero, Agerri, and Rigau, 2022), our results demonstrate the superiority of *data-transfer* for argument component detection in the medical domain.
- We show how *data-transfer* can easily and effectively be used to perform fully automatic data augmentation to improve results in the original English evaluation dataset.
- We establish that training in *data-transfer* with projected labels (with no manual correction) obtains similar performance to training with manually-checked labels, outperforming also *model-transfer* results. This is a highly encouraging outcome of our work, as it means that we should be able to discard any manual correction of the projected labels, making the whole *data-transfer* process fully automatic.
- The generated Spanish annotated data (both with automatically projected labels and its manually revised version) are publicly available to encourage crosslingual research in argument mining and to facilitate the reproducibility of results¹.

2 Related Work

One of the pioneer works with a clear influence in computational argumentation proposed a number of functional roles for arguments based on the manner in which the conclusion is made based on textual evidence: evidence, warrant, backing, qualifier, rebuttal, and claim (Toulmin, 1958). Furthermore, another important contribution argued that the most important relationship types between argument components are those of *support* and *attack* (Peldszus and Stede, 2013). Interestingly, they also defined five types of argumentation graphs according to the number and type of existing relations between the argument components: one claim having relations with multiple premises, a claim followed by another claim, etc.

Aiming to develop a fully computational AM approach, Stab and Gurevych (2014;

2017) proposed a model in which the objective of AM was to generate hierarchical structures to represent the argument components (*claims* and *premises*) in a manually annotated dataset of Persuasive Essays and their relations (*support* and *attack*). Their model has been widely adopted in NLP, also for the medical domain (Stab and Gurevych, 2017; Eger et al., 2018; Accuosto, Neves, and Saggion, 2021; Sousa et al., 2021; Mayer et al., 2021).

2.1 AM in the Medical Domain

Argument mining can be very valuable in the medical domain, particularly in Evidence-based Medicine where argumentative explanations are given in order to justify or discard a given diagnosis or treatment (Mayer et al., 2021). However, the few previous works on argument mining in the medical domain are all focused on English. In fact, to the best of our knowledge, there is only one dataset, AbstRCT, which provides annotated argument components and relations in the medical domain (Mayer et al., 2021). This is partially due to the inherent difficulty of obtaining medical data to start with, but also because of the cost and complexity of manually annotating argumentative structures. Mayer et al. (2021) present a large battery of experiments using various English MLMs specific to the medical domain (Beltagy, Lo, and Cohan, 2019; Lee et al., 2019). They also demonstrate that combining Transformer-based models with CRF, LSTM or GRU helped to improve results.

With respect to other proposals, Green et al. (2014) provided an analysis of arguments in biomedical data and defined argumentation schemes and inter-argument relationships. Alamri and Stevenson (2016) created a corpus using research abstracts of studies considered in systematic reviews related to cardiovascular diseases where the objective was solely to identify contradictory claims. Shankar, Tu, and Musen (2006) developed a tool for healthcare where one of the functionalities is extracting evidence for any treatment-related claims based on an argumentative structure defined by Toulmin (1958). Finally, Accuosto, Neves, and Saggion (2021) generated a dataset of arguments annotated at sentence level for bioscientific documents.

¹<https://huggingface.co/datasets/HiTZ/AbstRCT-ES>

2.2 Crosslingual Transfer

Crosslingual transfer techniques for sequence labelling aim at transferring knowledge from rich source languages (usually English) to other less resourced languages with the aim of mitigating the lack of annotated data in the target language (Yarowsky et al., 2001).

Most of the work was initially based on leveraging some kind of parallel data or representations like crosslingual embeddings for tasks such as POS tagging (Yarowsky et al., 2001; Gaddy et al., 2016), named entity recognition (NER) (Gaddy et al., 2016; Yang, Salakhutdinov, and Cohen, 2017; Agerri et al., 2018) and extractive question answering (Lewis et al., 2020; Artetxe, Ruder, and Yogatama, 2020), among others (Chen et al., 2019; Liu et al., 2020).

García-Ferrero, Agerri, and Rigau (2022) provide an exhaustive comparison between *model-transfer* and *data-transfer* approaches on a variety of sequence labelling tasks, datasets, and languages. They conclude that *model-transfer* using large MLMs such as XLM-RoBERTa-large (Conneau et al., 2020) outperforms *data-transfer* (translate and project) methods in every evaluation setting.

For AM there have also been works applying translation and label projection from English to other target languages. In this case, the objective was to compare the downstream task results when training with manually translated and projected data with respect to fully automatic one (Eger et al., 2018; Sousa et al., 2021).

In order to study crosslingual transfer for AM in medical texts, in this paper we develop the first Spanish corpus annotated with argument components and relations for the medical domain. This allows us to present the first experiments on *model-transfer* and *data-transfer* methods for AM. Surprisingly, and contrary to previous conclusions for other sequence labelling tasks (García-Ferrero, Agerri, and Rigau, 2022), our results demonstrate the superiority of *data-transfer* for AM in the medical domain.

3 Translate and Project

In this section we describe the steps undertaken in order to generate a Spanish version of AbstrCT (Mayer et al., 2021) by applying MT and label projection for argument components (*claims* and *premises*). The pro-

cess is shown in Figure 1. First, AbstrCT is translated and the argument component labels are automatically projected by exploiting word alignments generated by SimAlign (Sabet et al., 2020) and AWESOME (Dou and Neubig, 2021). Finally, the projected labels were manually checked to ensure trustworthy evaluation results.

3.1 AbstrCT Dataset

The dataset of Randomized Clinical Trials (RCT) consists of randomized controlled trials retrieved from the MEDLINE database via PubMed search². The trials are annotated with argument components and argumentative relations. The corpus contains paragraphs on five types of diseases: neoplasm, glaucoma, diabetes, hepatitis B, and hypertension. Overall there are 500 neoplasm, 100 glaucoma, and 100 mixed (20 from each of the 5 diseases) abstracts. The corpus includes the annotations for the argument components and relations in separate files. Table 1 provides the distribution of argument components and relations. Their definition follows previous work (Stab and Gurevych, 2017) but adapted to the medical domain.

Claims are concluding statements about the outcome of the study. In the medical domain it may be a diagnosis or a treatment.

Premises correspond to an observation or measurement in the study (ground truth), which supports or attacks another argument component, usually a claim. It is important that they are observed facts, therefore, credible without further evidence.

One document may contain multiple claims and premises that *support* or *attack* each other. Example (1) below presents an example of *claims* marked in bold with subscript C_n , and of *premises* in italics with subscript P_n .

Attack relations occur when the source component contradicts the target or when it states that some observation had no statistical significance. In Example (1), Premise₂ (P_2) (partially) attacks Claim₂ (C_2), namely, it qualifies the conditions of the study without fully objecting to it.

Support relations provide a justification from the source to the target. Example (1) contains several support links between arguments, for instance, Premise₁ (P_1) supports

²<https://pubmed.ncbi.nlm.nih.gov/7850561/>

Split	#prem	#claim	#support	#attack	#no-rel
Train - Neoplasm	1535	730	1194	200	12892
Dev - Neoplasm	218	108	185	30	1815
Test - Neoplasm	438	248	359	60	3961
Test - Glaucoma	404	190	317	29	2986
Test- Mixed	388	212	296	24	3012
Total	2983	1488	2351	343	24666

Table 1: Distribution of argument components and relations in the AbstrCT dataset.




Figure 1: The process of creating Spanish data from English.

Claim₁ (C_1) with numerical evidence to verify the statement.

(1) “Results of the pretest showed that many patients lacked knowledge about pain and pain management. The majority of pain topics had to be discussed. **[The Pain Education Program proved to be feasible]** C_1 : [75.0% of the patients had read the entire pain brochure, 55.7% had listened to the audio cassette, and 85.6% of pain scores were completed in the pain diary] P_1 . **[Results showed a significant increase in pain knowledge in patients who received the Pain Education Program and a significant decrease in pain intensity]** C_2 . [However, pain relief was mainly found in the intervention group patients without district nursing] P_2 . **[It can be concluded that the tailored Pain Education Program is effective for cancer patients in chronic pain]** C_3 .”

3.2 Machine Translation

Four MT systems were tested to translate the AbstrCT dataset into Spanish. Before proceeding to translate the full dataset we performed a pilot by translating a small set of sentences using m2m-100 (Fan et al., 2021), mBART (Tang et al., 2020), OPUS-MT (Tiedemann, Thottingal, and others,

2020) and DeepL³. In order to obtain a ranking of the MT systems, the output of each model was then blindly evaluated by two native Spanish-speaking experts in medical NLP. DeepL was, overall, the best-performing system, and OPUS-MT was the second-best-performing. The kappa inter-annotator agreement (IIA) in choosing the best MT system was around $\sim 70\%$.

A noticeable issue in translating scientific and medical data was the difficulty in dealing with abbreviations. Thus, the ability of MT systems to deal with them was one of the criteria in the selection of the most accurate model, besides linguistic coherence and correctness. Once the MT system (DeepL) was chosen, the next step consisted of projecting the argument component labels (*claims* and *premises*) from the original English annotated data to the automatically translated Spanish data. Note that label projection is only challenging for argument components as they consist of spans of texts that need to be aligned from the source to the target language (as illustrated in Figure 2), while *support* and *attack* relation classification is modeled as pairwise classification in which there is a 1-1 relationship which means that there is no need to perform label projection.

DC-treated patients experienced improved pain relief compared with VC (P=0.033) , whereas pain relief with DCb and VC was similar .

Los pacientes tratados con DC experimentaron un mayor alivio del dolor en comparación con la CV (P=0,033) , mientras que el alivio del dolor con DCb y CV fue similar .

Los pacientes tratados con DC experimentaron un mayor alivio del dolor en comparación con la CV (P=0,033) , mientras que el alivio del dolor con DCb y CV fue similar .

Figure 2: Alignment issues to project two *premise* argument component labels from English (top) to Spanish with AWESOME (middle) and SimAlign (bottom).

3.3 Label Projection

Following previous work on crosslingual transfer (Yarowsky et al., 2001; Agerri et al., 2018) we project the argument component labels via word alignments. In this work we use state-of-the-art methods for word alignment based on Transformer-based multilingual contextualized embeddings, namely, *SimAlign* (Sabet et al., 2020) and AWESOME (Dou and Neubig, 2021). Once you obtain the word alignments then we need to project the sequences of full argument components, an extra step required which is not covered by the word alignment algorithms. Instead of coming up with a new method, we apply the automatic annotation projection algorithm, Easy Label Projection, developed by García-Ferrero, Agerri, and Rigau (2022). However, as Easy Label Projection was developed for tasks such as NER or Opinion Target Extraction which are characterized by shorter and homogeneous sequences, we had to implement several changes to adapt it to the projection of long and heterogeneous argumentative sequences such as *claims* and *premises*.

While SimAlign allows projecting labels without any parallel data by extracting alignments from similarity matrices of multilingual embeddings, AWESOME requires parallel data, so we decided to use the English-Spanish parallel biomedical corpus.⁴

We generated three versions of AbstrCT in Spanish according to the projection method: (i) *automatic* projections, (ii) *post-processed* projections which programmatically corrected the *automatic* projections

³<https://www.deepl.com/>

⁴<https://github.com/biomedical-translation-corpora/corpora>

and, (iii) *manually* corrected projections.

3.3.1 Automatic projections

The middle alignment in Figure 2 produced by applying AWESOME and Easy Label Projection (García-Ferrero, Agerri, and Rigau, 2022) shows the difficulty for this method to correctly bracket the beginning of the *premise* sequences by including the two articles in the argument component. Another typical issue was erroneously including the final punctuation in the sentence as part of the argument component.

While SimAlign handled some articles at the beginning of argument component relatively better, it is difficult to find patterns in the misaligned cases. Thus, by looking at the last example in Figure 2, it can be seen that SimAlign correctly projected the first article, but not the ‘el’ article just after ‘mientras que’. Furthermore, the closing bracket and coma were also incorrectly projected although the final full stop was correctly left outside of the argument component.⁵

3.3.2 Post-processed projections

After identifying frequent mistakes according to some patterns, we were able to implement a simple method to automatically correct them. Furthermore, and taking into account that many sequences were full sentences, we applied a post-processing step to directly project the full sentence into Spanish, without requiring the use of word alignments and Easy Label Projection. This allowed us to automatically correct a large number of incorrectly projected sequences, as shown by #post-processed results in Tables 2 and 3.

⁵It is worth mentioning that in the original English AbstrCT punctuation is not consistently annotated.

	train_awesome	train_simalign	dev_awesome	dev_simalign
# total	4405	4405	680	680
# full O's	2345	2345	377	377
# full component	1752	703	257	257
# post-processed	800	88	95	11
# manual-corrections	140	194	20	25

Table 2: Train and development Neoplasm data in AbstrCT translated with DeepL. From top to down: Total number of argument components projected; number of sentences composed entirely of O tokens; number of sentences forming an argument component; number of post-processed sequences and number of manually corrected sequences.

	neoplasm_a	neoplasm_s	glaucoma_a	glaucoma_s	mixed_a	mixed_s
# total	1252	1252	1248	1248	1147	1147
# full O's	630	630	692	682	591	591
# full component	518	518	498	506	476	480
# post-processed	242	92	167	51	203	90
# manual-corrections	51	26	26	14	47	26

Table 3: Statistics of projected components for the DeepL translated test sets in AbstrCT; *_a* refers to AWESOME while *_s* corresponds to SimAlign.

3.3.3 Manual corrections

After the post-processing step, we manually checked the rest of the projected sequences to correct any possible misalignments. During this manual revision punctuation at the end of an argument component was systematically included as part of the sequence, even if this was not the case in the original English data.

The final number of manually corrected sequences is reported in the last row #manual-correction of Tables 2 and 3. As it can be seen, after post-processing the number of corrections are quite low, reducing the manual work to generate a Spanish dataset for AM in the medical domain to a minimum.

4 Experimental Setup

During the label projection described in the previous section we generated six versions of the Spanish AbstrCT dataset for argument component detection: three different projection types (automatic, post-processed and manual) with two different aligners (SimAlign and AWESOME). Furthermore, as no word alignments are required, we also directly generated one version for relation classification. In order to establish the optimal strategy to perform AM, namely, *model-transfer* or *data-transfer*, when no annotated

data is available for a target language, we run the following experiments:

- *model-transfer*: fine-tune multilingual MLMs in English and test in Spanish.
- *data-transfer*: fine-tune and test the MLMs using the AbstrCT Spanish dataset generated by MT and label projection. Evaluation is always performed on manually corrected data.
- *multilingual*: simple data augmentation using multilingual MLMs by combining the English and Spanish training data and evaluate in each of the target languages.

4.1 Models

For the experiments we use the architecture proposed in the original AbstrCT experiments (Mayer et al., 2021) although instead of using monolingual English MLMs as a backbone, for crosslingual transfer we will need to deploy a multilingual MLM. Mayer et al. (2021) experimented with adding sequence labelling components such as Conditional Random Fields (CRF), Recurrent Neural Network (RNN), or Gate Recurrent Units as a last layer on top of a Transformer model (Vaswani et al., 2017). The intuition

Model	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
SciBERT+GRU+CRF	82.41	75.84	91.11	83.97	82.89	91.68	82.40	78.21	91.35
BioBERT+GRU+CRF	80.85	73.99	90.59	83.95	83.52	91.72	82.41	78.26	92.02
BERT+GRU+CRF	82.68	76.23	89.90	82.22	79.07	89.07	82.68	77.98	89.61
mBERT+GRU+CRF	82.36	74.89	89.07	80.52	75.22	84.86	81.69	75.06	88.57
SciBert+LSTM+CRF	81.99	75.58	91.23	83.06	81.87	91.76	81.93	77.23	91.52

Table 4: Argument component detection of the source AbstrCT English data.

is that since the length of argument components is considerably long, adding a sequence labelling layer should help to better capture long dependencies between the tokens included in an argument component.

Results obtained showed that adding both a GRU and a CRF component to English MLMs specific to the medical domain such as BioBERT (Lee et al., 2019) and SciBERT (Beltagy, Lo, and Cohan, 2019) were the optimal options for argument component detection.

After the release of encoder-only MLMs for English caused a paradigm-shift in NLP research, multilingual MLMs models such as multilingual BERT and XLM-RoBERTa, trained to work on 100 languages, were published, obtaining excellent results on both monolingual and, especially, on multilingual and crosslingual settings (Pires, Schlinger, and Garrette, 2019; Wu and Dredze, 2020; Conneau et al., 2020). Therefore, as our objective is to study crosslingual transfer for AM, we decided to use a multilingual MLM as a backbone. Quite surprisingly, and contrary to previous results on Spanish and crosslingual transfer benchmarks for encoder-only MLMs (Agerri and Agirre, 2023; García-Ferrero, Agerri, and Rigau, 2022; Conneau et al., 2020), preliminary experimentation on argument component detection established that multilingual BERT (mBERT) (Devlin et al., 2019) obtained better results than XLM-RoBERTa (Conneau et al., 2020). Multilingual BERT was trained with a batch size of 256 and 512 sequence length for 1M steps, using both the MLM and Next Sentence Prediction (NSP) tasks.

After hyperparameter tuning on the English development data, the best results were obtained with 3 epochs, 5e-5 of learning rate, 32 batch size and 256 of maximum sequence length.

5 Results

Before proceeding to run the experiments on *model-* and *data-transfer*, we first tested mBERT as a backbone on the AM architecture originally designed for the English AbstrCT experiments (Mayer et al., 2021), as described above. We also fine-tuned the best English MLMs models following the method applied for English (Mayer et al., 2021). We report the results in Table 4 (F1 score is an average of F1-Claim (F1-C) and F1-Premise (F1-P) calculated at token level, as evaluated by the system used for the English AbstrCT).

F1-scores for Neoplasm show that mBERT obtains competitive results with respect to specialized monolingual models such as SciBERT and BioBERT. However, mBERT results are slightly worse when evaluated out-of-domain in Glaucoma and Mixed test data. In any case, this baseline shows that mBERT performs competitively on the English original data which makes it a good candidate to perform crosslingual and multilingual experiments.

Model	Neoplasm	Glaucoma	Mixed
BERT	66.97	57.04	69.32
SciBERT	70.31	65.75	71.31
BioBERT	55.84	59.23	56.17
mBERT	65.71	59.92	67.88

Table 5: F-1 score for relation classification on the source AbstrCT English data.

We perform the same baseline experiments for relation classification and report the results in Table 5. While the differences are a bit larger than for argument component detection in Table 4, especially with respect to SciBERT, mBERT remains competitive when compared to most monolingual English models.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
model-transfer									
DeepL + SimAlign (ES)	80.50	71.56	86.73	77.60	72.33	81.60	79.62	68.99	85.96
DeepL + Awesome (ES)	80.34	71.54	86.77	77.51	72.30	81.59	79.57	69.36	85.92
data-transfer									
DeepL + SimAlign (ES)	83.57	75.95	90.01	80.83	75.44	86.11	82.62	74.62	88.81
DeepL + Awesome (ES)	83.40	77.11	89.18	81.11	76.16	87.35	81.88	73.71	88.50
multilingual									
English	83.51	73.42	89.38	85.31	81.05	86.73	83.63	74.98	89.25
DeepL + SimAlign (ES)	84.35	76.64	88.43	84.54	78.67	87.24	83.90	73.46	88.87
DeepL + Awesome (ES)	84.58	76.77	88.61	84.62	78.71	87.23	84.03	73.99	88.92

Table 6: Results for argument component detection of *model-transfer*, *data-transfer* and *multilingual* experiments. Spanish data corresponds to manually corrected projections; in red best overall results per test set.

Crosslingual Transfer for Argument Components

Table 6 compares the performance in the different evaluation settings. The first noticeable result is that for argument component detection, *data-transfer* (training on data obtained by MT and label projection from English to Spanish) clearly outperforms *model-transfer* (training on English and predicting in Spanish). This result differs from previous works on crosslingual transfer for sequence labelling in which *model-transfer* was clearly superior (García-Ferrero, Agerri, and Rigau, 2022). We hypothesized that MT in the medical domain may introduce less translation artifacts than for other more open-ended domains.

Secondly, simple data augmentation by training on the English data together with their translated and projected counterpart in Spanish results in important performance gains. In fact, it outperforms the original English results reported in Table 4 (also against the monolingual domain-specific models) and it also helps to improve results over the monolingual Spanish *data-transfer* setting. It should also be noted that the improvements are particularly large when mBERT is evaluated out-of-domain, namely, on the Glaucoma and Mixed test sets.

Another aspect worth mentioning is that we obtain roughly the same results with both types of word alignments, which means that the process of *post-processing* and *manual correction* was robust enough to guarantee the same results regardless of the aligner used.

Summarizing, the best results for both languages are obtained with the multilingual data augmentation strategy, also surpassing the models trained with gold-standard English data. Therefore, *data-transfer* proves to be the most effective strategy to perform argument component detection in the medical domain whenever no annotated data for the target language is available.

Relation classification All the experiments that were done for argument component detection were also applied to the classification of argument relations. Results are reported in Table 7.

Here, similar to the argument components, the models trained on multilingual data also obtain the best overall results across the three test sets and *data-transfer* outperforms *model-transfer*, especially on the Glaucoma set. Furthermore, although very close, in this case, the multilingual results do not improve over the results in English obtained with gold standard data (reported in Table 5).

An important point to mention is that, as shown in Table 1, the relation classification corpus is extremely imbalanced, which may be why classification results are substantially lower. Another reason is the lack of context in the task itself, given that in many cases it is extremely difficult to distinguish relations given only two sentences, without further context. Thus, we believe that relation classification in the medical domain cannot be straightforwardly determined based only

Model	Neoplasm	Glaucoma	Mixed
train: EN+ES → test: EN	65.55	58.79	67.82
train: EN+ES → test: ES	62.55	58.60	65.74
train: EN → test: ES	62.45	55.92	65.02
train and test: ES	63.25	54.35	65.40

Table 7: F-1 scores for relation classification in crosslingual settings.

Model	Neoplasm			Glaucoma			Mixed		
	auto	post	manual	auto	post	manual	auto	post	manual
DeepL + SimAlign	75.87	79.99	80.50	75.64	77.30	77.60	75.25	79.25	79.62
DeepL + Awesome	70.07	79.12	80.34	71.39	76.77	77.51	69.85	78.21	79.57

Table 8: F1-macro for argument component detection in *model-transfer* setting; training in English and evaluating with projections (i) automatic; (ii) post-processed and (iii) manually corrected.

on local textual information. Instead, it may require more complex structures and additional insights from the data.

6 Discussion

As an additional analysis, we compare the results when evaluating *model-transfer* on the three types of projected data generated: (i) automatically projected labels, (ii) post-processed projection and (iii) manually checked labels. Table 8 reports the results. The most interesting outcome is that, while evaluating on automatic projections substantially degrades the final performance, the results obtained by evaluating on the post-processed data are practically the same as when evaluating on manually checked projections. This means that *data-transfer* can be performed fully automatically while obtaining reliable evaluation results, which is an extra argument in favour of this method to perform AM in the medical domain.

Overall, we can see how the results obtained by mBERT in this setting are high enough and they get better with each improvement introduced in the evaluation sets. The main issue with the automatic projections was that the alignment boundaries were different to those in the English gold training data (many errors were committed by not including a Spanish article as part of the argument), therefore the prediction scores, as expected, are noisy and unreliable.

Comparing the results of training in gold standard English and evaluating on the dif-

ferent types of projected data in Table 8 allowed us to learn that the *post-processed* version is slightly less reliable than testing on the manually-checked projections. Still, differences were minimal. A natural follow-up experiment is to compare the *data-transfer* results when training with the different projection versions and evaluating on the manually-checked data. Results of this experiment are reported in Table 9.⁶

First, it is interesting to observe that while SimAlign performs best when trained on automatic projections, AWESOME greatly benefits from the post-processing step, allowing it to outperform SigAlign in both post-processed and manual evaluation settings. This is coherent with the stats reported in Table 2 in which post-processing helped to greatly improve AWESOME automatic projections.

In any case, the most noteworthy aspect to be mentioned is that training on post-processed DeepL + AWESOME projections obtains practically the same results than when training on manually-checked data. This result is quite interesting as it confirms that *data-transfer* can be fully performed without any manual intervention. As this is the optimal approach in crosslingual transfer for AM, these results suggest that as long as we can obtain quality machine translation and alignments then *data-transfer* should be the choice

⁶Appendix A provides all the detailed results of performing this experiment with different MT systems and word aligners.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
<i>Train: DeepL + SimAlign - Automatic Projections</i>									
DeepL + SimAlign	82.33	75.90	87.64	80.94	75.36	85.97	81.81	74.99	88.07
DeepL + Awesome	81.98	75.27	86.99	80.84	75.31	85.96	81.73	75.36	88.03
<i>Train: DeepL + Awesome - Automatic Projections</i>									
DeepL + SimAlign	72.78	72.66	87.03	74.21	72.82	84.74	72.40	70.17	87.99
DeepL + Awesome	72.50	72.62	87.00	74.20	72.77	84.75	72.43	70.60	87.95
<i>Train: DeepL + SimAlign - Post-processed Projections</i>									
DeepL + SimAlign	82.98	75.79	88.92	81.85	75.68	87.55	82.21	73.04	89.26
DeepL + Awesome	82.77	75.50	88.89	81.74	75.63	87.53	82.14	73.46	89.22
<i>Train: DeepL + Awesome - Post-processed Projections</i>									
DeepL + SimAlign	83.09	75.19	89.07	81.43	75.58	87.0	82.31	73.69	89.11
DeepL + Awesome	83.04	75.29	89.05	81.32	75.53	86.98	82.23	74.09	89.07
<i>Train: DeepL + SimAlign - Manually-checked Projections</i>									
DeepL + SimAlign	83.57	75.95	90.01	80.83	75.44	86.11	82.62	74.62	88.81
DeepL + Awesome	83.49	76.27	90.05	80.73	75.43	86.09	82.56	75.13	88.81
<i>Train: DeepL + Awesome - Manually-checked Projections</i>									
DeepL + SimAlign	83.49	76.84	89.09	81.22	76.20	87.37	81.96	73.32	88.54
DeepL + Awesome	83.39	77.11	89.18	81.11	76.16	87.36	81.88	73.71	88.50

Table 9: Data-transfer results with different types of projected data for training. Evaluation tests are manually corrected.

to perform AM in the medical domain when no manually annotated data is readily available.

7 Conclusion

In this paper, we investigate crosslingual transfer techniques to perform AM on medical data for a language for which manual annotations are not available. Furthermore, we explore this in a real case scenario in which the only existing dataset annotated for the medical domain was in English. We compare *model-transfer* (apply multilingual MLMs such as mBERT to learn on the available English data and predict in Spanish) with *data-transfer*, which consists of generating training data in Spanish by machine translating and projecting the annotations from English to Spanish.

Contrary to previous results on crosslingual transfer for sequence labelling (García-Ferrero, Agerri, and Rigau, 2022), *data-transfer* clearly is the optimal option for argument component detection. Furthermore, combining the generated Spanish data with the original English set helps to improve results across every evaluation setting, also

with respect to the domain-specific models trained with gold-standard English data. Finally, a quick comparison evaluating with the different types of label projections demonstrates that *data-transfer* can be fully performed without manual intervention, a very important result of this study.

Results in relation classification were similar to those of argument component detection except that the multilingual setting does not outperform the results obtained when training on English gold data.

Finally, apart from the scientific findings, we should stress that in this work we have created the first dataset in Spanish to perform Argument Mining in the medical domain. The dataset, fine-tuned models and code will be publicly available upon publication. Additionally, we would like to further explore the method presented in this project to experiment with computational approaches to argumentation for other specific domains and languages for which no annotated data is available.

Acknowledgments

This work has been supported by the following MCIN/AEI/10.13039/501100011033 projects: (i) Antidote (PCI2020-120717-2) and EU NextGenerationEU/PRTR (ii) DeepKnowledge (PID2021-127777OB-C21) and by FEDER, EU; (iii) DeepMinor (CNS2023-144375) and European Union NextGenerationEU/PRTR. Anar Yeginbergen’s PhD contract is part of the PRE2022-105620 grant, financed by MCIN/AEI/10.13039/501100011033 and by the FSE+.

References

- Accuosto, P., M. Neves, and H. Saggion. 2021. Argumentation mining in scientific literature: From computational linguistics to biomedicine. In *Frommholz I, Mayr P, Cabanac G, Verberne S, editors. BIR 2021: 11th International Workshop on Bibliometric-enhanced Information Retrieval; 2021 Apr 1; Lucca, Italy. Aachen: CEUR; 2021. p. 20-36*, pages 20–36. CEUR Workshop Proceedings.
- Agerri, R. and E. Agirre. 2023. Lessons learned from the evaluation of Spanish Language Models. *Proces. del Leng. Natural*, 70:157–170.
- Agerri, R., Y. Chung, I. Aldabe, N. Aranberri, G. Labaka, and G. Rigau. 2018. Building named entity recognition taggers via parallel corpora. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Alamri, A. and M. Stevenson. 2016. A corpus of potentially contradictory research claims from cardiovascular research abstracts. *Journal of biomedical semantics*, 7(1):1–9.
- Artetxe, M., S. Ruder, and D. Yogatama. 2020. On the cross-lingual transferability of monolingual representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637. Association for Computational Linguistics.
- Beltagy, I., K. Lo, and A. Cohan. 2019. Scibert: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.
- Chen, X., A. H. Awadallah, H. Hassan, W. Wang, and C. Cardie. 2019. Multi-source cross-lingual model transfer: Learning what to share. In A. Korhonen, D. Traum, and L. Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3098–3112. Association for Computational Linguistics.
- Conneau, A., K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *ACL*.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*, abs/1810.04805.
- Dou, Z.-Y. and G. Neubig. 2021. Word alignment by fine-tuning embeddings on parallel corpora. *arXiv preprint arXiv:2101.08231*.
- Eger, S., J. Daxenberger, C. Stab, and I. Gurevych. 2018. 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. Association for Computational Linguistics.
- Fan, A., S. Bhosale, H. Schwenk, Z. Ma, A. El-Kishky, S. Goyal, M. Baines, O. Celebi, G. Wenzek, V. Chaudhary, et al. 2021. Beyond english-centric multilingual machine translation. *Journal of Machine Learning Research*, 22(107):1–48.
- Gaddy, D. M., Y. Zhang, R. Barzilay, and T. S. Jaakkola. 2016. Ten pairs to tag: multilingual pos tagging via coarse mapping between embeddings. Association for Computational Linguistics.
- García-Ferrero, I., R. Agerri, and G. Rigau. 2022. Model and data transfer for cross-lingual sequence labelling in zero-resource settings. In *In Findings of EMNLP*.
- Green, N., E. Cabrio, S. Villata, and A. Wyner. 2014. Argumentation for scientific claims in a biomedical research article. In *ArgNLP*, pages 21–25.

- Lee, J., W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang. 2019. Biobert: A pre-trained biomedical language representation model for biomedical text mining. *bioinformatics*, btz682.
- Lewis, P., B. Oguz, R. Rinott, S. Riedel, and H. Schwenk. 2020. MLQA: Evaluating cross-lingual extractive question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330. Association for Computational Linguistics.
- Li, M., S. Geng, Y. Gao, S. Peng, H. Liu, and H. Wang. 2017. Crowdsourcing argumentation structures in chinese hotel reviews. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 87–92. IEEE.
- Liu, Z., G. I. Winata, S. Cahyawijaya, A. Madotto, Z. Lin, and P. Fung. 2020. On the importance of word order information in cross-lingual sequence labeling. In *AAAI Conference on Artificial Intelligence*.
- Mayer, T., S. Marro, E. Cabrio, and S. Villata. 2021. Enhancing evidence-based medicine with natural language argumentative analysis of clinical trials. *Artificial Intelligence in Medicine*, 118:102098.
- Mochales, R. and A. Ieven. 2009. Creating an argumentation corpus: do theories apply to real arguments? a case study on the legal argumentation of the echr. In *Proceedings of the 12th international conference on artificial intelligence and law*, pages 21–30.
- Peldszus, A. and M. Stede. 2013. From argument diagrams to argumentation mining in texts: A survey. *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)*, 7(1):1–31.
- Pires, T. J. P., E. Schlinger, and D. Garrette. 2019. How Multilingual is Multilingual BERT? In *ACL*.
- Sabet, M. J., P. Dufter, F. Yvon, and H. Schütze. 2020. Simalign: High quality word alignments without parallel training data using static and contextualized embeddings. *arXiv preprint arXiv:2004.08728*.
- Sackett, D. L., W. M. C. Rosenberg, J. A. M. Gray, R. B. Haynes, and W. S. Richardson. 1996. Evidence based medicine: what it is and what it isn't. *BMJ*, 312:71 – 72.
- Shankar, R. D., S. W. Tu, and M. A. Musen. 2006. Medical arguments in an automated health care system. In *AAAI Spring Symposium: Argumentation for Consumers of Healthcare*, pages 96–104.
- Sousa, A., B. Leite, G. Rocha, and H. L. Cardoso. 2021. Cross-lingual annotation projection for argument mining in portuguese. In *Portuguese Conference on Artificial Intelligence*.
- Stab, C. and I. Gurevych. 2014. Annotating argument components and relations in persuasive essays. In *Proceedings of COLING 2014, the 25th international conference on computational linguistics: Technical papers*, pages 1501–1510.
- Stab, C. and I. Gurevych. 2017. Parsing argumentation structures in persuasive essays. *Computational Linguistics*, 43(3):619–659, September.
- Tang, Y., C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan. 2020. Multilingual translation with extensible multilingual pre-training and finetuning. *arXiv preprint arXiv:2008.00401*.
- Tiedemann, J., S. Thottingal, et al. 2020. Opus-mt—building open translation services for the world. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*. European Association for Machine Translation.
- Toulmin, S. E. 1958. *The uses of argument*. Cambridge university press.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Wu, S. and M. Dredze. 2020. Are All Languages Created Equal in Multilingual BERT? In *Workshop on Representation Learning for NLP*.

Yang, Z., R. Salakhutdinov, and W. W. Cohen. 2017. Transfer learning for sequence tagging with hierarchical recurrent networks. *arXiv preprint arXiv:1703.06345*.

Yarowsky, D., N. Grace, W. Richard, et al. 2001. Inducing multilingual text analysis tools via robust projection across aligned corpora. In *Proceedings of the First International Conference on Human Language Technology Research*, pages 1–8.

A *Results comparison with different types of projections and MT systems*

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
DeepL + SimAlign	75.87	71.21	85.70	75.64	72.08	81.18	75.25	68.84	85.08
DeepL + Awesome	70.07	70.05	84.83	71.39	71.26	80.29	69.85	67.37	84.46
OPUS-MT + SimAlign	77.59	71.09	86.24	75.46	69.86	80.37	76.63	70.26	85.50
OPUS-MT + Awesome	71.36	69.74	84.97	71.22	69.70	79.59	77.98	68.94	84.70

Table 10: Model-transfer English to Spanish results of argument components evaluated on automatic projection labels.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
DeepL + SimAlign	79.99	71.69	86.67	77.30	72.29	81.51	79.25	69.36	85.85
DeepL + Awesome	79.12	71.33	86.48	76.77	72.08	81.32	78.21	68.95	85.65
OPUS-MT + SimAlign	80.68	71.54	86.69	76.81	70.08	80.59	79.89	70.84	85.91
OPUS-MT + Awesome	80.21	71.13	86.48	76.42	69.92	80.62	79.23	70.59	85.78

Table 11: Model-transfer English to Spanish results of argument components evaluated on post-processed projection labels.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
<i>Train: OPUS-MT + SimAlign</i>									
DeepL + SimAlign	82.90	75.34	87.51	83.37	78.52	86.80	82.81	75.86	87.96
DeepL + Awesome	82.70	75.55	87.22	83.26	78.48	86.79	82.65	75.87	87.92
OPUS-MT + SimAlign	82.63	74.12	87.84	83.29	78.42	86.25	82.46	74.64	88.22
OPUS-MT + Awesome	82.60	74.27	87.90	83.13	78.39	86.17	82.29	74.51	88.09
<i>Train: OPUS-MT + Awesome</i>									
DeepL + SimAlign	73.32	73.67	86.87	75.36	73.44	85.11	74.28	73.56	87.64
DeepL + Awesome	73.15	73.74	86.73	75.32	73.40	85.11	74.26	73.99	87.63
OPUS-MT + SimAlign	74.51	73.11	87.21	75.97	73.87	84.85	74.71	73.22	87.96
OPUS-MT + Awesome	74.41	72.91	87.19	75.89	73.85	84.79	74.68	73.49	87.82
<i>Train: DeepL + SimAlign</i>									
DeepL + SimAlign	82.33	75.90	87.64	80.94	75.36	85.97	81.81	74.99	88.07
DeepL + Awesome	81.98	75.27	86.99	80.84	75.31	85.96	81.73	75.36	88.03
OPUS-MT + SimAlign	81.27	74.33	87.74	80.20	74.40	85.69	80.20	71.98	87.44
OPUS-MT + Awesome	80.99	74.15	87.71	80.04	74.37	85.61	80.11	72.24	87.30
<i>Train: DeepL + Awesome</i>									
DeepL + SimAlign	72.78	72.66	87.03	74.21	72.82	84.74	72.40	70.17	87.99
DeepL + Awesome	72.50	72.62	87.00	74.20	72.77	84.75	72.43	70.60	87.95
OPUS-MT + SimAlign	73.74	71.53	86.85	74.18	72.89	84.14	73.25	70.49	87.58
OPUS-MT + Awesome	73.53	71.28	86.84	74.13	72.86	84.09	73.22	70.76	87.45

Table 12: Data-transfer results with automatically projected data. Evaluation tests are manually corrected.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
<i>Train: OPUS-MT + SimAlign</i>									
DeepL + SimAlign	83.35	74.49	87.59	82.29	75.99	86.10	83.09	74.94	87.87
DeepL + Awesome	83.07	74.81	87.02	82.19	75.95	86.09	83.02	75.31	87.83
OPUS-MT + SimAlign	83.00	73.46	88.18	81.84	76.27	85.84	82.32	72.64	88.81
OPUS-MT + Awesome	83.08	73.99	88.34	81.72	76.25	85.77	82.19	72.68	88.68
<i>Train: OPUS-MT + Awesome</i>									
DeepL + SimAlign	83.50	75.72	88.15	82.69	74.99	87.03	82.43	72.85	88.31
DeepL + Awesome	83.38	75.80	88.11	82.58	74.95	87.01	82.26	72.85	88.28
OPUS-MT + SimAlign	83.68	75.03	88.64	82.94	77.41	87.68	83.21	73.55	89.59
OPUS-MT + Awesome	83.71	75.34	88.76	82.78	77.39	87.61	83.03	73.39	89.45
<i>Train: DeepL + SimAlign</i>									
DeepL + SimAlign	82.98	75.79	88.92	81.85	75.68	87.55	82.21	73.04	89.26
DeepL + Awesome	82.77	75.50	88.89	81.74	75.63	87.53	82.14	73.46	89.22
OPUS-MT + SimAlign	81.14	73.66	87.80	80.74	74.74	87.13	81.03	72.38	88.32
OPUS-MT + Awesome	81.07	73.91	87.86	80.58	74.72	87.06	80.94	72.66	88.19
<i>Train: DeepL + Awesome</i>									
DeepL + SimAlign	83.09	75.19	89.07	81.43	75.58	87.0	82.31	73.69	89.11
DeepL + Awesome	83.04	75.29	89.05	81.32	75.53	86.98	82.23	74.09	89.07
OPUS-MT + SimAlign	81.88	73.29	88.86	80.14	74.84	85.88	81.06	73.39	88.55
OPUS-MT + Awesome	81.74	73.04	88.86	79.99	74.82	85.82	80.97	73.65	88.42

Table 13: Data-transfer results with post-processed projected labels. Evaluation tests are manually corrected.

Test	Neoplasm			Glaucoma			Mixed		
	F1	F1-C	F1-P	F1	F1-C	F1-P	F1	F1-C	F1-P
<i>Train: OPUS-MT + SimAlign</i>									
DeepL + SimAlign	83.79	75.81	88.95	81.80	74.78	86.57	82.60	74.08	88.73
DeepL + Awesome	83.50	75.90	88.61	81.69	74.73	86.55	82.53	74.46	88.69
OPUS-MT + SimAlign	83.03	74.68	88.69	82.06	75.77	87.36	82.64	72.94	89.31
OPUS-MT + Awesome	82.80	74.58	88.69	81.90	75.73	87.28	82.51	72.94	89.18
<i>Train: OPUS-MT + Awesome</i>									
DeepL + SimAlign	83.51	75.45	88.72	82.38	76.01	87.02	82.76	74.81	88.86
DeepL + Awesome	83.34	75.64	88.72	82.27	75.97	86.99	82.69	75.18	88.83
OPUS-MT + SimAlign	82.87	74.24	88.71	82.60	76.59	87.53	82.64	73.44	89.32
OPUS-MT + Awesome	82.66	74.07	88.69	82.44	76.55	87.45	82.55	73.70	89.18
<i>Train: DeepL + SimAlign</i>									
DeepL + SimAlign	83.57	75.95	90.01	80.83	75.44	86.11	82.62	74.62	88.81
DeepL + Awesome	83.49	76.27	90.05	80.73	75.43	86.09	82.56	75.13	88.81
OPUS-MT + SimAlign	81.82	73.98	89.49	80.14	75.70	85.75	81.06	73.71	88.67
OPUS-MT + Awesome	81.58	73.75	89.46	80.02	75.72	85.69	80.99	74.01	88.54
<i>Train: DeepL + Awesome</i>									
DeepL + SimAlign	83.49	76.84	89.09	81.22	76.20	87.37	81.96	73.32	88.54
DeepL + Awesome	83.39	77.11	89.18	81.11	76.16	87.36	81.88	73.71	88.50
OPUS-MT + SimAlign	82.02	75.12	88.97	80.47	75.23	86.77	81.27	73.16	88.88
OPUS-MT + Awesome	81.86	74.97	88.97	80.35	75.21	86.71	81.19	73.42	88.75

Table 14: Data-transfer results with manually corrected projections. Evaluation tests are manually corrected.