

Resolving Legalese: A Multilingual Exploration of Negation Scope Resolution in Legal Documents

Ramona Christen^{1*}, Anastassia Shaitarova², Matthias Stürmer^{1,3}, Joel Niklaus^{1,3,4*}

¹University of Bern, ²University of Zurich, ³Bern University of Applied Sciences, ⁴Stanford University
ramona.christen@students.unibe.ch, shaita@cl.uzh.ch, {joel.niklaus, matthias.stuermer}@unibe.ch

Abstract

Resolving the scope of a negation within a sentence is a challenging NLP task. The complexity of legal texts and the lack of annotated in-domain negation corpora pose challenges for state-of-the-art (SotA) models when performing negation scope resolution on multilingual legal data. Our experiments demonstrate that models pre-trained without legal data underperform in the task of negation scope resolution. We release a new set of annotated court decisions in German, French, and Italian and use it to improve negation scope resolution in both zero-shot and multilingual settings. We achieve word-level F1-scores of up to 86.7% in our zero-shot cross-lingual experiments, where the models are trained on two languages of our legal datasets and evaluated on the third. Our multilingual experiments, where the models were trained on all available negation data and evaluated on our legal datasets, resulted in F1-scores of up to 91.1%.

Keywords: Negation Scope Resolution, Legal, Multilingual, Dataset, Transformers

1. Introduction

Negation scope resolution is an important research problem in the field of Natural Language Processing (NLP). It describes the detection of words that are affected by a negation cue (e.g. no or not) in a sentence, which is important for understanding its true meaning. Although this task is far from trivial, deep learning approaches have shown promising results (Khandelwal and Sawant, 2020; Shaitarova et al., 2020; Shaitarova and Rinaldi, 2021).

As with many NLP tasks, the largest amount of annotated data is available in English.¹ Multilingual datasets are less common and often not easily accessible. For example, on the [huggingface hub](#), hosting most important open-source datasets, 4559 datasets are tagged as English. The next most common language is Chinese with 10 times fewer datasets for a total of 469.² In addition, much of the work conducted in the area of negation scope resolution has been done in the medical domain in order to automatically process clinical reports and discharge summaries (Szarvas et al., 2008). Other datasets consist of literary texts (Morante and Blanco, 2012) or more informal data such as online reviews (Konstantinova et al., 2012). The legal domain differs from all of the above in that it is often very complex (i.e., legalese) and uses highly specific vocabulary and knowledge that is not common outside the legal domain

*Equal contribution.

¹Mielke (2016) analyzed all ACL conference proceedings from 2004, 2008, 2012, and 2016 and found that 58% to 69% of papers only evaluated in English.

²Numbers extracted from <https://huggingface.co/datasets> on 13.08.2023.

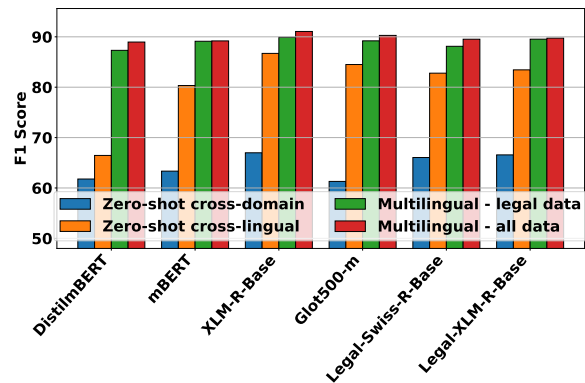


Figure 1: Main results from select models.

(Friedrich, 2021; Ruhl et al., 2017). This poses a challenge to any model tackling tasks in the legal domain. While a large amount of legal data is publicly available and has been annotated for various tasks (Chalkidis et al., 2021; Rasiah et al., 2023; Niklaus et al., 2021, 2023a; Brugger et al., 2023; Niklaus et al., 2023b; Chalkidis et al., 2022), *inter alia*, to the best of our knowledge there exists no legal negation corpus.

We annotate four new datasets containing legal judgments from Swiss and German courts in German, French and Italian for negation cues and scopes. We find that these legal documents contain on average longer sentences as well as longer annotated negation scopes, compared to existing datasets. Our experiments show that the legal domain poses a significant challenge to models attempting negation scope resolution. The results achieved by models pre-trained in different domains and evaluated on legal data are lower than those

seen in other cross-corpus experiments (Khandelwal and Sawant, 2020; Shaitarova and Rinaldi, 2021). Using our newly annotated datasets, we can improve these results. We conduct experiments where the models are fine-tuned on two languages of the legal data and evaluated on the third. In these zero-shot cross-lingual experiments, our models achieve higher F1-scores than the models pre-trained only on different domains. By training on all available data, we are able to further improve these results, achieving F1-scores around 90% for our multilingual experiments. Our results provide an interesting insight into how even smaller datasets can make a valuable contribution to improving the performance of language models (LMs) on a specific downstream task such as negation scope resolution.

Contributions

The contributions of this paper are three-fold:

- We annotate new datasets of legal documents for negation in German, French, and Italian each containing around 1000 sentences.
- We train and evaluate models on the task of negation scope resolution on the newly annotated datasets to provide a reference point and achieve word-level F1-scores in the mid eighties for cross-lingual zero-shot experiments and up to 91% in multilingual experiments.
- We publicly release the annotation guidelines, the data, the models and the experimentation code as resources and for reproducibility.³

2. Related Work

Different approaches have been used to address the issue of negation detection and negation scope resolution. Early research focused mainly on rule-based approaches. NegEx, a simple regular expression algorithm developed by Chapman et al. (2001), was successfully able to identify negations in the medical domain. Morante et al. (2008) first took a machine learning approach to negation scope resolution. They used two memory-based classifiers, one to identify the negation cue in a sentence, and one to identify the scope of the negation. On the negation scope resolution task, they achieved an F1-score of 81% on the BioScope corpus (Szarvas et al., 2008). These results were later surpassed by Fancellu et al. (2017), achieving

³The annotation guidelines as well as the code to fine-tune our models can be found on GitHub: https://github.com/RamonaChristen/Multilingual_Negation_Scope_Resolution_on_Legal_Data. Our best model (<https://huggingface.co/rcds/neg-xml-roberta-base>) and dataset (<https://huggingface.co/rcds/MultiLegalNeg>) are published on huggingface.

an F1-score of 92% by using neural networks for scope detection. Khandelwal and Sawant (2020) achieved the best results on the BioScope corpus, as well as on two other publicly available negation corpora, the SFU Review Corpus (Konstantinova et al., 2012) and the ConanDoyle-neg corpus (Morante and Blanco, 2012). Their NegBERT model uses Bidirectional Encoder Representation from Transformers (BERT) (Devlin et al., 2019) and applies a transfer learning approach for negation detection and scope resolution.

Only a limited amount of work has been conducted on negation scope resolution across different languages. Fancellu et al. (2018) developed a cross-lingual system, trained on English data and tested on a Chinese corpus. By employing cross-lingual universal dependencies in English they were able to achieve an F1-score of 72% on the Chinese data. Shaitarova et al. (2020) investigated cross-lingual zero-shot negation scope resolution between English, Spanish, and French. They built on NegBERT but used the multilingual BERT (mBERT) model. Shaitarova and Rinaldi (2021) built on this using NegBERT with mBERT and XLM-R_{Large} (Conneau et al., 2020), and were able to achieve a word-level F1-score of 87% on zero-shot transfer from Spanish to Russian.

The sparse amount of cross-lingual research can be explained by the lack of annotated data in languages other than English. There are few corpora annotated with negations in German and Italian (Jiménez-Zafra et al., 2020). The only German corpus annotated for negation and speculation contains medical data and clinical notes (Cotik et al., 2016). However, the corpus is not publicly available and no annotation guidelines have been published. For Italian, Altuna et al. (2017) presented a framework for the annotation of negations and applied it to a corpus of news articles and tweets, parts of which are publicly available. In French, Dalloux et al. (2020) annotated a medical corpus, available on request. To our knowledge, no legal corpus annotated with negations currently exists.

3. Data

3.1. Legal Data

We use court decisions in our legal datasets, also often referred to as judgments. The judgments from German courts were collected from *Bayern.Recht*⁴ and include a variety of legal domains and structures (Glaser et al., 2021). The Swiss court decisions in French, Italian, and German (CH) were collected from the Federal Supreme Court of Switzerland (FSCS). The FSCS is the highest legal authority in Switzerland and oversees federal criminal,

⁴<https://www.gesetze-bayern.de/>

administrative, patent, and cantonal courts.

Judgments published by the FSCS usually consist of four sections: 1) The introduction gives information about the date, chamber, involved judge(s) and parties, and the topic of the court decision. 2) The facts outline the important case information. 3) The considerations form the basis for the final ruling by providing relevant case law and other cited rulings. 4) The rulings gives the final decision made by the court.

3.2. Datasets

We annotated four new datasets in three languages for negation cues and scopes, and standardized the existing French and English datasets to make them more accessible. Our datasets consist of publicly available legal judgments from Swiss and German courts. Since negation scope resolution is a sentence-level task, we first split the data into sentences using sentence boundary annotations. The French (fr) and Italian (it) datasets consist of a subset of Swiss court decisions from the Swiss-Judgment-Prediction (SJP) dataset (Niklaus et al., 2022) and the Multi-Legal-Pile (Niklaus et al., 2023b) which were annotated for sentence spans by Brugger et al. (2023). The main German data (de (DE)) is a subset of judgments from German courts collected by Glaser. et al. (2021). Only judgments were included in our dataset because they include a variety of sources and legal areas, they also have a higher density of negation cues compared to other legal texts. To validate that the negation scope prediction also works on German court data from Switzerland, we curated a small dataset of German-Swiss court decisions (de (CH)) that is also a subset of the SJP corpus. We separated each dataset into a train (70%), test (20%), and validation (10%) split.

To ensure that sufficient negation data is available in each dataset, a negation score was assigned to each document based on a simple word search for the most common negation words in each language (see Appendix B). The documents with the highest negation scores were then selected to be annotated. Table 1 shows the amount of data and the distribution of negations for the newly created datasets in comparison to the existing datasets in English and French. Our datasets contain a slightly higher ratio of negated sentences compared to the other datasets. This can be attributed to the nature of legal data and our pre-selection procedure. Because we annotated only a subset of an existing corpus we were able to exclude documents without or only few negations while other corpora like ConanDoyle-neg and SFU annotated complete existing datasets or stories.

Annotations were done by native-language human annotators using the tool Prodigy. All anno-

	Dataset	Total	Negated	%neg
legal	fr	1059	382	36.07
	it	1001	418	41.76
	de (DE)	1098	454	41.35
	de (CH)	208	112	53.85
external	SFU	17672	3528	19.96
	BioScope	14700	2095	14.25
	ConanDoyle-neg	5714	1421	24.87
	Dalloux	11032	1817	16.47

Table 1: Total number of sentences, and number and percentage of sentences containing at least one negation.

tators are university students but not part of a legal study program. The annotations were cross-checked by one annotator, who has a linguistic background, with the help of an online translator to ensure that they adhere to the annotation guidelines and are consistent across all three languages. The annotation guidelines are based on existing guidelines for the English datasets, and have been extended to cover all three languages included in our data, as well as the characteristics of the legal domain. Key guidelines are summarized below.

Negation Cues Cues were not annotated as part of the negation scope following the annotation guidelines for the ConanDoyle-neg corpus (Morante et al., 2011). We excluded affixal cues⁵ in our annotations and kept all annotations to the word as the level of the minimal syntactic unit.

Multiple negations Annotators were instructed to annotate one negation per sentence. Sentences with multiple negations were duplicated before annotation based on the most common negation cues. To ensure that the same cue was not annotated twice, duplicates were displayed next to each other in the annotation tool to allow annotators to see which clues had yet to be annotated.

Maximum scope strategy As with BioScope, we used a maximum scope strategy. This means that the scope extends to the largest possible unit. If a negated clause has subordinate clauses providing additional information to the clause, the scope extends over the negated clause and all of its subordinate clauses, as illustrated in example 1. This sentence structure is very common in our set of legal data. In all following examples we mark the cue in **bold** and underline the scope. We provide an English translation for clarity.

⁵Affixal cues are cues within a word such as impossible

1 Vorliegend ginge es **nicht** darum, dass ein Arbeitgeber über Fristen oder Pflichten nicht aufgeklärt habe, somit eine blosse Untätigkeit des Arbeitgebers [...]

EN: *In the present case, it was **not** a matter of an employer not having provided information about deadlines or obligations, thus a mere inactivity on the part of the employer [...].*

Case citations Our dataset contains two main types of citations: inline citations and parenthesized citations. Inline citations, as in example 2 were annotated as part of the scope, while parenthesized citations, as in example 3, were excluded from the negation scope.

2 Da der Kläger **kein** ähnlicher leitender Angestellter i.S.d 14 Abs. 2Satz 2 KSchG ist [...]

EN: *Since the plaintiff is **not** a similar executive employee in the sense of 14 Abs. 2Satz 2 KSchG [...]*

3 Seit dem 06.02.2017 ist der Kläger im Handelsregister **nicht mehr** als Geschäftsführer eingetragen (vgl. Auszug aus dem Handelsregister in Anlage K9, Bl 75 ff. d.A).

EN: *Since 06.02.2017 the plaintiff is **no longer** registered in the commercial register as managing director (see extract from the commercial register in annex K9, Bl 75 ff. d.A)*

Punctuation Punctuation marks, such as periods or exclamation points, were excluded from the scope, unless the scope spans multiple clauses separated by commas.

Table 2 shows the average number of tokens in a sentence for all datasets, as well as the average length of the annotated scopes as a ratio between annotated and not annotated tokens. On average, the sentences in our legal datasets are slightly longer than in other datasets. Furthermore, the mean length of the annotated scopes in our data is higher than in all other datasets. For de (DE), more than 50% of tokens were annotated as scope, which is around twice as much as with the biomedical, literary, and review corpora. This is due to the legal domain’s sentence structure and our annotation guidelines, which include the subject in the scope. Additionally, nested sentences with multiple subordinate clauses are common in our dataset. This, combined with our maximum scope strategy, leads to longer scopes compared to other datasets.

	Dataset	Tokens	Scope
legal	fr	48.52	37.96%
	it	40.84	30.17%
	de (DE)	31.14	50.18%
	de (CH)	27.65	36.03%
external	SFU	24.46	21.87%
	BioScope	28.49	25.91%
	ConanDoyle-neg	22.11	32.37%
	Dalloux	25.96	19.82%

Table 2: The average number of tokens per sentence. Scopes are shown as a percentage of negated tokens.

4. Experimental setup

We performed experiments to assess negation scope resolution model performance on our multilingual legal data. We integrated the NegBERT architecture (Khandelwal and Sawant, 2020), successful in this task on prior datasets, with various pre-trained multilingual LMs outlined in Table 3. We ran each experiment five times with different random seeds and report the mean word-level F1-score averaged over random seeds, together with the standard deviation. For each token, we predict probabilities of this token belonging to either class (scope/not scope). These probabilities are averaged over all tokens belonging to one word and the class with the highest probability gives us the label for this word. We compute the F1 score over all words in the sentence with each word being one prediction. All experiments were conducted with the same hyperparameters for all models, optimized with a search over learning rate (5e-7, 1e-6, 3e-6, 1e-5, 3e-5, 5e-5) and batch size (4, 8, 16, 32, 64, and 128). We optimized the hyperparameters for mBERT and XLM-R and concluded that the best results can be achieved with an initial learning rate of 1e-5 and a batch size of 16. To avoid overfitting, we used early stopping with patience set to 8 as a compromise between the patience of 6 used in the original NegBERT experiments (Khandelwal and Sawant, 2020) and 9 used in the multilingual experiments of Shaitarova and Rinaldi (2021). We extended the maximum input length to 252 tokens for our data. Experiments ran on an NVIDIA A100 GPU via Google Colab, totaling around 105 hours of training time.

First, we evaluated ChatGPT in zero- and few-shot experiments for negation scope resolution. For all subsequent experiments, we used the NegBERT architecture. In the first NegBERT experiment, models were fine-tuned on all existing French and English datasets and evaluated on our new legal datasets, representing a Zero-shot cross-domain

Model	Source	InLen	Params	Vocab	NumTokens	Corpus	Langs
DistilmBERT	Sanh et al. (2020)	512	134M	120K	n/a	Wikipedia	104
mBERT	Devlin et al. (2019)	512	177K	120K	n/a	Wikipedia	104
XLM-R _{Base/Large}	Conneau et al. (2020)	512	278M/560M	250K	6'291B	2.5TB CC100	100
Glott500-m	ImaniGooghari et al. (2023)	512	395M	401K	94B	glot500-c	511
Legal-Swiss-R _{Base/Large}	Rasiah et al. (2023)	512	184M/435M	128K	262B/131B	CH Rulings/Legislation	3
Legal-XLM-R _{Base/Large}	Niklaus et al. (2023b)	512	184M/435M	128K	262B/131B	CH Rulings/Legislation	3

Table 3: Model stats. InLen: max input length during pre-training. Params: total parameter count. NumTokens: Batch size \times Steps \times InLen

transfer. For a second series of zero-shot experiments, we attempted a [Zero-shot cross-lingual transfer](#) within our legal datasets. In each cross-lingual experiment, models were trained on two dataset languages and evaluated on the third. We also executed [Multilingual experiments](#) using our datasets and all available data.

5. Results

ChatGPT We evaluated the performance of ChatGPT-3.5 (Brown et al., 2020), one of the leading LMs, in the task of negation scope resolution on our legal datasets. Other researchers have found that ChatGPT performs well on simple annotation tasks such as text classification (Gilardi et al., 2023). To analyze ChatGPT’s understanding of negation scopes, we conducted a small test over the chat interface (see Appendix A) which showed that it was able to correctly identify the negation scope of a simple German sentence. For the same request with an example sentence from our legal dataset, ChatGPT was not able to accurately identify the negation scope. To evaluate the performance on the whole dataset, we used the ChatGPT API with ‘gpt-3.5-turbo-16k’ to accommodate longer inputs. We set the temperature to 0 to reduce randomness and receive a coherent output in json format. Similar to the experiments with the NegBERT architecture, we gave the sentence as well as the negation cues as input and prompted ChatGPT to return the sentence annotated for negation scopes. In a zero-shot experiment, we did not give any annotated examples and only provided a short definition of negation scopes. The results show that ChatGPT’s performance on our datasets is subpar (Table 4). In an effort to increase the performance, we conducted some few-shot experiments where 1, 5 or 10 examples of annotated sentences were provided with the prompt, but it did not lead to improvement. The results of the 1-shot experiments averaged lower than the 0-shot experiments. Overall the standard deviation is very high which can be explained by the fact that a random set of annotated examples was selected for each of the five experiment runs. Overall we can conclude, that ChatGPT is currently not suited to solve negation

scope resolution in the legal domain without fine-tuning.

Zero-shot cross-domain transfer The results for our zero-shot cross-domain transfer experiments are presented in Table 5. The best results over all datasets were achieved by the Legal-XLM-R_{Large} model, scoring an F1-score of 71.6%. Overall, the LMs pre-trained on legal data demonstrated a 4-percentage point advantage, with a mean F1 of 68.3% averaged over all four legal models, compared to the other models pre-trained on different domains. Furthermore, we notice that the standard deviation for the experiments conducted with the LMs pre-trained on legal data is higher compared to the other models. A possible explanation is that pre-training on legal data improved negation predictions in some areas but adversely affected others, likely due to bias in the legal models, thereby increasing standard deviations across experiments. Generally, cross-domain transfer to the legal domain is less successful than other zero-shot experiments across languages and domains (i.e., Shaitarova et al. (2020); Khandelwal and Sawant (2020)). This suggests that transferring from non-legal to legal domains is challenging.

Zero-shot cross-lingual transfer Table 6 presents the results of our zero-shot cross-lingual experiments conducted with only our legal data. Although these datasets are considerably smaller than the existing English and French datasets, we were able to increase the F1-score by an average of 15.6% across all models and datasets. The legal models still performed well in these experiments, but they no longer showed an advantage over the other LMs. XLM-R_{Base} achieved the best results. All models, except for DistilmBERT, performed significantly better than in the previous experiment across all datasets. DistilmBERT performed worse on the German datasets than in the previous experiment. One explanation for this might be that DistilmBERT is the only cased model used in our experiments. While cased models usually outperform uncased models, this does not seem to apply to cross-lingual experiments. Similar results were found by Macková and Straka (2020), who

Test Dataset	0-shot	1-shot	5-shot	10-shot	Mean F1 by Dataset
fr	13.00 \pm 2.1	16.63 \pm 10.3	14.90 \pm 7.5	22.53 \pm 10.7	16.77 \pm 8.5
it	25.11 \pm 1.5	18.22 \pm 6.5	31.07 \pm 7.1	26.10 \pm 3.8	25.12 \pm 6.7
de (DE)	16.47 \pm 2.6	22.45 \pm 9.1	17.34 \pm 2.7	24.48 \pm 10.7	20.18 \pm 7.5
de (CH)	32.91 \pm 7.9	21.20 \pm 5.8	36.89 \pm 18.6	19.83 \pm 10.3	27.71 \pm 13.1
Mean F1 by experiment	21.87 \pm 8.9	19.62 \pm 7.8	25.05 \pm 13.6	23.23 \pm 8.9	

Table 4: Results for zero- and few-shot experiments conducted over the ChatGPT API.

Model	Test Dataset				Mean F1 by Model
	fr	it	de (DE)	de (CH)	
DistilmBERT	61.43 \pm 1.9	63.40 \pm 2.6	63.50 \pm 4.3	58.78 \pm 4.5	61.78 \pm 3.8
mBERT	66.39 \pm 2.1	68.49 \pm 0.8	64.17 \pm 3.1	54.31 \pm 4.8	63.34 \pm 6.2
XLM-R _{Base}	66.80 \pm 1.9	71.40 \pm 0.8	67.29 \pm 3.7	62.44 \pm 2.9	66.98 \pm 4.0
XLM-R _{Large}	72.30 \pm 2.0	70.30 \pm 0.9	73.81 \pm 4.2	63.72 \pm 4.6	70.03 \pm 5.0
Glott500-m	63.78 \pm 0.8	65.54 \pm 1.1	61.38 \pm 4.0	54.51 \pm 2.5	61.30 \pm 4.9
Legal-Swiss-R _{Base}	69.48 \pm 2.3	68.64 \pm 1.0	71.81 \pm 3.8	54.26 \pm 4.9	66.05 \pm 7.7
Legal-Swiss-R _{Large}	74.66 \pm 2.4	72.68 \pm 1.5	76.5 \pm 1.6	51.75 \pm 6.6	68.89 \pm 10.8
Legal-XLM-R _{Base}	71.50 \pm 3.1	71.48 \pm 2.2	71.35 \pm 5.4	51.93 \pm 3.5	66.57 \pm 9.3
Legal-XLM-R _{Large}	74.52 \pm 2.1	74.48 \pm 3.3	76.06 \pm 3.3	61.30 \pm 8.9	71.59 \pm 7.7
<i>ChatGPT</i>	13.00 \pm 2.1	25.11 \pm 1.5	16.47 \pm 2.6	32.91 \pm 7.9	21.87 \pm 8.9
Mean F1 by Dataset	68.99 \pm 4.9	69.60 \pm 3.7	69.54 \pm 6.4	57.00 \pm 6.4	66.28 \pm 7.6

Table 5: Cross-domain zero-shot results from existing datasets to our new legal datasets. All models except for ChatGPT were pre-trained on all external datasets, ChatGPT did not receive any training data. The bottom right entry shows the average across all datasets and models except ChatGPT.

conducted cross-lingual reading comprehension experiments from English to Czech and found that the uncased models outperformed the cased models in these experiments. They theorized that the overlap of sub-words is larger between English and Czech for uncased models because they disregard diacritical marks, which are common in Czech. A similar argument could be made for the cross-lingual transfer between Italian, French, and German because German includes a lot of casing information while Italian and French do not.

Multilingual experiments The best results for negation scope resolution on our legal datasets were achieved by training our models on the entirety of the available data (Table 7). This multilingual approach achieved an average F1-score of 90% across all models and datasets and outperformed all of the previous setups. This indicates that a relatively small amount of training data in the domain and language of the test dataset can significantly improve the performance of a LM. It is also notable that there seems to be no substantial difference in the performance of the different LMs in

this experiment, with a standard deviation of only ± 3.6 over all models and datasets. Although DistilmBERT obtained the lowest scores in this experiment, its performance is not significantly inferior to that of the mBERT model. This could be attributed to the fact that the training data also included German examples which might have mitigated the advantage of the uncased models with regard to shared vocabulary. We also conducted multilingual experiments only using our new datasets which achieved very similar results with an overall F1-score of 89.1 \pm 4 (see Appendix D).

5.1. Error analysis

We investigated the length of the predicted negation scopes as well as random samples of the predictions on the French and German test data to identify some common error cases.

Predicted scope length As expected, our cross-domain zero-shot experiments without legal training data achieved the lowest F1-scores overall. This can mostly be attributed to the differences in an-

Model	Test Dataset				Mean F1 by Model
	fr	it	de (DE)	de (CH)	
DistilmBERT	79.56 \pm 1.0	74.94 \pm 1.7	58.74 \pm 9.6	52.59 \pm 11.3	66.46 \pm 13.3
mBERT	87.22 \pm 1.6	81.94 \pm 1.3	81.39 \pm 3.6	70.78 \pm 6.7	80.33 \pm 7.1
XLM-R _{Base}	88.70 \pm 0.8	86.43 \pm 2.2	88.00 \pm 1.9	83.71 \pm 4.8	86.71 \pm 3.3
XLM-R _{Large}	90.55 \pm 0.9	84.93 \pm 1.7	91.36 \pm 0.8	76.65 \pm 4.5	85.87 \pm 6.4
Glott500-m	86.77 \pm 2.3	83.41 \pm 1.3	90.10 \pm 2.0	77.73 \pm 4.6	84.50 \pm 5.4
Legal-Swiss-R _{Base}	87.42 \pm 1.2	84.54 \pm 1.6	88.24 \pm 1.0	70.95 \pm 3.6	82.79 \pm 7.4
Legal-Swiss-R _{Large}	84.63 \pm 1.0	83.88 \pm 1.9	88.47 \pm 3.9	70.33 \pm 6.0	81.83 \pm 7.8
Legal-XLM-R _{Base}	86.40 \pm 2.1	83.28 \pm 1.4	89.56 \pm 2.5	74.52 \pm 8.0	83.44 \pm 7.0
Legal-XLM-R _{Large}	85.51 \pm 1.7	85.76 \pm 0.3	89.58 \pm 1.8	80.16 \pm 4.0	85.25 \pm 4.1
Mean F1 by dataset	86.31 \pm 3.2	83.23 \pm 3.5	85.05 \pm 10.4	73.05 \pm 10.3	81.91 \pm 9.3

Table 6: Multilingual zero-shot experiments within our legal datasets. Each column represents a different set of test and train data where the test data includes all legal datasets in languages that are not the language of the test dataset i.e. models evaluated on fr were trained with it and de (DE,CH).

Model	Test Dataset				Mean F1 by Model
	fr	it	de (DE)	de (CH)	
DistilmBERT	87.54 \pm 0.6	82.90 \pm 1.3	94.63 \pm 0.5	90.77 \pm 1.2	88.96 \pm 4.5
mBERT	89.98 \pm 2.1	83.72 \pm 1.0	95.21 \pm 0.5	87.83 \pm 1.0	89.10 \pm 4.4
XLM-R _{Base}	91.31 \pm 1.2	88.81 \pm 1.1	94.74 \pm 0.7	89.39 \pm 1.8	91.06 \pm 2.6
XLM-R _{Large}	90.77 \pm 1.8	87.44 \pm 0.5	93.40 \pm 1.1	90.20 \pm 3.9	90.45 \pm 3.0
Glott500-m	89.65 \pm 1.0	85.54 \pm 2.3	94.94 \pm 0.7	91.00 \pm 2.7	90.28 \pm 3.8
Legal-Swiss-R _{Base}	89.08 \pm 1.6	87.40 \pm 1.9	94.60 \pm 1.0	87.02 \pm 1.5	89.52 \pm 3.4
Legal-Swiss-R _{Large}	89.07 \pm 1.4	86.72 \pm 1.5	95.94 \pm 0.2	89.39 \pm 0.9	90.28 \pm 3.7
Legal-XLM-R _{Base}	90.71 \pm 0.5	86.67 \pm 0.5	95.41 \pm 0.7	86.17 \pm 2.4	89.74 \pm 4.0
Legal-XLM-R _{Large}	90.75 \pm 1.4	89.46 \pm 0.8	93.87 \pm 0.8	89.18 \pm 1.0	90.82 \pm 2.1
Mean F1 by Dataset	89.87 \pm 1.2	86.52 \pm 2.4	94.74 \pm 1.0	88.99 \pm 2.4	90.03 \pm 3.6

Table 7: Results from multilingual experiments over all available data.

notation for each dataset, as well as the different domains. Although the external corpora included French data, this did not improve the performance on the French dataset compared to the other legal datasets. A possible reason is that the subject was not annotated as part of the scope in the Dalloux dataset opposed to the French legal dataset.

Analyzing the predicted scope length compared to the actual scope length reveals one main issue with the zero-shot transfer from the external datasets of different domains to our legal datasets. Figure 2 shows the analysis of the predicted scopes by the Legal-XLM-R_{Large} model. In our cross-domain zero-shot experiment, the predicted scope length is significantly shorter than the actual annotated scope length. This is clarified by Table 2, revealing the external datasets have a shorter annotated scope length (24%) compared to our legal datasets (38.6%). Sample predictions confirm that

the model often omits the subject from the annotated scope.

Annotation : Es sei festzustellen, dass der Rück-
erstattungsanspruch **nicht** verjährt sei.

EN: *It should be noted that the claim for restitution is **not** forfeited.*

Prediction : Es sei festzustellen, dass der Rücker-
stattungsanspruch **nicht** verjährt sei.

EN: *It should be noted that the claim for restitution is **not** forfeited.*

As soon as some legal data is added to our training sets, the predicted scope length as well as the F1-score increases. An inspection of the predictions made by the legal and multilingual models shows that the additional training data helps to predict the subject as part of the scope. One exception where the subject was not annotated in the prediction is for subjects represented by an initial instead of a pronoun or a full name, which is common in




Figure 2: Actual scope length and scope length predicted by Legal-XLM- R_{Large} for each experiment. X marks the scope length of the train data.

legal documents for anonymization reasons. We suspect that in these cases the models were not able to identify the initial as the subject because these kinds of subjects might be more uncommon outside of the legal domain.

Annotation: E. ne disposait d'aucune autonomie budgétaire;

EN: E. had no budgetary autonomy

Prediction: E. ne disposait d'aucune autonomie budgétaire;

EN: E. had no budgetary autonomy

Non-continuous scopes Another error case is sentences where the scope is not continuous because it is interrupted by an interjection or contrasting statement. These kinds of sentences are more complex than the average sentence and not very common in the training data. A larger amount of training data containing similar sentence structures could improve accuracy.

Annotation: Eine ordentliche Kündigung ist während der vereinbarten Laufzeit beiderseits nur zum Vertragsende und nicht zu einem früheren Zeitpunkt zulässig.

EN An ordinary termination during the agreed term is only permissible on both sides at the end of the contract and not at an earlier time

Prediction: Eine ordentliche Kündigung ist während der vereinbarten Laufzeit beiderseits nur zum Vertragsende und nicht zu einem früheren Zeitpunkt zulässig.

EN An ordinary termination during the agreed term is only permissible on both sides at the end of the contract and not at an earlier time

6. Conclusions and Future Work

6.1. Conclusion

We released new legal datasets in German, French and Italian, annotated for negation cues and scopes and showed that the legal domain does pose a challenge for models in negation scope resolution. Cross-domain zero-shot experiments showed that models without legal training data do not perform as well on multilingual legal datasets as they do on other domains. The task is also too complex for ChatGPT, which was not able to reach F1-scores above 37%. Using our new datasets we fine-tuned different models on the legal domain, significantly improving the results and showing that even relatively small amounts of training data in a specific domain and language can improve the performance of multilingual LMs for negation scope resolution.

6.2. Future Work

Negation scope resolution models in the legal domain could benefit from more training data to increase the accuracy of predictions of more complex sentence structures such as non-continuous scopes. More diverse data from different legal fields could further improve the performance of negation scope models in the legal domain.

With our new datasets we were able to show that existing systems performing well on datasets across different domains are not necessarily able to perform as well on legal data. This should motivate future work to focus on this complex domain and evaluate the performance of existing systems in diverse NLP tasks.

Limitations

Due to resource constraints, our datasets are relatively small compared to other publicly available corpora. A larger set of legal data across a diverse set of sources, annotated with negations could further improve the performance of LMs for negation scope resolution in this field. We also did not investigate the potential of cross-lingual cue detection since this is the more trivial part of negation research and can easily be replaced by a list of negation cues for each language.

Ethics Statement

The goal of our work was to improve the performance of negation scope resolution systems in the legal domain. These improved systems could be used to support legal professionals in processing and analysing legal texts. These systems should only be used as an assistance to human experts with considerations to their limitations and possible

biases. To the best of our knowledge there is currently no real world application of a negation scope resolution system in the legal domain.

The legal data that we annotated and used to train our models is all publicly available and has all been anonymized. It should therefore not include any sensitive information.

Acknowledgments

We would like to thank our annotators from the University of Bern for their work. Furthermore, we want to express gratitude to the anonymous reviewers of this paper for their valuable feedback.

7. Bibliographical References

- Begoña Altuna, Anne-Lyse Minard, and Manuela Speranza. 2017. [The scope and focus of negation: A complete annotation framework for Italian](#). In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pages 34–42, Valencia, Spain. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Tobias Brugger, Matthias Stürmer, and Joel Niklaus. 2023. [MultiLegalSBD: A Multilingual Legal Sentence Boundary Detection Dataset](#). ArXiv:2305.01211 [cs].
- Ilias Chalkidis, Manos Fergadiotis, and Ion Androutsopoulos. 2021. [MultiEURLEX - a multilingual and multi-label legal document classification dataset for zero-shot cross-lingual transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6974–6996, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [LexGLUE: A Benchmark Dataset for Legal Language Understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330.
- Wendy W. Chapman, Will Bridewell, Paul Hanbury, Gregory F Cooper, and Bruce G Buchanan. 2001. [A simple algorithm for identifying negated findings and diseases in discharge summaries](#). *Journal of biomedical informatics*, 34(5):301–310.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised Cross-lingual Representation Learning at Scale](#). *arXiv:1911.02116 [cs]*. ArXiv: 1911.02116.
- Viviana Cotik, Roland Roller, Feiyu Xu, Hans Uszkoreit, Klemens Budde, and Danilo Schmidt. 2016. [Negation detection in clinical reports written in German](#). In *Proceedings of the Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM2016)*, pages 115–124, Osaka, Japan. The COLING 2016 Organizing Committee.
- Clément Dalloux, Vincent Claveau, Natalia Grabar, Lucas Oliveira, Claudia Moro, Yohan Gumiel, and Deborah Carvalho. 2020. [Supervised learning for the detection of negation and of its scope in french and brazilian portuguese biomedical corpora](#). *Natural Language Engineering*, 27:1–21.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *arXiv:1810.04805 [cs]*. ArXiv:1810.04805.
- Federico Fancellu, Adam Lopez, and Bonnie Webber. 2018. [Neural networks for cross-lingual negation scope detection](#).
- Federico Fancellu, Adam Lopez, Bonnie Webber, and Hangfeng He. 2017. [Detecting negation scope is easy, except when it isn't](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 58–63, Valencia, Spain. Association for Computational Linguistics.
- Roland Friedrich. 2021. [Complexity and entropy in legal language](#). *Frontiers in Physics*, 9:671882.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [Chatgpt outperforms crowd](#)

- workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Ingo Glaser., Sebastian Moser., and Florian Matthes. 2021. [Sentence boundary detection in german legal documents](#). In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*, pages 812–821. INSTICC, SciTePress.
- Ayyoob ImaniGooghari, Peiqin Lin, Amir Hossein Kargaran, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André Martins, François Yvon, and Hinrich Schütze. 2023. [Glot500: Scaling multilingual corpora and language models to 500 languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1082–1117, Toronto, Canada. Association for Computational Linguistics.
- Salud María Jiménez-Zafra, Roser Morante, María Teresa Martín-Valdivia, and L. Alfonso Ureña-López. 2020. [Corpora annotated with negation: An overview](#). *Computational Linguistics*, 46(1):1–52.
- Aditya Khandelwal and Suraj Sawant. 2020. [NegBERT: A Transfer Learning Approach for Negation Detection and Scope Resolution](#). *arXiv:1911.04211 [cs]*. ArXiv: 1911.04211.
- Natalia Konstantinova, Sheila CM De Sousa, Noa P Cruz Díaz, Manuel J Mana López, Maite Taboada, and Ruslan Mitkov. 2012. [A review corpus annotated for negation, speculation and their scope](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3190–3195.
- Kateřina Macková and Milan Straka. 2020. Reading comprehension in czech via machine translation and cross-lingual transfer. In *Text, Speech, and Dialogue*, pages 171–179, Cham. Springer International Publishing.
- Sabrina J. Mielke. 2016. [Language diversity in ACL 2004 - 2016](#).
- Roser Morante and Eduardo Blanco. 2012. [* sem 2012 shared task: Resolving the scope and focus of negation](#). In ** SEM 2012: The First Joint Conference on Lexical and Computational Semantics—Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 265–274, Montréal, Canada. Association for Computational Linguistics.
- Roser Morante, Anthony Liekens, and Walter Daelemans. 2008. [Learning the scope of negation in biomedical texts](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 715–724, Honolulu, Hawaii. Association for Computational Linguistics.
- Roser Morante, Sarah Schrauwen, and Walter Daelemans. 2011. Annotation of negation cues and their scope: Guidelines v1. *Computational linguistics and psycholinguistics technical report series, CTRS-003*, pages 1–42.
- Joel Niklaus, Ilias Chalkidis, and Matthias Stürmer. 2021. [Swiss-Judgment-Prediction: A Multilingual Legal Judgment Prediction Benchmark](#). In *Proceedings of the Natural Legal Language Processing Workshop 2021*, pages 19–35, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Joel Niklaus, Veton Matoshi, Pooja Rani, Andrea Galassi, Matthias Stürmer, and Ilias Chalkidis. 2023a. [LEXTREME: A Multi-Lingual and Multi-Task Benchmark for the Legal Domain](#). ArXiv:2301.13126 [cs].
- Joel Niklaus, Veton Matoshi, Matthias Stürmer, Ilias Chalkidis, and Daniel E. Ho. 2023b. [Multi-LegalPile: A 689GB Multilingual Legal Corpus](#). ArXiv:2306.02069 [cs].
- Joel Niklaus, Matthias Stürmer, and Ilias Chalkidis. 2022. [An empirical study on cross-x transfer for legal judgment prediction](#). ArXiv:2209.12325.
- Vishvaksenan Rasiah, Ronja Stern, Veton Matoshi, Matthias Stürmer, Ilias Chalkidis, Daniel E. Ho, and Joel Niklaus. 2023. [SCALE: Scaling up the Complexity for Advanced Language Model Evaluation](#). ArXiv:2306.09237 [cs].
- JB Ruhl, Daniel Martin Katz, and Michael J Bommarito. 2017. Harnessing legal complexity. *Science*, 355(6332):1377–1378.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv:1910.01108 [cs]*. ArXiv: 1910.01108.
- Anastassia Shaitarova, Lenz Furrer, and Fabio Rinaldi. 2020. [Cross-lingual transfer-learning approach to negation scope resolution](#). In *CEUR Workshop proceedings*. CEUR-WS.
- Anastassia Shaitarova and Fabio Rinaldi. 2021. [Negation typology and general representation models for cross-lingual zero-shot negation](#)

scope resolution in Russian, French, and Spanish. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 15–23, Online. Association for Computational Linguistics.

György Szarvas, Veronika Vincze, Richárd Farkas, and János Csirik. 2008. [The bioscope corpus: annotation for negation, uncertainty and their scope in biomedical texts](#). In *Proceedings of the workshop on current trends in biomedical natural language processing*, pages 38–45. Association for Computational Linguistics.

Appendix

A. ChatGPT negation scope example

User: In the following German sentence give me a list of all words that are part of a negation scope. Negation scope describes all words affected by a negation in a sentence. Format the output as json. Sentence: Ich weiss nicht was eine Kartoffel ist.

ChatGPT: Certainly, here's the JSON output for the words that are part of the negation scope in the provided German sentence:

```

1 {
2   "Neg_scope": [
3     "weiss",
4     "was",
5     "eine",
6     "Kartoffel",
7     "ist"
8   ]

```

Figure 3: ChatGPT annotation of a simple sentence.

User: In the following German sentence give me a list of all words that are part of a negation scope. Negation scope describes all words affected by a negation in a sentence. Format the output as json. Sentence: Die Gewährung der tariflichen Leistungen, somit auch die Versorgungszusage, die laut Tarifvertrag für Auszubildende nicht gelte, hätte überhaupt nur dann Sinn gemacht, wenn die Versorgungszusage entgegen des Geltungsbereichs des Tarifvertrages auch dem Auszubildenden erteilt hätte werden sollen.

ChatGPT: Certainly, here's the JSON output for the words that are part of the negation scope in the provided German sentence:

B. Negation Words

The following lists of words were used to preselect documents for annotation which contain a large

```

1 {
2   "Neg_scope": [
3     "gelte"
4   ]
5 }

```

Figure 4: ChatGPT annotation of a more complex legal sentence.

amount of negations. It consists of a self-compiled list based on our and our native speaking annotators knowledge of each language.

German: kein, keiner, nein, nicht, nichts, nie, niemals, niemand, nirgends, nirgendwo, nirgendwoher, nirgendwohin, keinesfalls, keineswegs, mitnichten, mangels, nur, weder

Italian: non, mai, né, nemmeno, neppure, neanche, nessuno, nessuna, niente, nulla, meno, senza

French: pas, plus, jamais, guère, rien, aucun, aucune

C. Mean results over all models, datasets and experiments

Model \ Test Dataset	Zero-shot cross-domain	Zero-shot cross-lingual	Multilingual - legal data	Multilingual - all data	Mean F1 by model
DistilmBERT	61.78 \pm 3.77	66.46 \pm 13.33	0.87 \pm 0.05	88.96 \pm 4.51	72.40 \pm 14.54
mBERT	63.34 \pm 6.24	80.33 \pm 7.10	0.89 \pm 0.04	89.19 \pm 4.41	77.62 \pm 12.33
XLM-R _{Base}	66.98 \pm 4.02	86.71 \pm 3.25	0.90 \pm 0.03	91.06 \pm 2.63	81.59 \pm 11.07
XLM-R _{Large}	70.03 \pm 4.98	85.87 \pm 6.44	0.90 \pm 0.04	90.45 \pm 3.00	82.12 \pm 10.10
Glott500-m	61.30 \pm 4.85	84.50 \pm 5.36	0.89 \pm 0.04	90.28 \pm 3.84	78.70 \pm 13.46
Legal-Swiss-R _{Base}	66.05 \pm 7.72	82.79 \pm 7.41	0.88 \pm 0.05	89.52 \pm 3.41	79.45 \pm 11.82
Legal-Swiss-R _{Large}	68.89 \pm 10.80	81.83 \pm 7.81	0.90 \pm 0.03	90.28 \pm 3.66	80.33 \pm 11.84
Legal-XLM-R _{Base}	66.57 \pm 9.33	83.44 \pm 7.01	0.90 \pm 0.04	89.74 \pm 3.99	79.92 \pm 12.10
Legal-XLM-R _{Large}	71.59 \pm 7.73	85.25 \pm 4.07	0.89 \pm 0.04	90.82 \pm 2.12	82.55 \pm 9.61
Mean F1 by experiment	66.28 \pm 7.64	81.91 \pm 9.26	0.89 \pm 0.04	90.03 \pm 3.57	

Table 8: Mean Results over all models and experiments

D. Multilingual results legal data

Model \ Test Dataset	fr	it	de (DE)	de (CH)	Mean F1 by Model
DistilmBERT	86.06 \pm 0.76	81.82 \pm 0.79	93.98 \pm 0.82	87.40 \pm 2.36	87.32 \pm 4.65
mBERT	90.16 \pm 1.33	84.56 \pm 1.63	94.95 \pm 0.80	86.81 \pm 2.06	89.12 \pm 4.25
XLM-R-Base	90.26 \pm 0.96	88.05 \pm 1.81	94.12 \pm 0.59	87.21 \pm 2.66	89.91 \pm 3.16
XLM-R-Large	90.23 \pm 1.40	86.93 \pm 0.73	94.56 \pm 0.85	86.44 \pm 3.56	89.54 \pm 3.80
Glott500-m	88.81 \pm 1.47	85.62 \pm 1.12	94.23 \pm 1.40	88.13 \pm 2.60	89.20 \pm 3.59
Legal-Swiss-R-Base	87.98 \pm 1.46	89.53 \pm 0.54	93.15 \pm 0.44	81.82 \pm 3.88	88.12 \pm 4.62
Legal-Swiss-R-Large	88.35 \pm 0.88	88.20 \pm 1.13	95.30 \pm 0.37	89.39 \pm 1.37	90.31 \pm 3.13
Legal-XLM-R-Base	88.89 \pm 1.58	88.41 \pm 1.84	95.56 \pm 0.88	85.27 \pm 3.83	89.53 \pm 4.39
Legal-XLM-R-Large	88.86 \pm 0.95	87.98 \pm 0.64	94.46 \pm 0.69	85.30 \pm 3.12	89.15 \pm 3.76
Mean F1 by dataset	88.84 \pm 1.70	86.79 \pm 2.55	94.48 \pm 1.01	86.42 \pm 3.36	89.13 \pm 3.97

Table 9: Results of multilingual experiments using only our legal datasets.

E. Precision Scores

Model \ Experiment	Multilingual - all data	Multilingual - legal data	Zero-shot cross-domain	Zero-shot cross-lingual
XLM-R-Large	87.9 \pm 5.2	85.8 \pm 5.1	95.2 \pm 3.3	89.1 \pm 5.7
XLM-R-Base	89.1 \pm 4.2	88.3 \pm 2.5	94.8 \pm 2.2	90.2 \pm 5.8
Glott500-m	90.9 \pm 3.6	90.1 \pm 4.1	95.8 \pm 2.1	89.0 \pm 5.5
mBERT	87.8 \pm 5.9	89.1 \pm 4.1	96.1 \pm 2.2	90.9 \pm 5.1
DistilMBERT	88.9 \pm 5.0	90.3 \pm 3.1	94.0 \pm 4.0	90.8 \pm 6.5
Legal-XLM-R-Large	89.3 \pm 3.9	89.0 \pm 4.0	94.6 \pm 2.6	88.7 \pm 5.5
Legal-Swiss-R-Large	90.6 \pm 5.0	89.5 \pm 3.8	92.3 \pm 5.0	88.4 \pm 6.0
Legal-XLM-R-Base	88.5 \pm 4.9	88.0 \pm 4.1	94.5 \pm 3.7	87.4 \pm 5.5
Legal-Swiss-R-Base	87.9 \pm 4.8	87.8 \pm 3.6	89.2 \pm 6.3	87.4 \pm 3.9

Table 10: Precision scores over all experiments.

F. Recall Scores

Model \ Experiment	Multilingual - all data	Multilingual - legal data	Zero-shot cross-domain	Zero-shot cross-lingual
XLM-R-Large	93.5 \pm 4.1	93.8 \pm 3.3	55.7 \pm 6.0	83.8 \pm 10.4
XLM-R-Base	93.4 \pm 4.0	91.7 \pm 5.1	52.0 \pm 5.2	84.2 \pm 7.4
Glott500-m	89.9 \pm 5.6	88.5 \pm 5.6	45.3 \pm 5.4	80.9 \pm 7.7
mBERT	91.0 \pm 6.0	89.4 \pm 6.2	47.6 \pm 6.8	72.9 \pm 10.7
DistilMBERT	89.5 \pm 7.5	84.7 \pm 6.9	46.2 \pm 4.5	55.2 \pm 17.4
Legal-XLM-R-Large	92.6 \pm 3.1	89.6 \pm 6.1	58.1 \pm 9.1	82.8 \pm 7.8
Legal-Swiss-R-Large	90.2 \pm 5.0	91.3 \pm 4.1	55.7 \pm 11.6	77.3 \pm 12.2
Legal-XLM-R-Base	91.2 \pm 4.5	91.1 \pm 5.3	52.3 \pm 10.9	80.4 \pm 9.9
Legal-Swiss-R-Base	91.4 \pm 3.9	88.6 \pm 6.5	53.1 \pm 9.0	79.2 \pm 10.8

Table 11: Recall scores over all experiments.