

Document downloaded from:

<http://hdl.handle.net/10251/184749>

This paper must be cited as:

Ruiz-Dolz, R.; Alemany-Bordera, J.; Heras, S.; García-Fornes, A. (2021). Transformer-Based Models for Automatic Identification of Argument Relations: A Cross-Domain Evaluation. *IEEE Intelligent Systems*. 36(6):62-70.
<https://doi.org/10.1109/MIS.2021.3073993>



The final publication is available at

<https://doi.org/10.1109/MIS.2021.3073993>

Copyright Institute of Electrical and Electronics Engineers

Additional Information

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Transformer-Based Models for Automatic Identification of Argument Relations: A Cross-Domain Evaluation

Ramon Ruiz-Dolz

Universitat Politècnica de València, Valencian Research Institute for Artificial Intelligence (VRAIN)

Jose Alemany

Universitat Politècnica de València, Valencian Research Institute for Artificial Intelligence (VRAIN)

Stella Heras

Universitat Politècnica de València, Valencian Research Institute for Artificial Intelligence (VRAIN)

Ana García-Fornes

Universitat Politècnica de València, Valencian Research Institute for Artificial Intelligence (VRAIN)

Abstract—Argument Mining is defined as the task of automatically identifying and extracting argumentative components (e.g., premises, claims, etc.) and detecting the existing relations among them (i.e., support, attack, rephrase, no relation). One of the main issues when approaching this problem is the lack of data, and the size of the publicly available corpora. In this work, we use the recently annotated *US2016* debate corpus. *US2016* is the largest existing argument annotated corpus, which allows exploring the benefits of the most recent advances in Natural Language Processing in a complex domain like Argument (relation) Mining. We present an exhaustive analysis of the behavior of transformer-based models (i.e., BERT, XLNET, RoBERTa, DistilBERT and ALBERT) when predicting argument relations. Finally, we evaluate the models in five different domains, with the objective of finding the less domain dependent model. We obtain a macro F1-score of 0.70 with the *US2016* evaluation corpus, and a macro F1-score of 0.61 with the *Moral Maze* cross-domain corpus.

■ **COMPUTATIONAL ARGUMENTATION** has proved to be a very solid way to approach several problems such as fake news detection [5], recommendation systems [14] or debate analysis [4]

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

among others. However, in almost every domain, it is of great importance to be able to automatically extract the arguments and their relations from the input source. Argument Mining (AM) is the Natural Language Processing (NLP) task by which this problem is addressed. The Transformer model architecture [17] and its subsequent pre-training approaches have been a turning point in the NLP research area. Thanks to its architecture, it has been possible to capture longer-

range dependencies between input structures, and thus the performance of systems developed for the most general NLP tasks (i.e., translation, text generation or language understanding) improved significantly. Therefore, the Transformer architecture has laid the foundations on which newer models and pre-training approaches have been proposed, defining the state-of-the-art in NLP. In this work, we analyze the behavior of BERT [3], XLNET [20], RoBERTa [9], DistilBERT [15] and ALBERT [6] when facing the hardest AM task: identifying relational properties between arguments.

Argument Mining was formally defined in [13] as the task that aims to automatically detect arguments, relations and their internal structure. As pointed out in [8], due to the complexity of AM, the whole task can be decomposed into three main sub-tasks depending on their argumentative complexity. First, the identification of argument components consists in distinguishing argumentative propositions from non-argumentative propositions. This allows to segment the input text into arguments, making it possible to carry out the subsequent sub-tasks. Second, the identification of clausal properties is the part of AM that focuses on finding premises or conclusions among the argumentative propositions. Third, the last sub-task is the identification of relational properties. Two different argumentative propositions are considered at a time, and the main objective is to identify which type of relation links both propositions. Different relations can be observed in argument analysis, from the classical attack/support binary analysis [2], to the identification of complex patterns of human reasoning (i.e., argumentation schemes [19]). Therefore, the identification of argumentative relations is the most complex part of AM [8], but its complexity may vary depending on how the problem is instantiated.

One of the main problems when addressing any AM task is the lack of high quality annotated data. In fact, as the argumentative complexity of the task increases, it gets harder to find large enough corpus to do experiments that match the latest NLP advances. An important feature that characterizes the transformer-based models is that large corpora are needed to achieve the performance improvement mentioned above. Re-

cently, in [18], a new argument annotated corpus of the United States 2016 debate (*US2016*) was published. This corpus contains data from the transcripts of the televised political debates and from internet debates generated around the same context. This is the first publicly available corpus with enough data to begin exploring the benefits of the most recent contributions in NLP, when applied to the identification of argumentative relations. Additionally, the *US2016* corpus has been annotated using Inference Anchoring Theory (IAT¹), a standard argument annotation guideline that provides more information than the classic attack/support binary annotation. Learning a model to automatically annotate with the use of IAT, makes it possible to evaluate its performance not only with the test samples of the corpus, but also with other different corpus already analyzed and tagged using this standard (e.g., *Moral Maze* corpus). This way, it is our objective to both: evaluate the performance of these new models in the identification of argument relations task; and to find out which one is more robust to variations in the application domain.

In this work, we explore the benefits of the most recent advances in NLP applied to relation prediction in the AM domain. For this purpose we use the recently published *US2016* corpus, since it is to the best of our knowledge, the largest annotated corpus containing information about argumentative relations, and the *Moral Maze* cross-domain corpus. Then we do: (i) a pre-processing of the corpus in order to clean and structure the data for the requirements of our experiments; (ii) an analysis of the performance of the most relevant transformer-based models (i.e., BERT, XLNET, RoBERTa, DistilBERT and ALBERT) when learning to predict the relations between argumentative propositions defined by the IAT standard; and (iii) an evaluation of the obtained models in five different domains (*Moral Maze* corpus) with the objective of analyzing the domain dependency of the transformer-based models when facing this AM task.

RELATED WORK

Argument Mining is one of the main research areas in Computational Argumentation. AM has

¹<https://typo.uni-konstanz.de/add-up/wp-content/uploads/2018/04/IAT-CI-Guidelines.pdf>

caught the attention of many researchers since it is considered to be the first step towards autonomous argumentative systems. We identified many different approaches to the Argument (relation) Mining problem, which depend on the proposed methods (i.e., Parsing algorithms, Textual Entailment Suites, Logistic Regression, Support Vector Machines and Neural Networks), and the available corpus at each moment. Initial research on automatic identification of argument relations was done in [13] where parsing algorithms were used to determine the type of relation existing between two argument propositions. Some years later, AM started to gain relevance in the NLP community. We can observe the popularization of machine learning techniques for NLP purposes in [11], [16] and [10]. Support Vector Machines (SVMs) seemed to be the best performing machine learning technique for the purpose of argument relation identification. With the advent of Neural Networks (NNs), a performance gap between previous works and this new approach could be observed. In [2] the empirical results obtained by Recurrent Neural Network (RNN) models for AM were significantly better. However, there is an interesting observation to make emphasis on, which makes it hard to compare AM works. As it can be pointed out after looking at the results depicted in works like [12] or [2], the corpus used in each work has a lot of influence in the results. This is due to many different factors such as class distributions, variable language complexity (e.g., use of irony, enthymemes, etc.) or the own size of the corpus. Therefore, misleading results may be observed if the generalization of the model is not properly evaluated. On the other hand, deep learning algorithms require much more data to significantly increase its performance compared to classic neural, machine learning or statistical methods. Therefore, from all these past years of argument relation identification works, the performance has been improved not only with the use of new models or techniques, but also with the creation of better corpora. In [18], a new argument annotated corpus (*US2016*) was published, with enough data to begin exploiting the benefits of the most recent advances in NLP (i.e., transformer-based models) in the AM domain. To the best of

our knowledge, this is the first work addressing the Argument (relation) Mining problem using Transformer-based models in more than a unique domain.

DATA

Two different corpora have been used in this work: the *US2016* debate corpus and the *Moral Maze* multi-domain corpus. Both corpora can be downloaded from the Argument Interchange Format Database (AIFdb), an initiative of researchers from the ARG-tech² with the objective of creating a standard formatted argument corpus database [7]. This database contains 193 different argumentative corpora structured using the AIF standard. Each corpus is divided into several argument maps (**Figure 1**), and each argument map contains a set of Argument Discourse Units (ADUs) with its argumentative relations annotated using the Inference Anchoring Theory (IAT). This annotation method considers the most important three argumentative relations: inference (RA), conflict (CA) and rephrase (MA). An inference relation between two propositions determines that one is used to support or justify the other; a conflict relation indicates that two propositions have contradictory information; and a rephrase between two propositions means that they are equivalent from an argumentative point of view.

In order to adapt the AIFdb corpus to the needs of this task, we did some pre-processing. Each argument map is stored in a JSON file, and represented as a graph following the AIF standard. We generate a unique tab-separated values file per corpus containing three different values: proposition1, proposition2 and label. In addition to the existing IAT relation labels, we decided to generate an additional relation: the no relation (NO) label. Since most of the pairs of propositions found in a debate are not related, we decided to generate a 65% of samples belonging to this new class. For this purpose, we mixed up the propositions that were not annotated with any of the IAT relation classes. This way, the resulting model will also be able to discriminate between related or not related propositions.

²<https://www.arg-tech.org/>




Figure 1. US2016 Argument Map Sample. ADUs are bounded by rectangles. Relation types are contained in the rhombuses.

Table 1. Class distribution of the US2016 corpus, Train and Test partitions.

	US2016	Train	Test
RA	2744	2195	549
CA	888	710	178
MA	705	564	141
NO	8055	6444	1611
Total	12392	9913	2479

Table 2. Multi-domain evaluation corpus (Moral Maze) class distribution.

	MM2012	B	E	M	P	W
RA	833	128	121	205	192	187
CA	200	26	36	30	45	63
MA	156	3	25	48	41	39
NO	2209	292	339	526	517	537
Total	3398	449	521	810	795	826

US2016 Debate Corpus

The *US2016*³ corpus is an argument annotated corpus of the electoral debate carried out in 2016 in the United States. It contains both, transcriptions of the different rounds of TV debate, and discussions from the Reddit forums as detailed in [18]. The class distribution of the processed *US2016* corpus is depicted in **Table 1**. Since it is the largest publicly available argument annotated corpus in the literature, we used it to train the models. We decided to split the corpus with the 80% of the proposition pairs for training, and the remaining 20% for the evaluation.

Moral Maze Multi-Domain Corpus

The *Moral Maze*⁴ multi-domain corpus is an argument annotated corpus obtained from the transcriptions of the 2012 Moral Maze BBC discussion show. This corpus has been built from samples collected in five different broadcasts. The class distribution of the processed *Moral Maze* corpus is depicted in **Table 2**. This corpus is used to evaluate the domain robustness of the trained models across five different domain corpus: Bank (B), Empire (E), Money (M), Problem (P) and

Welfare (W); each one focused on a specific debate topic and with a different distribution of classes.

AUTOMATIC IDENTIFICATION OF RELATIONAL PROPERTIES

The problem addressed in this paper can be seen as an instance of the sentence pair classification problem. The sentence pair classification problem consists of assigning the most likely class to two text inputs at a time. In Argument Mining, after segmenting the text and defining the argument components, the argument graph must be built by identifying the relational properties between every two argument components. Therefore, given two argument components (i.e., sentences): $x_1^N = x_1, x_2, \dots, x_N$ of length N where x_n is each word of the first component; and $y_1^M = y_1, y_2, \dots, y_M$ of length M where y_m is each word of the second component, the classification problem can be modeled as defined in Equation 1,

$$\hat{c} = \arg \max_{c \in C} p(c | x_1^N, y_1^M) \quad (1)$$

where $C = [RA, CA, MA, NO]$, so the four different relation types existing in the IAT labeling are considered: inference (RA), conflict (CA),

³<http://corpora.aifdb.org/US2016>

⁴<http://corpora.aifdb.org/mm2012>

rephrase (MA) and no relation (NO). To approach this problem, we decided to use transformer-based neural architectures. The most recent works in the literature tackle the Argument Mining problem using Recurrent Neural Networks (e.g., LSTMs, BiLSTMs, etc.). However, the Transformer architecture presents several interesting improvements with respect to the RNNs. The Transformer architecture uses multiple attention modules, which allow to capture longer range dependencies between words in a sentence. Given the nature of this work’s task, we expect to have long input sentences since argumentative text is, generally, more complex than others. Therefore, we think attention mechanisms can be very useful for the identification of relational properties between argument components.

In this work, we apply Inductive Transfer Learning combined with different Transformer pre-training methods that allow us to learn our task not from scratch but using previously calculated weights. We decided to use the pre-training methods that performed the best in other NLP tasks such as Natural Language Understanding, Question Answering or Text Generation: BERT [3], XLNET [20], RoBERTa [9], DistilBERT [15] and ALBERT [6]. All these models have in common that they are based on the Transformer architecture, however different approaches have been considered in order to compute the initial weights. BERT, also known as Bidirectional Encoder Representations from Transformers, is pre-trained on masked language model and next sentence prediction tasks. The model is designed to be able to fine-tune its weights on other different tasks by adding an additional output layer. XLNet is proposed after identifying a potential problem in BERT: the language modeling of the existing dependencies between the masked positions. XLNet combines both auto-regressive language modeling and auto-encoding techniques in order to overcome the detected potential issues. RoBERTa is a strong optimization of the BERT pre-training approach. After researchers did a thorough analysis on the impact of the most important hyper-parameters, this new model was able to obtain interesting results in most of the evaluated tasks. Finally, both DistilBERT and ALBERT were proposed as smaller and faster

Table 3. Transformer-based architectures configuration.

Model	TBlocks	HSize	AH	Params.
BERT-base [3]	12	768	12	110M
BERT-large [3]	24	1024	16	340M
XLNet-base [20]	12	768	12	110M
XLNet-large [20]	24	1024	16	340M
RoBERTa-base [9]	12	768	12	125M
RoBERTa-large [9]	24	1024	16	335M
DistilBERT-base [15]	6	768	12	66M
ALBERT-base [6]	12	768	12	11M
ALBERT-xxlarge [6]	12	4096	64	223M

versions of the previous approaches. We find it interesting to also analyze and evaluate the behavior of these smaller versions, which have been designed to democratize the use of transformer-based pre-training methods without significant loss of performance.

EVALUATION

Experimental Setup

All the experiments carried out in this work have been run in a double NVIDIA Titan V computer with an Intel Xeon W-2123 CPU and 62Gb of RAM. This way, we can evaluate not only the performance of the models in the classification task, but also their training computational cost in our specific task. The number of parameters of each model is directly related to the training computational cost. **Table 3** summarises the most relevant features that define each Transformer architecture considered in this research. The Transformer blocks (TBlocks) stand for the number of layers; the hidden size (HSize) represents the number of hidden states in each layer; the attention heads (AH) indicate the number of pointers used by the attention layers; finally, the last feature is the total number of parameters (Params.) of each architecture.

In our experiments, we explore the benefits of transfer learning applied to the argument relation mining task. For that purpose, during our training phase, we use the pre-trained encoder of each model with a linear layer on its top. The output size of the linear layer coincides with the number of classes considered in our instance of the problem (i.e., 4). With the *softmax* function, we are able to model the probability of belonging to one class or another for each pair of arguments (Equation 1).

We adapted the maximum sequence length and the batch size of our inputs in each experiment. These parameters were configured in order to use the whole available GPU memory. When training BERT-base models, we defined a maximum sequence length of 256 and a batch size of 64. When training BERT-large models, we halved those values to a maximum sequence length of 128 and a batch size of 32. We trained XLNet-base with a maximum sequence length of 256 and a batch size of 32, and XLNet-large with a maximum sequence length of 256 and a batch size of 8. RoBERTa-base was trained with a maximum sequence length of 256 and batch size of 32, and for training RoBERTa-large we used the same maximum sequence length but a batch size of 16. For DistilBERT we used a maximum sequence length of 256 and a batch size of 128. Finally, ALBERT-base was trained defining a maximum sequence length of 256 and batch size of 64, but in order to fit ALBERT-xxlarge in our available memory we had to define a maximum sequence length of 128 and a batch size of 4. We trained all these models for 50 epochs in our corpus. The best results (depicted in the following section) were obtained with a $1e-5$ learning rate.

Results

In this section we present the empirical results obtained after running the experiments on all the previously defined models. In addition to the Transformer-based architectures, we have also trained a Recurrent Neural Network (RNN) as a baseline in our task. We used the best performing RNN architecture in argument relation mining proposed in [2], consisting of two Long Short-Term Memory (LSTM) networks working in parallel with each pair of arguments. We trained the baseline model for 50 epochs in our data, as the authors did in the original publication. In order to measure the performance of the different models, we have evaluated them using the macro F1-score metric. Due to the huge class imbalance in our corpora, the use of the macro F1-score makes possible to avoid misleading results during the evaluation. Additionally, we also measured the training time required by each model when learning the task proposed in this work, in order to analyze if it can be worthwhile to sacrifice their

performance in pursuit of faster training times or availability in lower resource environments.

The macro F1-scores obtained by each model are depicted in **Table 4**. In the first column, we can see every trained model. The second column represents the macro F1 obtained by each model when evaluated with the test partition of the same corpus used for training (i.e., *US2016*). The third column contains the scores obtained when the evaluation is performed on a different corpus (i.e., *MM2012*) containing a mixture of five domains. Finally, the last five columns are the macro F1-scores of the models when using each one of the five domain specific corpora (i.e., *Bank*, *Empire*, *Money*, *Problem* and *Welfare*) for evaluation.

With most of the models, we achieved state-of-the-art macro F1-scores for relation identification in Argument Mining [1]. Here, it is important to make emphasis that the way we considered to represent argumentative relations (i.e., IAT labelling) make this task harder than most of the previous work (i.e., attack/support) in this area. We obtained a 0.70 macro F1-score with *RoBERTa-large*, outperforming the LSTM baseline used as a reference of previous research in argument relation identification. Furthermore, in order to have a more strong reference to compare with previous published results, we carried out an experiment using the same parameters but considering a binary instance of the problem (i.e., only attack and support relations). This way, *RoBERTa-large* achieved a macro F1-score of 0.81 highlighting the mentioned complexity gap between the two instances of the same problem. In general, we can observe that *RoBERTa* has performed very well in this task. When looking at the cross-domain evaluation, *RoBERTa-large* has also performed the best. We obtained a 0.61 macro F1-score when doing the evaluation with a different domain corpus. Moreover, the model has been able to keep a good performance with each one of the five domain specific corpora, even having different class distributions. With *ALBERT-xxlarge-v2* it has been possible to obtain a slightly better performance when evaluating with the *Empire* corpus. It is possible to observe how the scores obtained on the *Bank* and *Empire* corpus are slightly lower than the rest. This is mainly due to their smaller size, combined with the strong imbalance between classes. We also

Table 4. Performance of the models in the automatic identification of argument relations, given in macro F1-scores.

Experiment	US2016-test	MM2012	Bank	Empire	Money	Problem	Welfare
LSTM (baseline)	.26	.24	.25	.22	.24	.25	.23
BERT-base-cased	.62	.53	.40	.45	.54	.47	.53
BERT-base-uncased	.65	.56	.42	.48	.54	.50	.54
BERT-large-cased	.61	.55	.45	.49	.53	.47	.51
BERT-large-uncased	.66	.57	.47	.49	.56	.49	.57
XLNet-base	.65	.56	.44	.49	.51	.54	.55
XLNet-large	.69	.57	.44	.51	.53	.53	.54
RoBERTa-base	.68	.58	.51	.52	.54	.52	.58
RoBERTa-large	.70	.61	.53	.53	.59	.56	.59
DistilBERT	.55	.42	.33	.39	.40	.43	.39
ALBERT-base-v2	.60	.54	.49	.45	.53	.47	.51
ALBERT-xxlarge-v2	.67	.59	.50	.54	.56	.48	.59

did experiments with cased and uncased models, in order to see the relevance of cased text in the relation identification task. As we can observe, the uncased models performed significantly better than the cased models, so we can point out that cased text did not help to improve the performance of the models in our task.

On the other hand, we obtained the worst results with *DistilBERT* and *ALBERT-base-v2*, as one might expect. We decided to use these models in order to see if the observed performance sacrifice was worthwhile in exchange for more feasible training times. Table 5 contains the training times required by each model under our experimental setup. With *DistilBERT*, it was possible to achieve a significant reduction of training time in exchange for a huge drop in performance. However, with *ALBERT-base-v2* we could not observe a significant reduction of training time. From our experiments, we have not seen any significant advantage in using these *lite* models. We also observed that the computational cost of training *XLNet-large* and *ALBERT-xxlarge-v2* in our task was very expensive. *XLNet-large* was 5.1 times slower to train than *BERT-large*. As for *ALBERT-xxlarge-v2*, the training time was 7.1 times higher than *BERT-large*. This is due to its hidden size of 4096 with respect to the 1024 sized large models. Thus, observing the performance of the models in means of their macro F1-score and the required training time, we still think that *RoBERTa* is the best approach to tackle both,

Table 5. Training time of 50 epochs running in a double NVIDIA Titan V computer.

Experiment	Time
BERT-base	39m 11s
BERT-large	2h 19m 57s
XLNet-base	1h 52m 38s
XLNet-large	11h 51m 09s
RoBERTa-base	43m 17s
RoBERTa-large	4h 44m 33s
DistilBERT	16m 15s
ALBERT-base-v2	38m 04s
ALBERT-xxlarge-v2	16h 20m 22s

Table 6. Distribution of the misclassified samples per class using the *RoBERTa-large* model. Each column indicates the real class of the samples, each row indicates the assigned class by our model.

Pred. \ Real	RA	CA	MA	NO
RA	-	0.512	0.603	0.730
CA	0.200	-	0.138	0.226
MA	0.100	0.075	-	0.044
NO	0.700	0.412	0.259	-

domain-specific and cross-domain identification of relational properties between arguments. Even the *RoBERTa-base* version performed well in this task and it was 6.6 times faster than its large version on training.

Error Analysis

In an effort to conduct a thorough evaluation, we decided to analyze the errors made

by *RoBERTa-large*, the best performing model. For this purpose we measured the volume of misclassifications found on each one of the four classes considered in this work. **Table 6** shows the error distribution detected when analyzing the results. Two important remarks can be pointed out when looking at the obtained error distributions. First of all, it is possible to observe how most of the misclassified argument pairs labeled with an inference relation were assigned the no relation class. Similarly, most of the misclassified argument pairs without relation were assigned the inference class by our model. We observed that many of these errors were due to a loss of contextual information. In an argumentative discourse, it is very common to refer to past concepts without explicitly mentioning them (i.e., enthymemes) or simplifying them with the use of pronouns. The lack of dialogical context can make the automatic identification of argument relations a harder task. For a better understanding of this problem we present the following example with two argument components labeled with inference relation:

- P1: *I think it's not going to help change the culture*
 P2: *In banking we've a totally different situation*

Our system classified the pair as no related samples. In fact, by only reading the sentence pair, one may think there is not any argument relation between them. In these situations, it is evident that the key to avoid any possible error is to give additional information about the uttered propositions. In this case, depending on the background meaning of the “*it*” and “*we*” pronouns, the sentences may be related or not. The only way of considering this proposition pair related as an inference, is assuming that the *it* pronoun refers to the banking system. We also detected that in these situations the softmax outputs of our model gives very close probabilities for both, *RA* and *NO* classes. Another indicator of the existing model misunderstandings presented before, are the similar error distributions that conflicting arguments show with both inference and no relation classes. On the other hand, we also pointed out that the rephrased argument pairs were mainly misclassified as inference related arguments. However, when analyzing them we observed that most of the relations could also be considered as inference related arguments de-

pending on the interpretation. For example:

- P1: *We do need curriculum reform*
 P2: *RUBIO too believes in curriculum reform*

In this case, the sentence pair can be interpreted as a rephrase, assuming that “*We*” and “*RUBIO too*” are equivalent subjects. But it can also be interpreted as an argument from authority, with P2 supporting (inference relation) P1. In some situations the line that differences rephrase from inference may not be as clear as desired, and both types of relation can be considered correct. Additionally, with these second type of significant detected errors, it is also possible to observe the problem mentioned before. Therefore, the loss of information caused by the use of pronouns or enthymemes in the discourse can be determinant when approaching a task of this complexity.

CONCLUSION

The automatic identification of argument relations is an essential task in the whole computational argumentation process. It allows to automatically generate the argumentative structure from argument discourse units. In this work, we present how the automatic identification of argument relations, based on Inference Anchoring Theory labeling, can be approached using the latest advances in natural language processing. To the best of our knowledge, this is the first work using transformer-based pre-trained models to learn this task. For this purpose, we have used the largest publicly available argument annotated corpus to the date. Most of the trained models have been able to outperform the state-of-the-art baselines in argument relation mining ([1]), even with a more complex instance of the task. We observed a significant better performance with *RoBERTa* than other models, the best results were achieved with *RoBERTa-large*. We also made a cross-domain evaluation of the models, in order to find out their domain robustness. Even there was a small drop in performance (most probably because of the significant variations of linguistic and class distributions between different domain corpora), the scores on different domains were still close to previous AM reports on a unique domain. This way, it is our objective to contribute on paving the way for finding models that do a better generalisation of this task. Finally, we

analyzed the errors made by our best performing model. We have seen that two important groups of errors are caused by the loss of contextual information. We also pointed out that another important group of errors made by the model was due to possible multiple interpretations of the relations. We think that significant improvements in model performance can be achieved after analyzing the most common errors detected in this work. As future work, we propose the following modifications to the automatic identification of argument relations task: (i) pronoun replacement, to solve the loss of contextual information in some propositions; (ii) consider the possible classification ambiguity, in some cases, by accepting multiple correct relations if the interpretation leads to this conclusion; and (iii) incorporation of external information. In argumentation theory, an enthymeme is known as the omission of a claim or a support of an argument. In order to make the discourse more fluid, it is very common to use enthymemes in situations where the omitted information is considered to be known by all the participants. Therefore, without external information, the model may not be able to fully understand relations between enthymemes.

ACKNOWLEDGMENT

This work is partially supported by the Spanish Government project TIN2017-89156-R, the FPI grant BES-2015-074498, and the Valencian Government project PROMETEO/2018/002. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan V GPUs used for this research.

REFERENCES

1. Oana Cocarascu, Elena Cabrio, Serena Villata, and Francesca Toni. A dataset independent set of baselines for relation prediction in argument mining. *CoRR*, abs/2003.04970, 2020.
2. Oana Cocarascu and Francesca Toni. Identifying attack and support argumentative relations using deep learning. In Martha Palmer, Rebecca Hwa, and Sebastian Riedel, editors, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1374–1379. Association for Computational Linguistics, 2017.
3. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Tamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019.
4. Ria Jha, Francesco Belardinelli, and Francesca Toni. Formal verification of debates in argumentation theory. In Chih-Cheng Hung, Tomás Cerný, Dongwan Shin, and Alessio Bechini, editors, *SAC '20: The 35th ACM/SIGAPP Symposium on Applied Computing, online event, [Brno, Czech Republic], March 30 - April 3, 2020*, pages 940–947. ACM, 2020.
5. Neema Kotonya and Francesca Toni. Gradual argumentation evaluation for stance aggregation in automated fake news detection. In Benno Stein and Henning Wachsmuth, editors, *Proceedings of the 6th Workshop on Argument Mining, ArgMining@ACL 2019, Florence, Italy, August 1, 2019*, pages 156–166. Association for Computational Linguistics, 2019.
6. Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
7. John Lawrence and Chris Reed. Aifdb corpora. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Computational Models of Argument - Proceedings of COMMA 2014, Atholl Palace Hotel, Scottish Highlands, UK, September 9-12, 2014*, volume 266 of *Frontiers in Artificial Intelligence and Applications*, pages 465–466. IOS Press, 2014.
8. John Lawrence and Chris Reed. Argument mining: A survey. *Comput. Linguistics*, 45(4):765–818, 2019.
9. Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019.
10. Stefano Menini, Elena Cabrio, Sara Tonelli, and Serena Villata. Never retreat, never retract: Argumentation analysis for political speeches. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications*

- of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018, pages 4889–4896. AAAI Press, 2018.
11. Nona Naderi and Graeme Hirst. Argumentation mining in parliamentary discourse. In Matteo Baldoni, Cristina Baroglio, Floris Bex, Floriana Grasso, Nancy Green, Mohammad-Reza Namazi-Rad, Masayuki Numao, and Merlin Teodosia Suarez, editors, *Principles and Practice of Multi-Agent Systems - International Workshops: IWECC 2014, Gold Coast, QLD, Australia, December 1-5, 2014, and CMNA XV and IWECC 2015, Bertinoro, Italy, October 26, 2015, Revised Selected Papers*, volume 9935 of *Lecture Notes in Computer Science*, pages 16–25. Springer, 2015.
 12. Vlad Niculae, Joonsuk Park, and Claire Cardie. Argument mining with structured svms and rnns. In Regina Barzilay and Min-Yen Kan, editors, *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 985–995. Association for Computational Linguistics, 2017.
 13. Raquel Mochales Palau and Marie-Francine Moens. Argumentation mining: the detection, classification and structure of arguments in text. In *The 12th International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 8-12, 2009, Barcelona, Spain*, pages 98–107. ACM, 2009.
 14. Antonio Rago, Oana Cocarascu, and Francesca Toni. Argumentation-based recommendations: Fantastic explanations and how to find them. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 1949–1955. ijcai.org, 2018.
 15. Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108, 2019.
 16. Christian Stab and Iryna Gurevych. Parsing argumentation structures in persuasive essays. *Comput. Linguistics*, 43(3):619–659, 2017.
 17. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
 18. Jacky Visser, Barbara Konat, Rory Duthie, Marcin Koszowy, Katarzyna Budzynska, and Chris Reed. Argumentation in the 2016 US presidential elections: annotated corpora of television debates and social media reaction. *Lang. Resour. Evaluation*, 54(1):123–154, 2020.
 19. Douglas Walton, Chris Reed, and Fabrizio Macagno. *Argumentation Schemes*. Cambridge University Press, 2008.
 20. Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764, 2019.
- Ramon Ruiz-Dolz** received the Master’s Degree in Artificial Intelligence, Pattern Recognition and Digital Imaging from the Universitat Politècnica de València (UPV) in 2019. He currently works as a researcher at the Valencian Research Institute for Artificial Intelligence (VRAIN) and is pursuing a Ph.D. degree in Computer Science. His research interests are focused on Argument Mining, Computational Argumentation and Persuasion Technologies. Contact him at the Universitat Politècnica de València, 46022, Valencia, Spain, raruidol@dsic.upv.es.
- Jose Alemany** received his PhD degree in Computer Science from the Universitat Politècnica de València (UPV) in 2020. He is currently working as a Postdoctoral Researcher at the Valencian Research Institute for Artificial Intelligence (VRAIN). He is also working as an Assistant Professor at Florida Universitària. His research interests include information dissemination, privacy-preserving, content analysis, and complex networks. Contact him at the Universitat Politècnica de València, 46022, Valencia, Spain, jalemany1@dsic.upv.es.
- Stella M. Heras Barberá** holds a PhD in Computer Science (extraordinary prize Cum Laude) from the Polytechnic University of Valencia (UPV, 2011). She obtained the Executive Master in Project Management by the University of Valencia (2013) (certified PMP-1558995) and has the title of University Special-

ist in University Pedagogy (UPV, 2011). She currently works as a researcher at the Valencian Research Institute for Artificial Intelligence (VRAIN) and as an associate professor in the Department of Languages and Computer Systems at the UPV, where she has been teaching since 2007. Her research area is focused on the development of artificial intelligence systems (computational argumentation, persuasion technologies, educational recommender systems). Contact her at the Universitat Politècnica de València, 46022, Valencia, Spain, stehebar@upv.es.

Ana García-Fornes received her PhD degree in Computer Science from the Universitat Politècnica de València (UPV, 1996). She currently works as a Full Professor at the Department of Information Systems and Computation at the UPV, Spain and as a researcher at the Valencian Research Institute for Artificial Intelligence (VRAIN). Her research interests focus on real-time systems, multi-agent systems, agreement technologies and privacy in social media. Contact her at the Universitat Politècnica de València, 46022, Valencia, Spain, agarcia@dsic.upv.es.