

Perturbations and Subpopulations for Testing Robustness in Token-Based Argument Unit Recognition

Jonathan Kamp¹ Lisa Beinborn¹ Antske Fokkens^{1,2}

¹Computational Linguistics and Text Mining Lab, Vrije Universiteit Amsterdam

²Dept. of Mathematics and Computer Science, Eindhoven University of Technology

{j.b.kamp, l.beinborn, antske.fokkens}@vu.nl

Abstract

Argument Unit Recognition and Classification aims at identifying argument units from text and classifying them as *pro* or *against*. One of the design choices that need to be made when developing systems for this task is what the unit of classification should be: segments of tokens or full sentences. Previous research suggests that fine-tuning language models on the token-level yields more robust results for classifying sentences compared to training on sentences directly. We reproduce the study that originally made this claim and further investigate what exactly token-based systems learned better compared to sentence-based ones. We develop systematic tests for analysing the behavioural differences between the token-based and the sentence-based system. Our results show that token-based models are generally more robust than sentence-based models both on manually perturbed examples and on specific subpopulations of the data.

1 Introduction

Identifying argumentation units is difficult, both for humans and machines. The challenge starts with the question of what it means for a segment to be argumentative towards a given topic in the first place (Trautmann et al., 2020; Habernal et al., 2014, e.g.). Trautmann et al. (2020) propose a pragmatic approach for defining arguments and ask annotators to identify segments that can be placed in the <argument span> slot of the following template: “<TOPIC> should be supported/opposed, because <argument span>”. They compare models that are trained to label tokens as being part of argumentative segments to models that classify full sentences as containing argumentative segments (ARG) or not (non-ARG) (see Figure 1), ultimately arguing that token-based training is preferable. Their experiments suggest that a token-based approach is more robust when sentence boundaries are unknown or not precisely given.

Argumentative segments provide reasons for taking a positive (pro) or negative (against) stance on a topic. These arguments are highly topic-specific, but the decent accuracy of cross-topic models indicates that there are also topic-independent cues. Niven and Kao (2019) previously showed that transformer-based models learn to map specific cue words to a label and learn little about argumentation reasoning. It could be that this is the most we can expect in a cross-topic scenario. The question then remains where these cues are found: are they in the ARG segments themselves or are they also provided by the non-ARG context? When comparing token-based and sentence-based models, we expect token-based models to be better at picking up cues that are specific to ARG segments themselves, whereas sentence-based models may be more susceptible to cues from the non-ARG context, in particular, when these appear to announce an argumentation (e.g. *because I think that...*). Reliance on (non-ARG) cues is a particularly strong signal that general cues rather than reasoning are used.

In this paper, we dive further into this line of research. We rerun experiments with the best models of Trautmann et al. (2020) to ensure a fair basis of comparison, reproducing most of the original results and coming close for the rest. We then design multiple robustness tests comparing the behavior of token- and sentence-based models in mixed-segment sentences, i.e. sentences that contain at least one ARG segment and one non-ARG segment. We expect token-based models to be more robust because they are trained to distinguish between ARG and non-ARG segments within sentence boundaries and thus have access to more precise information as to what makes up an ARG segment during training. This hypothesis is confirmed by our perturbation tests, which also show different behavior on subpopulations of the data. We thus show that a relatively small, curated dataset of ad-

segments	non-ARG									ARG						non-ARG	
seq. labeling	non-ARG	non-ARG	non-ARG	non-ARG	non-ARG	non-ARG	non-ARG	non-ARG	non-ARG	ARG	ARG	ARG	ARG	ARG	ARG	ARG	non-ARG
seq. classification	ARG																
	People	who	support	gun	control	claim	or	pretend	that	gun	control	laws	significantly	reduce	violent	crime	.

Figure 1: An example sentence from the AURC-8 (Trautmann et al., 2020) topic *gun control*. Each sentence in the dataset has one vector of token-wise gold labels (in- and output in a sequence *labeling* approach, i.e. **token-based**) as well as one sentence-wise gold label (in- and output in a sequence *classification* approach, i.e. **sentence-based**). ARG and non-ARG gold *segments* are sequences of tokens that carry the same label.

versarial examples can provide systematic insights into model behavior. Additional robustness tests with subpopulations of the data surprisingly do not yield clear differences between the two approaches.

2 Background and Related Work

In this section, we first present related work on argument unit recognition (§2.1) and then dive further into the concept of robustness tests (§2.2).

2.1 Argument Unit Recognition

Argumentation theory is about identifying how humans reach common ground and compromise, how societal information is exchanged, what the degree of subjectivity in viewpoints is and how polarised different stances can be. In the digital era, arguments from a wide range of sources are analysed. These sources range from debates on social media and (online) fora to technical documents used by professionals in the legal domain. Arguments roughly reflect the rationale behind a stance or decision, in relation to a certain topic or proposition. The field of computational argumentation attempts to model the argument patterns that are present in human language. Lauscher et al. (2021) distinguish between different tasks in argument modeling: \sim mining, \sim assessment, \sim reasoning, and \sim generation. Argument Unit Recognition and Classification is a task that can be positioned within argument *mining*, as argumentative from non-argumentative expressions are first distinguished, and a stance is then attributed to each of the identified arguments. Ajjour et al. (2017) show that the task of segmenting a text into argument units of different types remains particularly challenging in a cross-domain setting.

The first part of this research aims to reproduce the Argument Unit Recognition and Classification experiments by Trautmann et al. (2020), who train multiple transformer-based models on a novel argumentation dataset that is labeled at the token

level: spans of tokens are then predicted as being pro, against or non-argumentative towards a given topic.

Argument mining has been thoroughly approached by (Bi-)LSTM modeling (Eger et al., 2017), SVMs and RNNs (Niculae et al., 2017). Apart from Trautmann et al. (2020), however, transformer-based architectures have been deployed more rarely. Poudyal et al. (2020) show how RoBERTa (Liu et al., 2019) can successfully be applied on the legal ECHR dataset on a claim-premise task. Ruiz-Dolz et al. (2021) test several flavors of BERT models (Devlin et al., 2019) on the same task, but on a less domain-specific debate corpus. Mayer et al. (2020a) compare different domain-generic and -specific transformer-based models in combination with CRF and GRU layers on medical texts. Similarly to Trautmann et al. (2020), they experiment with both sequence labeling and sequence classification, applying the former to a component detection task, and using the latter to classify relations between argument components.

The next subsection provides background and related work on the second contribution of our paper: testing robustness.

2.2 Robustness Testing

Goel et al. (2021) describe three ways of testing robustness: (1) testing on **subpopulations** of the test data the model is expected to perform poorly on; (2) **perturbing** the test data by creating adversarial examples (Zhang et al., 2020) that are expected to shed light on weaknesses of the model; (3) assessing model performance on pre-existing evaluation sets to establish **scalability** and **cross-domain** validity. We briefly discuss the role of each of these three in our work.

We design three **subpopulation** tests, two of which are based on similarity between training and test instances and one based on the ratio of argumentative tokens in a sentence. We also design

three **perturbation** tests. The first design choice involves the level of granularity, which is usually on the word- or phrase-level. In our case, perturbation units are aligned with the granularity of the annotated spans, i.e. the ARG or non-ARG segments remain intact but are combined in different ways. A second point of attention in creating perturbations is the risk of altering the grammar or semantics in an unintended way. Automatic metrics have been utilised to determine whether linguistic aspects are preserved after perturbation, such as the Jaccard similarity coefficient, grammar and syntax related measurements and edit-based measurements (Zhang et al., 2020). These may be relatively fast to use, especially on a large scale, but might fall short in tasks where generating adversarial candidates goes beyond relatively simple, single word substitutions. We therefore opt for manual verification of our samples. Trautmann et al. (2020) include the third type of robustness test already in that they apply **cross-topic** evaluation. Since we are mostly interested in the models’ generic ability in identifying argumentative segments, we apply our robustness tests in the cross-topic setting only.

A handful of studies have applied robustness tests to transformer-based models on an argument mining task. Schiller et al. (2021) apply paraphrases, spelling alterations and negation stress tests on a stance detection task. Niven and Kao (2019) apply a negation stress test on a Argument Reasoning Comprehension Task, where negating a warrant (i.e. a type of argument) should result in predicting the inverted label. Mayer et al. (2020b) protract robustness testing into adversarial training: by inserting or replacing simple linguistic elements in the original data, such as nouns, scalar adverbs and punctuation, they use the perturbed examples for retraining the model, achieving higher performance. Mayer et al. (2020b) show the effectiveness of single, token-level perturbations, while aiming to control for *same-meaning* preservation between the original and perturbed example pairs. Instead, we focus on the different argumentative load that different parts of a sentence carry to guide our perturbations, and ensure that the result of each perturbation is semantically sound (yet not unaltered). Finally, the in-domain versus cross-domain comparison is a more frequent type of testing, but it is often approached from a generalisability perspective (*how well does the model perform on cross-domain data?*), which has a slightly different connotation

from a robustness perspective (*how well does the model defend itself from specific adversaries in cross-domain data?*). Our work can be seen as an extension to Trautmann et al. (2020), who find that token-based models are more robust against sentence segmentation errors than sentence-based models. Our robustness tests go beyond their work in that they show that token-based models are also more robust compared to sentence-based models in well-formed sentences with manipulated combinations of ARG and non-ARG segments. Furthermore, the phenomena that we are testing robustness on are more likely to occur than scenarios in which sentence boundaries are not given.

3 Reproduction Experiments

This section describes our reproduction study, including the dataset used (§3.1), the experimental setup (§3.2), the model evaluation metrics (§3.3), and the requirements for a successful reproduction together with our results (§3.4).

3.1 Dataset Description

The AURC-8 dataset developed by Trautmann et al. (2020) is divided over eight topics: 1. *abortion* 2. *cloning* 3. *marijuana legalization* 4. *minimum wage* 5. *nuclear energy* 6. *death penalty* 7. *gun control* 8. *school uniforms*. In their manual labeling process, annotators were presented with candidate sentences in which arguments related to one given topic were possibly present. Argument spans were annotated according to the slot-filling template “<TOPIC> *should be supported/opposed, because <argument span>*”. This results in spans annotated as *PRO* (a supporting argument) or *CON* (an opposing argument). Spans that remain unlabeled are assigned *NON* (a non-argumentative segment). As an example, both underlined spans in the following sentence about *death penalty* are labeled as *CON* segments: ‘*It does not deter crime and it is extremely expensive to administer .*’ Instead, the first underlined span in the following sentence about *gun control* is labeled as a *CON* segment whereas the second span is labeled as *PRO*: ‘*Yes , guns can be used for protection but laws are meant to protect us , too .*’ In both example sentences, the spans of adjacent non-underlined tokens form the *NON* segments. The dataset consists of 1,000 example sentences per topic. Of the 8,000 total sentences, 3,500 (43.75%) are annotated as ARG and 4,500 (56.25%) as non-ARG. The por-

tion of ARG sentences is divided over 658 examples (14.62%) containing exclusively *PRO* segments, 621 examples (13.80%) containing exclusively *CON* segments and 3,221 (71.58%) containing any combination of *PRO*, *CON* and *NON* segments.¹

The models are run on two different splits of the data: in-domain and cross-domain. In the in-domain setup, the first 70% of the examples from each of the Topics 1-6 is assigned to training, the next 10% to the development set, and the last 20% to the test set. The cross-domain setup assigns all sentences from Topics 1-5 to training, Topic 6 to development, whereas Topic 7 and 8 form the test set.²

3.2 Experimental Setup

We use two training approaches: token-based and sentence-based.

Token-based Models are trained on the sequence of token-wise gold labels, in a sequence-labeling fashion. The input to the model are tokenised sentences.

Sentence-based Models are trained on a sentence-level gold label, in a sequence-classification fashion. The sentence-level gold label is a modification of the token-level gold labels. Let t_L be the set of labels assigned to individual tokens in a sentence, and f_{PRO} and f_{CON} the number of tokens in the sentence that are labeled as *PRO* and *CON*, respectively. Then, the sentence label s_L is obtained as follows:

$$\begin{aligned} t_L &= \{NON\}, & s_L &:= NON \\ t_L &= \{NON, PRO\}, & s_L &:= PRO \\ t_L &= \{NON, CON\}, & s_L &:= CON \\ t_L &\supseteq \{PRO, CON\}: \\ &\text{if } f_{PRO} > f_{CON}, & s_L &:= PRO \\ &\text{if } f_{CON} > f_{PRO}, & s_L &:= CON \\ &\text{if } f_{PRO} = f_{CON}, & s_L &:= random^3 \end{aligned}$$

The input instance fed to the sentence-based model is the same tokenised sentence used as input in the token-based model. Instead of feeding along a sequence of token-wise labels, we feed its unique s_L . The output is a predicted s_L .

¹In this calculation, *NON* segments that are solely formed by punctuation marks are ignored.

²Visit Appendix A for additional details on dataset versioning and pre-processing of the data.

³A random choice from $\{PRO, CON\}$ is made.

We re-train the models based on the architecture that performed best in the original paper: BERT_{LARGE} (Devlin et al., 2019). We also train a token-based model with a CRF layer.⁴ In the original results, the CRF layer improved segmentation. The model without CRF more often broke segments up into multiple single-word segments.

For each domain split, for each model setup we carry out series of 5 training runs with a different random seed for each run. We report mean F1-scores and standard deviation for each series of runs. Hyperparameter settings are reported in Appendix A.

3.3 Model Evaluation

The models are evaluated on two metrics: token-F1 and sentence-F1.⁵ Token-F1 is calculated as the average over the per-class F1-scores for all tokens in the evaluation set. Sentence-F1 is the average over per-class F1-scores for all sentences in the evaluation set. Whereas token-F1 is straightforward for the token-based setup, and sentence-F1 is for the sentence-based setup, one extra step is needed to retrieve the sentence labels from the token-based predictions, and token labels from the sentence-based predictions. When using the token-based model, we obtain the sentence labels from the assigned tokens using the same approach as described in §3.2. After applying the sentence-based model, we obtain token labels by assigning the predicted sentence label to all tokens of the sentence.

3.4 Reproduction Results and Considerations

We consider the mean F1-scores over three runs from Trautmann et al. (2020) as the benchmark for the reproduction comparisons. We follow Moore and Rayson (2018) and provide F1-distributions reporting the mean and standard deviation from our experiments. It remains a methodological challenge to determine a threshold within which a score can be defined as successfully reproduced. We follow Reuver et al. (2021) and consider the reproduction successful if given a distribution of reproduced F1-scores D , the original mean F1-score falls within two standard deviations from the mean

⁴We were not able to re-implement the CRF layer for the sentence-based approach and could therefore not include this.

⁵Trautmann et al. (2020) also include a third metric: segment-F1. Given that the description of their implementation remains underspecified and since the metric is not strictly relevant to our work, we report the original segment-F1 results in Appendix A along with our own segment-F1 implementation.

		token-F1		sentence-F1		
setting	model	token -based setup	sentence -based setup	token -based setup	sentence -based setup	
in-domain	orig	BERT _{LARGE}	.683	.627	.709	.715
		BERT _{LARGE} +CRF	.696	.622	.711	.725
	repr	BERT _{LARGE}	.698 (.003)	.614 (.008)	.708 (.004)	.713 (.012)
		BERT _{LARGE} +CRF	.696 (.003)	-	.711 (.006)	-
cross-domain	orig	BERT _{LARGE}	.596	.544	.598	.602
		BERT _{LARGE} +CRF	.620	.519	.610	.573
	repr	BERT _{LARGE}	.587 (.008)	.529 (.011)	.604 (.009)	.566 (.017)
		BERT _{LARGE} +CRF	.578 (.008)	-	.609 (.007)	-

Table 1: Original results (white background) compared to reproduction results (non-white background) on the test set. Models are divided over an in- and cross-domain setting. Reproduction results show the mean scores from 5 runs, along with the standard deviation (within parentheses). The reproduction scores where the original score falls within two standard deviations from the mean are given in bold.

of D . We provide all individual decisions on the test set for a more accurate comparison, since F1-scores can still stem from different behavior on sub-populations of the data.

At a first glance, the differences between the original and replicated results are relatively small for both token-based and sentence-based models. One pattern from the original paper is not reproduced, namely, the positive effect on performance by the CRF layer on the token-based model. In the light of the threshold of two standard deviations, we observe in Table 1 that reproductions are partially successful. On the test set, 5 out of 6 F1-scores are reproduced in the in-domain setting, and 4 out of 6 in the cross-domain setting. Success rate of reproduction does not seem to depend on training setup either: 6 out of 8 for token-based versus 3 out of 4 for sentence-based. Scores that are not reproduced come close as they fall within half a decimal from the original.⁶

4 Robustness Testing

We test robustness in a cross-domain setting. By isolating this problem from topic-dependent content biases, the models are expected to focus more on indicators that are representative of a generic notion of argumentation. While a token-based model is explicitly instructed that there are fine-grained argumentative differences within a sen-

⁶See Table 5 in Appendix A for a complete overview of the reproduction results. It can be observed that none of the segment-F1 metrics are reproduced, probably caused by a slightly different implementation of how these scores are calculated.

tence, a sentence-based model is not. Therefore, we expect the sentence-based models to have more difficulty in predicting the cues that are argumentative on a micro-level (i.e. tokens, segments), which translates to difficulties at the macro-level (i.e. the sentence). Our robustness tests precisely operate at a micro-level: adding, replacing or removing segments should impact the sentence-based model more negatively than the token-based model.

We apply robustness tests to the two cross-domain token-based models (BERT_{LARGE}, BERT_{LARGE}+CRF) and the sentence-based model (BERT_{LARGE}). We investigate robustness for the task as a binary prediction problem (Argument Unit Recognition) and remove the stance component: ARG entails both labels *PRO* and *CON*, and *non-ARG* corresponds to *NON*. As anticipated in §2.2, we categorise the robustness tests according to two classes: *perturbations* on the test set (§4.1) and *subpopulations* of the test set (§4.2).

4.1 Perturbations on the Test Set

We craft a *before*-dataset and *after*-dataset in the following way. First, artificial candidate test sets are generated through deletion, recombination or label-based pre-selection of segments. The segments are sampled from the original test set. Second, we manually label or complete the candidate examples. We create three types of tests **T1**, **T2** and **T3**. We report on the impact of the perturbation through Δacc , i.e. the difference between the accuracy *before* and the accuracy *after* the perturbation has been applied. Hence, each example in

either the *before*- or *after*-dataset has one gold label (at the sentence-level) on which the models are evaluated.

T1 - Announcing Segments Observations in the original test set show that non-ARG segments can broadly be divided in segments that announce (ANN) an immediately subsequent ARG segment, and segments that do not (non-ANN). For instance, ANN segments are phrases that include literal argument indicators such as *evidence*, *claim*, *argument*, *reason* followed by a copula, and phrases that include reporting verbs. Examples:

ANN

... a major argument against this topic is. . .
 ... he thinks that. . .

non-ANN

... this document was written in 2022 and. . .
 ... but. . .

ANN segments are an example of information that is known to be non-ARG by the token-based model, but not by the sentence-based model where it falls under a coarse-grained, sentence-level ARG label. Since ANN segments mostly co-occur with ARG segments, the sentence-based model is likely to mix them up. The token-based model may also use an ANN segment as signal that an ARG is following, but has better chances of using information from the following segment itself to identify when this is not the case. We test this by creating counter-examples that concatenate ANN segments to a subsequent *non-ARG* segment. This results in non-ARG sentence-level labels, for instance, ‘*Pro-abortion politicians think that...*’ + ‘*...the debate has become very delicate.*’. If our theory is correct, the token-based model would generally be able to classify the two segments separately as non-ARG, resulting in a non-ARG label for the sentence, whereas the sentence-model is more prone to label the sentence as ARG based on the ANN segment.

	concatenation	sentence gold
<i>before</i>	ANN <u>non-ARG</u> seg. + ARG seg.	ARG
<i>after</i>	ANN <u>non-ARG</u> seg. + <u>non-ARG</u> seg.	<u>non-ARG</u>

We first extracted candidate $\langle a, b \rangle$ pairs, where a is an ARG segment, b is a non-ARG segment and a is immediately followed by b in the same sentence from the original AURC-8 dataset. Pairs that do not form a full sentence are manually discarded. Subsequently, we manually labeled the non-ARG segments as (non-)ANN, until reaching 100 ANN annotations for the *gun control* topic and 100 for *school uniforms*. Each ANN segment (e.g. ‘*Pro-abortion politicians think that...*’) is then manually completed with a novel non-ARG segment (e.g. ‘*...the debate has become very delicate.*’) to form a full non-ARG sentence and is added to the *after*-dataset. The respective $\langle a, b \rangle$ pairs are added to the *before*-dataset.

T2 - Concatenate Non-Argumentative Sentence

Here we test robustness by concatenating an ARG segment with a pure non-ARG sentence. In between the two segments, the connector ‘*and besides,*’ is used to create a well-formed sentence. This results in constructions where the ARG segment ends up in a context of a relatively high number of non-ARG tokens. Such a concatenation would result in e.g.: ‘*Uniforms force conformity*’ + ‘*and besides,*’ + ‘*it’s a great service for parents as I was able to pick up lots of good stuff for little money.*’ The token-based model is expected to classify the two segments as ARG and non-ARG respectively, resulting in a sentence-wise ARG label prediction. The sentence-based model might be more biased by the high ratio of non-ARG tokens that are present in the sentence, potentially resulting in a sentence-wise non-ARG prediction.

	concatenation	sentence gold
<i>before</i>	ARG seg.	ARG
<i>after</i>	ARG seg. + connector + <u>non-ARG</u> sent.	ARG

We populate a candidate dataset with concatenations of an ARG segment, the connector and a pure non-ARG sentence, in that order. The components in each concatenation (except for the connector, which is constant) are on the same topic and are sampled from the original test set. The ARG segment should be a full stand-alone sentence. From this candidate dataset, we then select 50 examples

for *gun control* and 50 for *school uniforms* that are sound, stand-alone sentences to be added to the *after*-dataset. The *before*-dataset consists of the respective ARG segments.

T3 - Remove Non-Argumentative Segment In this test, we remove the non-ARG context around the remaining ARG segment creating uncontextualised argument units. We expect this perturbation to have less impact on the token-based model, as its decision is potentially less informed by the missing non-ARG segments.

	concatenation	sentence gold
<i>before</i>	ARG seg.	
	+ <u>non-ARG</u> seg.	
	/	ARG
<i>after</i>	<u>non-ARG</u> seg.	
	+ ARG seg.	
	ARG seg.	ARG

We extract pairs consisting of an ARG and a non-ARG segment from the original corpus. Both elements of each pair stem from the same source sentence and are originally adjacent. We manually check them to ensure that both the pair and the ARG segment alone form well-formed sentences. We select a total of 200 examples with an approximate 50%-50% split of examples where non-ARG precedes or follows the ARG segment, as well as an approximate 50%-50% split between the two topics. The pairs form the *before*-dataset whereas the ARG segments alone form the *after*-dataset.

4.2 Subpopulations of the Test Set

A subpopulation is a group of test instances that is selected based upon a criterion that is expected to influence the performance of the model. We take it a step further: we consider each instance in the test set a subpopulation on its own and assign it a value from a continuous variable in the data. In our case, the continuous variable is a semantic similarity score between train and test data, and the ratio of noisy (non-argumentative) tokens per sentence, two aspects that generally impact language classification tasks. The point-biserial correlation coefficient r_{pb} is then calculated between this continuous variable and the dichotomous prediction correctness. Thus, r_{pb} is expected to be lower for

a model when the continuous variable forms less of a bias on its decisions compared to its effect on another model.

T4 - Similarity Train-Test Same Labels The outcome of this test provides an indication of the impact of semantic similarity between training and test data on the decision of the model. For each of the mixed-segment sentences in the test set, a pairwise semantic similarity coefficient is calculated in relation to each of the sentences in the training set. If the maximum semantic similarity coefficient for one test sentence corresponds to a training instance with the same label (ARG), the test sentence is stored in the *T4-set* along with its coefficient. The correlation between prediction correctness of mixed-segment sentences from the *T4-set* and their respective maximum similarity coefficients is then computed. We expect the sentence-based model to be more affected by it than the token-based model, given that semantic similarity at the macro-level of the sentence may be a more prominent indicator for the former model. A token-based setup, on the other hand, should be able to classify segments within the sentence as it can rely on explicit ARG vs non-ARG information. This translates to a correlation coefficient that is expected to be higher for the sentence-based model than for the token-based model.

The semantic vector representation of a sentence is given by its averaged token vectors.⁷ Sentence similarity corresponds to the cosine similarity between the two semantic vector representations of a sentence pair.

T5 - Similarity Train-Test Opposite Labels For T5, the maximum similarity coefficient is calculated in relation to the instances in the training set that have an opposite label (non-ARG) to the mixed-segment sentences. Similarly to the *T4-set*, a *T5-set* is created accordingly. This aspect is expected to have more impact on the sentence-based model than the token-based model, hence yielding a weaker correlation for the latter.

T6 - Argumentative Token Ratio Through T6, prediction correctness is correlated with the argumentative token ratio in mixed-segment sentences from the test set. This ratio is calculated as the number of ARG tokens over the number of all tokens. In line with the expectations in T4 and T5,

⁷spacy.io/models/en → en_core_web_lg v3.3.0.

the token-based model should be less affected by this sentence-level aspect, resulting in a weaker correlation coefficient compared to the sentence-based model.

4.3 Results Robustness Tests

The perturbation results of T1, T2, T3 are collected in Table 2, where Δacc quantifies the impact of each perturbation. Specifically, Δacc represents the difference in accuracy by the models on sentences before and after the perturbation has been applied. It can be observed that an overall negative Δacc pattern is present across the grid, which is expected behavior. The maximum absolute negative impact is $\Delta acc = -.077$, achieved through T3 on the token-based model without CRF layer. From a relative point of view, the sentence-based model is impacted most with -11.7% on T3.

As an answer to our initial expectations, the token-based model with CRF layer is more robust to perturbations than the sentence-based model on two out of three tests: T1 (Announcing Segments) and T3 (Remove Non-Argumentative Segment). This is quantified in terms of both absolute Δacc (-.022 on T1, -.068 on T3) and relative Δacc (-2.5% on T1, -9.2% on T3). In comparison, the token-based model without CRF is impacted more heavily than the sentence-based model in absolute terms (-.043 versus -.027 on T1; -.077 versus -.076 on T3), but is more robust in relative terms on T3 (-10.1% versus -11.7%). Although the CRF layer has not proven to clearly increase the token-based model performance (not observable in Table 1, nor in Table 2), it appears to improve the robustness of the model.

Interestingly, the token-based model without CRF layer is the only one to considerably improve performance on the T2 *after*-dataset. This behavior is unexpected since all *after*-sets were meant to *trick* the models rather than to help them. A possible explanation might be that the connector ‘*and besides,*’ is often included in the annotated ARG spans in AURC-8 training instances. This could represent a general downside of token-based models: picking up a small cue in the sentence as ARG, therefore predicting the sentence-wise label as ARG.

The results of subpopulation tests T4, T5 and T6 are given in Table 3. We hypothesised that continuous aspects in the data (such as semantic similarity between full sentences in training and test or the

argument token ratio of a sentence) would correlate more strongly with predictions by the sentence-based model compared to the token-based model. This hypothesis could not be confirmed. Apart from being close to 0, which indicates no correlation, the *rpb* coefficients for the token-based model are also close to the coefficients for the sentence-based model on the same tests T4-6, which indicates no difference in bias between the models. The perturbation results (T1-3), however, provide an indication that there are differences between subpopulations. Specifically, both token-based models achieve a higher accuracy on each single *before*- and *after*-dataset, which are specific subpopulations of the data. This clear difference in performance can be explained by the fact that these tests do not cover pure, non-argumentative sentences on which the sentence-based model might be stronger (as can be inferred from the comparable sentence-F1 scores between the two types of models in Table 1). We therefore believe more research on subpopulations is needed. In particular, we may investigate alternative implementations of the continuous variables, such as using Sentence-bert (Reimers and Gurevych, 2019) for representing the semantics of individual instances or also looking at the number of semantically similar examples in the training data.

5 Conclusion

In this study, we partially reproduced the results of Trautmann et al. (2020) and introduced new robustness tests that showed how token-based models are generally more robust than models trained at a sentence level on an Argument Unit Recognition task. We applied two type of tests: perturbations and subpopulations. With regards to the perturbations, we found that 1) removing the non-ARG segment from a mixed-segment sentence, and 2) replacing the ARG segment with a non-ARG segment in *announcing* phrases such as ‘*Their main argument is <ARG>*’ or ‘*Most politicians against gun legislation think that <ARG>*’ negatively impact a sentence-based model more than a token-based model. We did not find a difference in bias among the two types of models with regards to semantic similarity between training and evaluation data, and high argumentative token ratios at the sentence level. Instead, we showed that the development of perturbation test sets itself can shed light on specific subpopulations of the data: our token-based

<i>model</i>	T1			T2			T3		
	<i>before</i>	<i>after</i>	Δacc	<i>before</i>	<i>after</i>	Δacc	<i>before</i>	<i>after</i>	Δacc
token-based BERT _{LARGE}	.875 (.018)	.832 (.033)	-0.043 -4.8%	.760 (.035)	.830 (.037)	+0.070 +9.2%	.760 (.022)	.683 (.037)	-0.077 -10.1%
token-based BERT _{LARGE} +CRF	.878 (.021)	.856 (.024)	-0.022 -2.5%	.790 (.015)	.760 (.028)	-0.030 -3.8%	.740 (.036)	.672 (.023)	-0.068 -9.2%
sentence-based BERT _{LARGE}	.835 (.034)	.808 (.053)	-0.027 -3.2%	.638 (.059)	.640 (.064)	+0.002 +0.3%	.652 (.047)	.576 (.034)	-0.076 -11.7%

Table 2: Impact perturbations on cross-domain token-based and sentence-based models. Mean accuracy and standard deviation (within parentheses) over 5 runs is reported for each model. Accuracy is calculated on the test set before applying the perturbation (*before*) and after applying the perturbation (*after*). Δacc represents the absolute and relative (%) difference between *before* and *after*.

<i>model</i>	T4	T5	T6
	r_{pb}	r_{pb}	r_{pb}
token-based BERT _{LARGE}	-.068 (.031)	.027 (.009)	.028 (.023)
token-based BERT _{LARGE} +CRF	-.031 (.016)	.013 (.029)	.046 (.010)
sentence-based BERT _{LARGE}	-.014 (.034)	-.037 (.044)	.042 (.037)

Table 3: Impact subpopulations on cross-domain token-based and sentence-based models. r_{pb} indicates the point-biserial correlation coefficient between prediction correctness and a given aspect of the sentence. The range of r_{pb} is $[-1, 1]$, where the two extremes indicate a perfect negative and positive correlation, respectively. The coefficients in the table are the means from 5 runs per model, along with the standard deviations (within parentheses).

models performed better on both mixed-segment sentences and single argumentative segments.

By approaching the task from a challenging, cross-domain perspective, we isolated the problem from model reliance on topic-dependent content. Our analyses reveal that it is difficult to define a common denominator for the notion of argumentativeness across topics. They highlighted the importance of the type of knowledge we expect to be learned by a computational model of argumentation. Structural choices in the annotation setup can lead to systematic gaps in the dataset that allow the model to take superficial shortcuts (Gardner et al., 2020). Robustness tests are a means to detect such gaps and, as a side effect, help in unraveling conceptual vagueness.

Acknowledgements

Jonathan Kamp’s and Antske Fokkens’ research was funded by the Dutch National Science Organ-

isation (NWO) through the project InDeep: Interpreting Deep Learning Models for Text and Sound (NWA.1292.19.399). Lisa Beinborn’s research was funded by NWO through the project CLARIAH-PLUS (CP-W6-19-005). We would like to thank Dietrich Trautmann and his co-authors for making their clear code and dataset publicly available, and for the email correspondence. We also thank the anonymous reviewers for their valuable contribution.

References

- Yamen Ajjour, Wei-Fan Chen, Johannes Kiesel, Henning Wachsmuth, and Benno Stein. 2017. Unit segmentation of argumentative texts. In *Proceedings of the 4th Workshop on Argument Mining*, pages 118–128.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language under-](#)

- standing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Steffen Eger, Johannes Daxenberger, and Iryna Gurevych. 2017. **Neural end-to-end learning for computational argumentation mining**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11–22, Vancouver, Canada. Association for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, et al. 2020. Evaluating models’ local decision boundaries via contrast sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323.
- Karan Goel, Nazneen Fatema Rajani, Jesse Vig, Zachary Taschdjian, Mohit Bansal, and Christopher Ré. 2021. **Robustness gym: Unifying the NLP evaluation landscape**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 42–55, Online. Association for Computational Linguistics.
- Ivan Habernal, Judith Eckle-Kohler, and Iryna Gurevych. 2014. Argumentation mining on the web from information seeking perspective. In *ArgNLP*.
- Anne Lauscher, Henning Wachsmuth, Iryna Gurevych, and Goran Glavaš. 2021. Scientia potentia est—on the role of knowledge in computational argumentation. *arXiv preprint arXiv:2107.00281*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv e-prints*, pages arXiv–1907.
- Tobias Mayer, Elena Cabrio, and Serena Villata. 2020a. Transformer-based argument mining for healthcare applications. In *ECAI 2020*, pages 2108–2115. IOS Press.
- Tobias Mayer, Santiago Marro, Elena Cabrio, and Serena Villata. 2020b. Generating adversarial examples for topic-dependent argument classification 1. In *Computational Models of Argument*, pages 33–44. IOS Press.
- Andrew Moore and Paul Rayson. 2018. **Bringing replication and reproduction together with generalisability in NLP: Three reproduction studies for target dependent sentiment analysis**. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1132–1144, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Vlad Niculae, Joonsuk Park, and Claire Cardie. 2017. **Argument mining with structured SVMs and RNNs**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 985–995, Vancouver, Canada. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. **Probing neural network comprehension of natural language arguments**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Prakash Poudyal, Jaromir Savelka, Aagje Ieven, Marie Francine Moens, Teresa Goncalves, and Paulo Quaresma. 2020. **ECHR: Legal corpus for argument mining**. In *Proceedings of the 7th Workshop on Argument Mining*, pages 67–75, Online. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992.
- Myrthe Reuver, Suzan Verberne, Roser Morante, and Antske Fokkens. 2021. **Is stance detection topic-independent and cross-topic generalizable? - a reproduction study**. In *Proceedings of the 8th Workshop on Argument Mining*, pages 46–56, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ramon Ruiz-Dolz, Jose Alemany, Stella M Heras Barberá, and Ana García-Fornes. 2021. Transformer-based models for automatic identification of argument relations: A cross-domain evaluation. *IEEE Intelligent Systems*, 36(6):62–70.
- Benjamin Schiller, Johannes Daxenberger, and Iryna Gurevych. 2021. Stance detection benchmark: How robust is your stance detection? *KI-Künstliche Intelligenz*, 35(3):329–341.
- Dietrich Trautmann, Johannes Daxenberger, Christian Stab, Hinrich Schütze, and Iryna Gurevych. 2020. Fine-grained argument unit recognition and classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9048–9056.
- Wei Emma Zhang, Quan Z Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41.

A Appendix

Software and Hyperparameters Our code and data are available at github.com/jbkamp/repo-Rob-Token-AUR. The implementation for the token-based model was

retrieved from github.com/trtm/AURC, and we adapted it to train a sentence-based model through the `transformers.BertForSequenceClassification` class, at huggingface.co. For both, we used the large cased pre-trained model with whole word masking at huggingface.co. We used the same settings across models: learning rate was kept at $1e-5$, dropout rate at 0.1 and the maximum length of the tokenised BERT input was set at 64 tokens. Optimizer adopted: AdamW. The batch size was set at 32 and models were trained for a maximum of 100 epochs with early stopping if the performance did not improve significantly after the 10th epoch.

Dataset Versioning Trautmann et al. (2020) published their results based on the AURC-8 dataset, requested and obtained via e-mail correspondence. A second version at github.com/trtm/AURC/tree/master/data of the AURC-8 dataset was uploaded in a later moment with, *cleaner parsing and encoding* (github.com/trtm/AURC#readme, last consulted on June 16th 2022) but with the same number of labels and sentences. The two datasets differ to a low degree: 4.91% of the sentences are not equal ($n = 393$). Of this subset, all elements show a better cleaning of punctuation tokens compared to the original. To the best of our knowledge, this is the only difference between the original and the updated dataset. Therefore, we prefer using the updated, cleaner version of the dataset. We remove duplicate sentences within and across training set, development set and test set, per split (see resulting counts in Table 4).

Segment-F1 In order to compute the segment-F1 score, we average over all sentence-wise segment-F1 scores, for each sentence in the evaluation set. To obtain a sentence-wise segment-F1 score we consider all pairs $\langle y, \hat{y} \rangle$, where y is the sequence of true labels for a segment and \hat{y} is the sequence of predicted labels for that segment. Let r be the overlap ratio between y and \hat{y} :

$$r = \frac{|y \cap \hat{y}|}{|y|} \quad (1)$$

We only compute r for segments where the label of y is *PRO* or *CON*. If $r > .5$ and labels are the

same, \hat{y} is considered a true prediction; otherwise, a false prediction. The sentence-wise segment-F1 is the number of true predictions over all predictions for that sentence. If the sentence does not contain *PRO* nor *CON* segments, and no *PRO* nor *CON* is predicted, the segment-F1 score for the sentence is 1.0. See Table 5 for a full overview of the results, including token-F1, segment-F1 and sentence-F1.

#	topic	in-domain			cross-domain		
		train	dev	test	train	dev	test
1	abortion	700	99	200	800	0	0
2	cloning	696	100	200	800	0	0
3	marijuana legalization	699	100	200	800	0	0
4	minimum wage	699	100	200	800	0	0
5	nuclear energy	699	100	200	800	0	0
6	death penalty	700	100	200	0	800	0
7	gun control	0	0	0	0	0	1,000
8	school uniforms	0	0	0	0	0	1,000

Table 4: The eight topics from the AURC-8 dataset (Trautmann et al., 2020) along with the number of sentence instances per data split after duplicates removal.

model	token-F1				segment-F1				sentence-F1			
	token		sentence		token		sentence		token		sentence	
	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup	-based setup
	dev	test	dev	test	dev	test	dev	test	dev	test	dev	test
in-domain BERT _{LARGE}	.732	.683	.671	.627	.749	.709	.599	.567	.738	.709	.759	.715
in-domain BERT _{LARGE} +CRF	.743	.696	.637	.622	.750	.724	.552	.547	.744	.711	.731	.725
in-domain BERT _{LARGE}	.717	.698	.628	.614	.776	.749	.514	.500	.715	.708	.726	.713
	(.004)	(.003)	(.005)	(.008)	(.011)	(.005)	(.009)	(.004)	(.008)	(.004)	(.007)	(.012)
in-domain BERT _{LARGE} +CRF	.716	.696	-	-	.766	.743	-	-	.718	.711	-	-
	(.003)	(.003)			(.003)	(.008)			(.008)	(.006)		
cross-domain BERT _{LARGE}	.604	.596	.550	.544	.653	.626	.487	.473	.606	.598	.628	.602
cross-domain BERT _{LARGE} +CRF	.615	.620	.505	.519	.681	.649	.456	.464	.627	.610	.569	.573
cross-domain BERT _{LARGE}	.581	.587	.515	.529	.630	.603	.424	.433	.591	.604	.596	.566
	(.011)	(.008)	(.012)	(.011)	(.007)	(.011)	(.014)	(.004)	(.016)	(.009)	(.010)	(.017)
cross-domain BERT _{LARGE} +CRF	.584	.578	-	-	.627	.593	-	-	.601	.609	-	-
	(.009)	(.008)			(.012)	(.004)			(.011)	(.007)		

Table 5: Full overview of the original results (white background) compared to reproduction results (non-white background). Models are divided over an in-domain setting and a cross-domain setting. Reproduction results show the mean scores from 5 runs, along with the standard deviation (within parentheses). The reproduction scores where the original score falls within two standard deviations from the mean are given in bold.