

An Empirical Study on Cross-X Transfer for Legal Judgment Prediction

Joel Niklaus^{†*} Matthias Stürmer[†] Ilias Chalkidis^{‡‡*}

[†] Institute of Computer Science, University of Bern, Switzerland

[‡] Department of Computer Science, University of Copenhagen, Denmark

[◊] Cognitiv+, Athens, Greece

Abstract

Cross-lingual transfer learning has proven useful in a variety of Natural Language Processing (NLP) tasks, but it is understudied in the context of legal NLP, and not at all in Legal Judgment Prediction (LJP). We explore transfer learning techniques on LJP using the trilingual Swiss-Judgment-Prediction dataset, including cases written in three languages. We find that cross-lingual transfer improves the overall results across languages, especially when we use adapter-based fine-tuning. Finally, we further improve the model’s performance by augmenting the training dataset with machine-translated versions of the original documents, using a $3\times$ larger training corpus. Further on, we perform an analysis exploring the effect of cross-domain and cross-regional transfer, i.e., train a model across domains (legal areas), or regions. We find that in both settings (legal areas, origin regions), models trained across all groups perform overall better, while they also have improved results in the worst-case scenarios. Finally, we report improved results when we ambitiously apply cross-jurisdiction transfer, where we further augment our dataset with Indian legal cases.

1 Introduction

Rapid development in Cross-Lingual Transfer (CLT) has been achieved by pre-training transformer-based models in large multilingual corpora (Conneau et al., 2020; Xue et al., 2021), where these models have state-of-the-art results in multilingual NLU benchmarks (Ruder et al., 2021). Moreover, adapter-based fine-tuning (Houlsby et al., 2019; Pfeiffer et al., 2020) has been proposed to minimize the misalignment of multilingual knowledge (alignment) when CLT is applied, especially in a zero-shot fashion, where the target language is unseen during training. CLT is severely understudied in legal NLP applications except for

* Equal contribution.

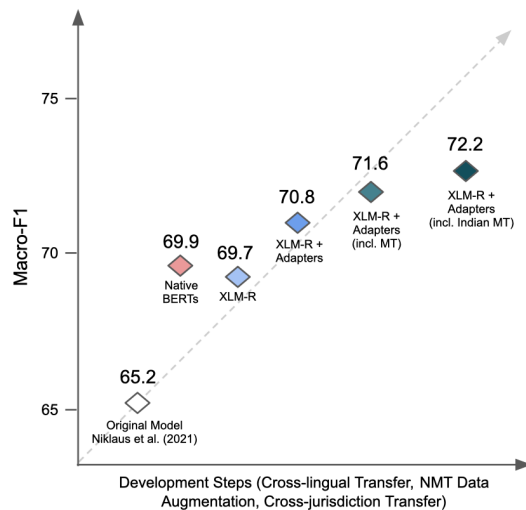


Figure 1: Incremental performance improvement through several development steps.

Chalkidis et al. (2021) who experimented with several methods for CLT on MultiEURLEX, a newly introduced multilingual legal topic classification dataset, including EU laws.

To the best of our knowledge, CLT has not been applied to the Legal Judgment Prediction (LJP) task (Aletras et al., 2016; Xiao et al., 2018; Chalkidis et al., 2019; Malik et al., 2021), where the goal is to predict the verdict (court decision) given the facts of a legal case. In this setting, positive impact of cross-lingual transfer is not as conceptually straight-forward as in other general applications (NLU), since there are known complications for sharing legal definitions and interpreting law across languages (Gotti, 2014; McAuliffe, 2014; Robertson, 2016; Ramos, 2021).

Following the work of Niklaus et al. (2021), we experiment with their newly released trilingual Swiss-Judgment-Prediction (SJP) dataset, containing cases from the Federal Supreme Court of Switzerland (FSCS), written in three official Swiss languages (German, French, Italian). The dataset covers four legal areas (public, penal, civil, and social law) and lower courts located in eight regions of Switzerland (Zurich, Ticino, etc.), which poses

interesting new challenges on model robustness / fairness and the effect of cross-domain and cross-regional knowledge sharing. In their experiments, Niklaus et al. (2021) find that the performance in cases written in Italian is much lower compared to the rest, while also performance varies a lot across regions and legal areas.

Main Research Questions

We pose and examine four main research questions:

RQ1: *Is cross-lingual transfer beneficial across all or some of the languages?*

RQ2: *Do models benefit or not from cross-regional and cross-domain transfer?*

RQ3: *Can we leverage data from another jurisdiction to improve performance?*

RQ4: *How does representational bias (wrt. language, origin region, legal area) affect model’s performance?*

Contributions

The contributions of this paper are fourfold:

- We explore, for the first time, the application of cross-lingual transfer learning in the challenging LJP task in several settings (Section 3.3). We find that a pre-trained language model fine-tuned multilingually, outperforms its monolingual counterparts, especially when we use adapter-based fine-tuning and augment the training data with machine-translated versions of the original documents (3× larger training corpus) with larger gains in a low-resource setting (Italian).
- We perform cross-domain and cross-regional analyses (Section 3.4) exploring the effects of cross-domain and cross-regional transfer, i.e., train a model across domains, i.e., legal areas (e.g., civil, penal law), or regions (e.g., Zurich, Ticino). We find that in both settings (legal areas, regions), models trained across all groups perform overall better and more robustly; while always improving performance in the worst-case (region or legal area) scenario.
- We also report improved results when we apply cross-jurisdiction transfer (Section 3.5), where we further augment our dataset with Indian legal cases originally written in English.
- We release the augmented dataset (incl. 100K machine-translated documents) and our code for replicability and future experimentation.¹

¹https://huggingface.co/datasets/swiss_judgment_prediction

The cumulative performance improvement amounts to 7% overall and 16+% in the low-resource Italian subset, compared to the best reported scores in Niklaus et al. (2021), while using cross-lingual and cross-jurisdiction transfer we improve for 2.3% overall and 4.6% for Italian over our strongest baseline (NativeBERTs).

2 Dataset and Task description

2.1 Swiss Legal Judgment Prediction Dataset

We investigate the LJP task on the Swiss-Judgment-Prediction (SJP) dataset (Niklaus et al., 2021). The dataset contains 85K cases from the Federal Supreme Court of Switzerland (FSCS) from the years 2000 to 2020 written in German, French, and Italian. The court hears appeals focusing on small parts of the previous (lower court) decision, where they consider possible wrong reasoning by the lower court. The dataset provides labels for a simplified binary (*approval, dismissal*) classification task. Given the facts of the case, the goal is to predict if the plaintiff’s request is valid or partially valid (i.e., the court *approved* the complaint).

Since the dataset contains rich metadata, such as legal areas and origin regions, we can conduct experiments on the robustness of the models (see Section 3.4). The dataset is not equally distributed; in fact, there is a notable representation disparity where Italian have far fewer documents (4K), compared to German (50K) and French (31K). Representation disparity is also vibrant with respect to legal areas and regions. We refer readers to the work of Niklaus et al. for detailed dataset statistics.

2.2 Indian Legal Judgment Prediction Dataset

The Indian Legal Documents Corpus (ILDC) dataset (Malik et al., 2021) comprises 30K cases from the Indian Supreme Court in English. The court hears appeals that usually include multiple petitions and rules a decision (*accepted vs. rejected*) per petition. Similarly to Niklaus et al. (2021), Malik et al. released a simplified version of the dataset with binarized labels. In effect, the two datasets (SJP, ILDC) target the very same task (partial or full approval of plaintiff’s claims), nonetheless in two different jurisdictions (Swiss Federation and India). Our main goal, when we use ILDC as a complement of SJP, is to assess the possibility of cross-jurisdiction transfer from Indian to Swiss cases (see Section 3.5), an experimental scenario

that has not been explored so far in the literature.

2.3 NMT-based Data Augmentation

In some of our experiments, we perform data augmentation using machine-translated versions of the original documents, i.e., translate a document originally written in a single language to the other two (e.g., from German to French and Italian). We performed the translations using the EasyNMT² framework utilizing the *many-to-many* Neural Machine Translation (NMT) model of Fan et al. (2020).³ A preliminary manual check of some translated samples showed sufficient translation quality to proceed forward. We release the machine-translated additional dataset for future consideration on cross-lingual experiments or quality assessment.

To the best of our knowledge, machine translation for data augmentation has not been studied in legal Natural Language Processing (NLP) applications, while it is generally a straight-forward, though under-studied idea. As we show in the experiments (see Section 3.3), the translations are effective, leading to an average improvement of 1.6% macro-F1 for standard fine-tuning and 0.8% for adapter-based one (see Table 1). For the low-resource Italian subset, the improvement even amounts to 3.2% and 1.6%, respectively.

3 Experiments

3.1 Hierarchical BERT

Since the examined dataset (SJP) contains many documents with more than 512 tokens (90% of the documents are up to 2048), we use Hierarchical BERT models (Chalkidis et al., 2019; Niklaus et al., 2021; Dai et al., 2022) to encode up to 2048 tokens per document (4×512 blocks).

We split the text into consecutive blocks of 512 tokens and feed the first 4 blocks to a shared standard BERT encoder. Then, we aggregate the block-wise CLS tokens by passing them through another 2-layer transformer encoder, followed by max-pooling and a final classification layer.

We re-use and expand the implementation released by Niklaus et al. (2021),⁴ which is based on the Hugging Face library (Wolf et al., 2020). Notably, we first improve the masking of the blocks.

²<https://github.com/UKPLab/EasyNMT>

³The *one-to-one* OPUS-MT (Tiedemann and Thottingal, 2020) models did not have any model available from French to Italian (fr2it) at the time of the experiments.

⁴<https://github.com/JoelNiklaus/SwissJudgementPrediction>

Specifically, when the document has less than the maximum number (4) of blocks, we pad with extra sequences of PAD tokens, without the use of special tokens (CLS, SEP), as was previously performed. This minor technical improvement seems to affect the model’s performance at large (group A1 Prior SotA vs. NativeBERTs — Table 1).

We experiment with monolingually pre-trained BERT models (aka NativeBERTs) and the multilingually pre-trained XLM-R of Conneau et al. (2020). Specifically, for monolingual experiments (Native BERTs), we use German-BERT (Chan et al., 2019) for German, CamemBERT (Martin et al., 2020) for French, and UmBERTo (Parisi et al., 2020) for Italian, similar to Niklaus et al. (2021).

In our multilingual experiments, we also assess the effectiveness of adapter-based fine-tuning (Houlsby et al., 2019; Pfeiffer et al., 2020), in comparison to standard full fine-tuning. In this setting, adapter layers are placed after all feed-forward layers of XLM-R and are trained together with the parameters of the layer-normalization layers. The rest of the model parameters remain untouched.

3.2 Experimental Set Up

We follow Niklaus et al. (2021) and report macro-averaged F1 score to account for the high class-imbalance in the dataset (approx. 20/80 approval/dismissal ratio). We repeat each experiment with 3 different random seeds and report the average score and standard deviation across runs (seeds). We perform grid-search for the learning rate and report test results, selecting the hyperparameters with the best development scores.⁵

3.3 Cross-lingual Transfer

We first examine *cross-lingual transfer*, where the goal is to share (transfer) knowledge across languages, and we compare models in three main settings: (a) *Monolingual* (see Section 3.3.1): fine-tuned per language, using either the documents originally written in the language, or an augmented training set including the machine-translated versions of all other documents (originally written in another language), (b) *Cross-lingual* (see Section 3.3.2): fine-tuned across languages with or without the additional translated versions, and (c) *Zero-shot cross-lingual* (see Section 3.3.3): fine-tuned across a subset of the languages excluding the target language at a time. We present the results in Table 1.

⁵Additional details on model configuration, training, and hyper-parameter tuning can be found in Appendix A.

Model	#D	#M	German \uparrow	French \uparrow	Italian \uparrow	All \uparrow	(Diff. \downarrow)
<i>A1. Monolingual: Fine-tune on the tgt training set (src = tgt) — Baselines</i>							
Prior SotA (Niklaus et al.)	3-35K	N	68.5 \pm 1.6	70.2 \pm 1.1	57.1 \pm 0.4	65.2 \pm 0.8	(13.1)
NativeBERTs	3-35K	N	<u>69.6</u> \pm 0.4	<u>72.0</u> \pm 0.5	<u>68.2</u> \pm 1.3	69.9 \pm 1.6	(3.8)
XLM-R	3-35K	N	68.2 \pm 0.3	69.9 \pm 1.6	65.9 \pm 1.2	68.0 \pm 2.0	(4.0)
<i>A2. Monolingual: Fine-tune on the tgt training set incl. machine-translations (src = tgt)</i>							
NativeBERTs	60K	N	<u>70.0</u> \pm 0.7	<u>71.0</u> \pm 1.3	<u>71.9</u> \pm 2.5	71.0 \pm 0.8	(0.9)
XLM-R	60K	N	68.8 \pm 1.4	70.7 \pm 2.1	71.9 \pm 2.6	70.4 \pm 1.3	(1.1)
<i>B1. Cross-lingual: Fine-tune on all training sets (src \subset tgt)</i>							
XLM-R	60K	1	68.9 \pm 0.3	71.1 \pm 0.3	68.9 \pm 1.4	69.7 \pm 1.0	(2.2)
XLM-R + Adapters	60K	1	<u>69.9</u> \pm 0.6	<u>71.8</u> \pm 0.7	<u>70.7</u> \pm 1.8	<u>70.8</u> \pm 0.8	(0.9)
<i>B2. Cross-lingual: Fine-tune on all training sets incl. machine-translations (src \subset tgt)</i>							
XLM-R	180K	1	70.2 \pm 0.5	71.5 \pm 1.1	72.1 \pm 1.2	71.3 \pm 0.7	(1.9)
XLM-R + Adapters	180K	1	70.3 \pm 0.9	72.1 \pm 0.8	72.3 \pm 2.1	71.6 \pm 0.8	(2.0)
<i>C. Zero-shot Cross-lingual: Fine-tune on all training sets excl. tgt language (src \neq tgt)</i>							
XLM-R	25-57K	1	58.4 \pm 1.2	58.7 \pm 0.8	<u>68.1</u> \pm 0.2	61.7 \pm 4.5	(9.7)
XLM-R + Adapters	25-57K	1	<u>62.5</u> \pm 0.6	<u>58.8</u> \pm 1.5	67.5 \pm 2.2	<u>62.8</u> \pm 3.7	(8.7)

Table 1: Test results for all training set-ups (monolingual w/ or w/o translations, multilingual w/ or w/o translations, and zero-shot) w.r.t source (src) and target (tgt) language. Best overall results are in **bold**, and best per setting (group) are underlined. #D is the number of training documents used. #M is the number of models trained/used. The mean and standard deviation are computed across random seeds and across languages for the last column. Diff. shows the difference between the best and the worst performing language. *The adapter-based multilingually fine-tuned XLM-R model including machine-translated versions (3 \times larger corpus) has the best overall results.*

3.3.1 Mono-Lingual Training

We observe that the baseline of *monolingually* pre-trained and fine-tuned models (NativeBERTs) have the best results compared to the *multilingually* pre-trained but *monolingually* fine-tuned XLM-R (group A1 – Table 1). Representational bias across languages (Section 2.1) seems to be a key part of performance disparity, considering the performance of the least represented language (Italian) compared to the rest (3K vs. 21-35K training documents). However, this is not generally applicable, i.e., French have better performance compared to German, despite having approx. 30% less training documents.

Translating the full training set provides a 3 \times larger training set (approx. 180K in total) that “equally” represents all three languages.⁶ Augmenting the original training sets with translated versions of the documents (group A2 – Table 1), originally written in another language, improves per-

⁶Representational equality with respect to number of training documents per language, but possibly not considering text quality, since we use NMT to achieve that goal.

formance in almost all (5/6) cases (languages per model). Interestingly, the performance improvement in Italian, which has the least documents (less than 1/10 compared to German), is the largest across languages with 3.7% for NativeBERT (68.2 to 71.9) and 6% for XLM-R (65.9 to 71.9) making Italian the best performing language after augmentation. Data augmentation seems more beneficial for XLM-R, which does not equally represent the three examined languages.⁷

3.3.2 Cross-Lingual Training

We now turn to the *cross-lingual transfer* setting, where we train XLM-R across all languages in parallel. We observe that cross-lingual transfer (group B1 – Table 1) improves performance (+4.5% p.p.) across languages compared to the same model (XLM-R) fine-tuned in a monolingual setting (group A1 – Table 1). This finding suggests that cross-lingual transfer (and the inherited benefit of using larger multilingual corpora) has a signifi-

⁷Refer to Conneau et al. (2020) for resources per language used to pre-train XLM-R (50% less tokens for Italian).

Origin Region	#D	#L	ZH	ES	CS	NWS	EM	RL	TI	FED	All
Region-specific fine-tuning with MT data augmentation											
Zürich (ZH)	26.4K	de	<u>65.5</u>	65.6	63.7	68.2	62.0	57.9	63.2	54.8	62.6
Eastern Switzerland (ES)	17.1K	de	62.9	<u>66.9</u>	62.8	65.2	62.2	60.2	57.8	55.1	61.6
Central Switzerland (CS)	14.4K	de	62.5	65.5	<u>63.2</u>	65.1	60.7	57.8	60.5	55.9	61.4
Northwestern Switzerland (NWS)	17.1K	de	66.0	68.6	<u>65.2</u>	<u>67.9</u>	61.6	57.0	57.1	55.5	62.4
Espace Mittelland (EM)	24.9K	de,fr	64.1	66.6	63.3	66.7	<u>64.0</u>	66.8	63.2	58.4	64.1
Région Lémanique (RL)	40.2K	fr,de	61.0	64.7	60.2	63.7	63.4	<u>69.8</u>	67.6	54.3	63.1
Ticino (TI)	6.9K	it	55.0	56.3	53.2	54.5	56.0	54.7	<u>66.0</u>	53.1	56.1
Federation (FED)	3.9K	de,fr,it	57.5	59.6	56.8	58.9	55.0	56.5	53.5	<u>54.9</u>	56.6
Cross-regional fine-tuning w/o MT data augmentation											
XML-R	60K	de,fr,it	68.5	71.3	67.7	71.2	69.0	71.4	67.4	64.6	68.9
XML-R + Adapters	60K	de,fr,it	69.2	73.9	67.9	72.6	69.0	72.1	70.1	64.2	69.9
Cross-regional fine-tuning with MT data augmentation											
NativeBERTs	180K	de,fr,it	69.0	72.1	68.6	72.0	69.9	71.9	68.8	64.8	69.6
XML-R	180K	de,fr,it	69.2	72.9	68.3	73.3	69.9	71.7	70.4	65.0	70.1
XML-R + Adapters	180K	de,fr,it	69.2	73.3	69.9	73.0	70.3	72.1	70.9	63.8	70.3

Table 2: Test results for models trained per region or across all regions. Best overall results are in **bold**, and in-domain are underlined. #D is the total number of training examples. #L are the languages covered. *Cross-regional transfer is beneficial for all regions and has the best overall results. The shared multilingual model trained across all languages and regions slightly outperforms the baseline (NativeBERTs).*

cant impact, despite the legal complication of sharing legal definitions across languages. Augmenting the original training sets with the documents translated across all languages, further improves performance (group B2 – Table 1).

3.3.3 Zero-Shot Cross-Lingual Training

We also present results in a *zero-shot cross-lingual* setting (group C – Table 1), where XML-R is trained in two languages and evaluated in the third one (unseen in fine-tuning). We observe that German has the worst performance (approx. 10% drop), which can be justified as German is a *Germanic* language, while both French and Italian are *Romance* and share a larger part of the vocabulary.

Contrarily, in case of Italian, the low-resource language in our experiments, the model strongly benefits from zero-shot cross-lingual transfer, leading to 2.2% p.p. improvement, compared to the monolingually trained XML-R. In other words, training XML-R with much more (approx 20×) out-of-language (57K in German and French) data is better compared to training on the limited (3K) in-language (Italian) documents (68.1 vs. 65.9).

3.3.4 Fine-tuning with Adapters

Across all cross-lingual settings (groups B-C – Table 1), the use of Adapters improves substantially the overall performance. The multilingual adapter-based XML-R in group B1 (Table 1) has compa-

rable performance to the NativeBERTs models of group A2, where the training dataset has been artificially augmented with machine translations. In a similar setting (group B2 – Table 1), the multilingual adapter-based XML-R in group B2 has the best overall results, combining the benefits of both cross-lingual transfer and data augmentation.

With respect to *cross-lingual performance parity*, the adapter-based XML-R model has also the highest performance parity (least diff. in the last column of Table 1), while augmenting the dataset with NMT translations leads to both the worst-case (language) performance and best performance for the least represented language (Italian).

In conclusion, cross-lingual transfer with an augmented dataset comprised of the original and machine-translated versions of all documents, has the best overall performance with a vibrant improvement (3% compared to our strong baselines – second part of Group A1 in Table 1) in Italian, the least represented language.

3.4 Cross-Domain/Regional Transfer Analysis

Further on, we examine the benefits of transfer learning (knowledge sharing) in other dimensions. Hence, we analyze model performance with respect to origin regions and legal areas (domains of law).

Legal Area	#D	Public Law	Civil Law	Penal Law	Social Law	All
Domain-specific fine-tuning with MT data augmentation						
Public Law	45.6K	<u>56.4</u> ± 2.2	52.2 ± 2.0	59.7 ± 4.9	60.1 ± 5.8	57.1 ± 3.2
Civil Law	34.5K	44.4 ± 7.9	<u>64.2</u> ± 0.6	45.5 ± 13.1	43.6 ± 5.2	49.4 ± 8.6
Penal Law	35.4K	40.8 ± 10.1	55.8 ± 2.9	84.5 ± 1.3	61.1 ± 7.5	60.6 ± 15.7
Social Law	29.1K	52.6 ± 4.2	56.6 ± 2.0	69.0 ± 5.5	<u>70.2</u> ± 2.0	62.1 ± 7.6
Cross-domain fine-tuning w/o MT data augmentation						
XLM-R	60K	57.4 ± 2.0	66.1 ± 3.1	81.4 ± 1.4	70.8 ± 2.0	68.9 ± 8.7
XLM-R + Adapters	60K	58.4 ± 2.5	66.1 ± 2.4	83.1 ± 1.2	71.1 ± 1.4	69.7 ± 9.0
Cross-domain fine-tuning with MT data augmentation						
NativeBERTs	180K	58.1 ± 3.0	64.5 ± 3.7	83.0 ± 1.3	71.1 ± 4.3	69.2 ± 9.2
XLM-R	180K	58.0 ± 3.0	67.2 ± 1.6	84.4 ± 0.2	70.2 ± 1.3	70.0 ± 9.5
XLM-R + Adapters	180K	58.6 ± 2.7	66.8 ± 2.8	83.1 ± 1.3	71.3 ± 2.4	69.9 ± 8.8

Table 3: Test results for models (XLM-R with MT unless otherwise specified) **fine-tuned** per legal area (domain) or across all legal areas (domains). Best overall results are in **bold**, and in-domain are underlined. The mean and standard deviations are computed across languages per legal area and across legal areas for the right-most column. #D is the total number of training examples. *Cross-domain transfer is beneficial for 3 out of 4 legal areas and has the best overall results*. The shared multilingual model trained across all languages and legal areas outperforms the baseline (monolingual BERT models).

3.4.1 Origin Regions

In Table 2 we present the results for *cross-regional* transfer. In the top section of the table, we present results with region-specific multilingual (XLM-R) models evaluated across regions (in-region on the diagonal, zero-shot otherwise). We observe that the cross-regional models (two lower groups of Table 2) always outperform the region-specific models. Moreover, cross-lingual transfer is beneficial across cases, while adapter-based fine-tuning further improves results in 5 out of 8 cases (regions). Data augmentation is also beneficial in most cases.

In the top part of Table 2, in 60% of the cases (regions: ZH, ES, CS, NWS, TI), a “zero-shot” model, i.e., trained in the cases of another region, slightly outperforms the in-region model. In other words, in almost every case (target region), there is another *monolingual* region-specific model that outperforms the in-region one.

We consider two main factors that may explain these results: (a) the region-wise *representational bias* considering the number of cases per region, and (b) the cross-regional *topical similarity* of the training and test subsets across different regions. To approximate the cross-regional topical similarity, we consider the distributional similarity (or dissimilarity) w.r.t. legal areas (Table 6 in Appendix C). None of these factors can fully explain

the results. Although in 3 out of 5 cases, the best performing (out-of-region) model has been trained on more data compared to the in-region one. There are also other confounding factors (e.g., language), i.e., models trained on the cases of either Espace Mittelland (EM) or Région Lémanique (RL), both bilingual with 8-10K cases, have the best results across all single-region models, hence a further exploration of the overall dynamics is needed.

3.4.2 Legal Areas

In Table 3 we present the results for *cross-domain* transfer between legal areas (domains of law). The results on the diagonal (underlined) are in-domain, i.e., fine-tuned and evaluated in the same legal area. We observe that for each domain, the models trained on in-domain data have the best results in the respective domain compared to the rest.

Interesting to note is that the best results (**bold**) are achieved in the cross-domain setting in 3 out of 4 legal areas. Such an outcome is not anticipated based on the current trends in law industry, where legal experts (judges, lawyers) over-specialize and excel in specific legal areas, e.g., criminal defense lawyers. Penal law poses the only exception where the domain-specific model is on par with the cross-domain model. Again, the results per area do not correlate with the volume of training data (*cross-*

Model	Training Dataset	#D	German \uparrow	French \uparrow	Italian \uparrow	All	(Diff. \downarrow)
Cross-lingual fine-tuning w/ or w/o MT data augmentation							
XLM-R	Original	60K	68.9 \pm 0.3	71.1 \pm 0.3	68.9 \pm 1.4	69.7 \pm 1.0	(2.2)
XLM-R + Adapters	Original	60K	<u>69.9</u> \pm 0.6	<u>71.8</u> \pm 0.7	<u>70.7</u> \pm 1.8	<u>70.8</u> \pm 0.8	(0.9)
XLM-R	+ MT Swiss	180K	70.2 \pm 0.5	71.5 \pm 1.1	<u>72.1</u> \pm 1.2	71.3 \pm 0.7	(1.9)
XLM-R + Adapters	+ MT Swiss	180K	<u>70.3</u> \pm 0.8	<u>72.1</u> \pm 0.8	<u>72.1</u> \pm 1.2	<u>71.5</u> \pm 0.9	(1.8)
Cross-jurisdiction fine-tuning w/ MT data augmentation							
XLM-R	+ MT {Swiss, Indian}	276K	70.5 \pm 0.4	71.8 \pm 0.3	73.5 \pm 1.4	72.0 \pm 0.9	(3.0)
XLM-R + Adapters	+ MT {Swiss, Indian}	276K	71.0 \pm 0.4	73.0 \pm 0.6	72.6 \pm 1.1	72.2 \pm 1.2	(2.0)
Cross-jurisdiction zero-shot fine-tuning w/ MT data augmentation							
XLM-R	MT Indian	96K	50.4 \pm 1.5	47.9 \pm 1.0	49.5 \pm 1.3	49.3 \pm 1.0	(2.5)
XLM-R + Adapters	MT Indian	96K	<u>51.6</u> \pm 2.9	<u>49.7</u> \pm 1.4	<u>50.1</u> \pm 1.4	<u>50.5</u> \pm 1.0	(1.9)

Table 4: Test results for cross-jurisdiction transfer. We present results in four settings: *standard* (Original) *augmented* (+ MT Swiss), *further augmented incl. cross-jurisdiction* (+ MT Swiss + MT Indian) and *zero-shot* (MT Indian). Best results are in **bold**. Diff. shows the difference between the best performing language and the worst performing language (max - min). **Further augmenting with translated Indian cases is overall beneficial.**

domain representational bias), and suggest that other qualitative characteristics (e.g., the idiosyncrasies of criminal law) affect the task complexity.

Similarly to the cross-regional experiments, the shared multilingual model (XLM-R) trained across all languages and legal areas with an augmented dataset outperforms the NativeBERTs models trained in a similar setting, giving another indication that the performance gains from cross-lingual transfer and data augmentation via machine translation are robust across domains as well.

3.5 Cross-Jurisdiction Transfer

We, finally, “ambitiously” stretch the limits of transfer learning in LJP and we apply *cross-jurisdiction* transfer, i.e., use of cases from different legal systems, another form of cross-domain transfer. For this purpose, we further augment the SJP dataset of FSCS cases, with cases from the Supreme Court of India (SCI), published by Malik et al. (2021).⁸ We consider and translate all (approx. 30K) Indian cases ruled up to the last year (2014) of our training dataset, originally written in English, to all target languages (German, French, and Italian).⁹

In Table 4, we present the results for two cross-jurisdiction settings: *zero-shot* (Only MT Indian), where we train XLM-R on the machine-translated

⁸Although the SCI rules under the Indian jurisdiction (law), while the FSCS under the Swiss one, we hypothesize that the fundamentals of law in two modern legal systems are quite common and thus transferring knowledge could potentially have a positive effect. We discuss this matter in Section 5.

⁹We do not use the original documents written in English, as English is not one of our target languages.

version of Indian cases, and *further augmented* (Original + MT Swiss + MT Indian), where we further augment the (already augmented) training set of Swiss cases with the translated Indian ones. While zero-shot transfer clearly fails; interestingly, we observe improvement for all languages in the further augmented setting. This opens a fascinating new direction for LJP research.

Similar to our results in Section 3.3 with respect to cross-lingual performance parity, the standard adapter-based XLM-R model has also the highest performance parity (least diff. on Table 4), while the same model trained on the fully augmented dataset leads to the worst-case (language; German) performance and best performance for the least represented language (Italian).

The cumulative improvement from all applied enhancements adds up to 7% macro-F1 compared to the XLM-R baseline and 16% to the best method by Niklaus et al. (2021) in the low-resource Italian subset, while using cross-lingual and cross-jurisdiction transfer we improve for 2.3% overall and 4.6% for Italian over our strongest baseline (NativeBERTs).

Since our experiments present several incremental improvements, we assess the stability of the performance improvements with statistical significance testing by comparing the most crucial settings in Appendix B.

4 Related Work

Legal Judgment Prediction (LJP) is the task, where given the facts of a legal case, a system

has to predict the correct outcome (legal judgment). Many prior works experimented with some forms of LJP, however, the precise formulation of the LJP task is non-standard as the jurisdictions and legal frameworks vary. Aletras et al. (2016); Medvedeva et al. (2018); Chalkidis et al. (2019) predict the plausible violation of European Convention of Human Rights (ECHR) articles of the European Court of Human Rights (ECtHR). Xiao et al. (2018, 2021) study Chinese criminal cases where the goal is to predict the ruled duration of prison sentences and/or the relevant law articles.

Another setup is followed by Sulea et al. (2017); Malik et al. (2021); Niklaus et al. (2021), which use cases from Supreme Courts (French, Indian, Swiss, respectively), hearing appeals from lower courts relevant to several fields of law (legal areas). Across tasks (datasets), the goal is to predict the binary verdict of the court (approval or dismissal of the examined appeal) given a textual description of the case. None of these works have explored neither cross-lingual nor cross-jurisdiction transfer, while the effects of cross-domain and cross-regional transfer are also not studied.

Cross-Lingual Transfer (CLT) is a flourishing topic with the application of pre-trained transformer-based models trained in a multilingual setting (Devlin et al., 2019; Lample and Conneau, 2019; Conneau et al., 2020; Xue et al., 2021) excelling in NLU benchmarks (Ruder et al., 2021). Adapter-based fine-tuning (Houlsby et al., 2019; Pfeiffer et al., 2021) has been proposed as an anti-measure to mitigate misalignment of multilingual knowledge when CLT is applied, especially in a zero-shot fashion, where the target language is unseen during training (or even pre-training).

Meanwhile, CLT is understudied in legal NLP applications. Chalkidis et al. (2021) experiment with standard fine-tuning, while they also examined the use of adapters (Houlsby et al., 2019) for zero-shot CLT on a legal topic classification dataset comprising European Union (EU) laws. They found adapters to achieve the best tradeoff between effectiveness and efficiency. Their work did not examine the use of methods incorporating translated versions of the original documents in any form, i.e., translate train documents or test ones. Recently, Xenouleas et al. (2022) used an updated, unparalleled version of Chalkidis et al. dataset to study NMT-augmented CLT methods. Other multilingual legal NLP resources (Galassi et al., 2020; Drawzeski

et al., 2021) have been recently released, although CLT is not applied in any form.

5 Motivation and Challenges for Cross-Jurisdiction Transfer

Legal systems vary from country to country. Although they develop in different ways, legal systems also have some similarities based on historically accepted justice ideals, i.e., the rule of law and human rights. Switzerland has a civil law legal system (Walther, 2001), i.e., statutes (legislation) is the primary source of law, at the crossroads between Germanic and French legal traditions.

Contrary, India has a hybrid legal system with a mixture of civil, common law, i.e., judicial decisions have precedential value, and customary, i.e., Islamic ethics, or religious law (Bhan and Rohatgi, 2021). The legal and judicial system derives largely from the British common law system, coming as a consequence of the British colonial era (1858-1947) (Singh and Kumar, 2019).

Based on the aforementioned, cross-jurisdiction transfer is challenging since the data (judgments) abide to different law standards. Although the Supreme Court of India (SCI) rules under the Indian jurisdiction (law), while the Federal Supreme Court of Switzerland (FSCS) under the Swiss one, we hypothesize that the fundamentals of law in two modern legal systems are quite common and thus transferring knowledge could potentially have a positive effect, and thus it is an experiment worth considering, while we acknowledge that from a legal perspective equating legal systems is deeply problematic, since the legislation, the case law, and legal practice are different.

Our empirical work and experimental results shows that cross-jurisdiction transfer in this specific setting (combination of Swiss and Indian decisions) has a positive impact in performance, but we cannot provide any profound hypothesis neither we are able to derive any conclusions on the importance of this finding on legal literature and practice. We leave these questions in the hands of those who can responsibly bear the burden, the legal scholars.

6 Conclusions and Future Work

6.1 Answers to the Research Questions

Following the experimental results (Section 3), we answer the original predefined research questions:

RQ1: *Is cross-lingual transfer beneficial across all or some of the languages?* In Section 3.3, we

find that vanilla CLT is beneficial in a low-resource setting (Italian), with comparable results in the rest of the languages. Moreover, CLT leveraging NMT-based data augmentation is beneficial across all languages. Overall, our experiments lead to a single multi-lingual cross-lingually “fairer” model.

RQ2: *Do models benefit or not from cross-regional and cross-domain transfer?* In Section 3.4, we find that models benefit from cross-regional transfer across all cases, since they are exposed to (trained in) many more documents (cases). We believe cross-regional diversity is not a significant aspect, compared to the importance of the increased data volume and language diversity. Cross-domain transfer is beneficial in three out of four cases (legal areas), with comparable results on penal (criminal) law, where the application of law seems to be more straight-forward / standardized (higher performing legal area). Cross-regional and cross-domain transfer lead to more robust models.

RQ3: *Can we leverage data from another jurisdiction to improve performance?* In Section 3.5, we find that cross-jurisdiction transfer in our specific setup, i.e., very similar LJP tasks, is beneficial. Again, we believe that this is mostly a matter of additional unique data (cases), rather than a matter of jurisdictional similarity. Cross-jurisdiction transfer leads to a better performing model.

RQ4: *How does representational bias (wrt. language, origin region, legal area) affect model’s performance?* We observe that representational bias – in non-extreme cases (e.g., w.r.t. language) – does not always explain performance disparities across languages, regions, or domains, and other characteristics also need to be considered.

6.2 Conclusions - Summary

We examined the application of Cross-Lingual Transfer (CLT) in Legal Judgment Prediction (LJP) for the very first time, finding a multilingually trained model to be superior when augmenting the dataset with NMT. Adapter-based fine-tuning leads to even better results. We also examined the effects of cross-domain (legal areas) and cross-regional transfer, which is overall beneficial in both settings, leading to more robust models. Cross-jurisdiction transfer by augmenting the training set with machine-translated Indian cases further improves performance.

6.3 Future Work

In future work, we would like to explore the use of a legal-oriented multilingual pre-trained model by either continued pre-training of XLM-R, or pre-training from scratch in multilingual legal corpora. Legal NLP literature (Chalkidis et al., 2022; Zheng et al., 2021) suggests that domain-specific language models positively affect performance.

In another interesting direction, we will consider other data augmentation techniques (Feng et al., 2021; Ma, 2019) that rely on textual alternations (e.g., paraphrasing, etc.). We would also like to further investigate cross-jurisdictional transfer, either exploiting data for similar LJP tasks, or via multi-task learning on multiple LJP datasets with dissimilar task specifications.

7 Ethics Statement

The scope of this work is to study LJP to broaden the discussion and help practitioners to build assisting technology for legal professionals and laypersons. We believe that this is an important application field, where research should be conducted (Tsarapatsanis and Aletras, 2021) to improve legal services and democratize law, while also highlight (inform the audience on) the various multi-aspect shortcomings seeking a responsible and ethical (fair) deployment of legal-oriented technologies.

In this direction, we study how we could better exploit all the available resources (from various languages, domains, regions, or even different jurisdictions). This combination leads to models that improve overall performance – more robust models –, while having improved performance in the worst-case scenarios across many important demographic or legal dimensions (low-resource language, worst performing legal area and region).

Nonetheless, irresponsible use (deployment) of such technology is a plausible risk, as in any other application (e.g., online content moderation) and domain (e.g., medical). We believe that similar technologies should only be deployed to assist human experts (e.g., legal scholars in research, or legal professionals in forecasting or assessing legal case complexity) with notices on their limitations.

The main examined dataset, Swiss-Judgment-Prediction (SJP), released by Niklaus et al. (2021), comprises publicly available cases from the FSCS, where cases are pre-anonymized, i.e., names and other sensitive information are redacted. The same applies for the second one, Indian Legal Documents Corpus (ILDC) of Malik et al. (2021).

Acknowledgements

This work has been supported by the Swiss National Research Program “Digital Transformation” (NRP-77)¹⁰ grant number 187477. This work is also partly funded by the Innovation Fund Denmark (IFD)¹¹ under File No. 0175-00011A. This research has been also co-financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH – CREATE – INNOVATE (T2EΔK-03849).

We would like to thank Thomas Lüthi for his legal advice, Mara Häusler for great discussions regarding the evaluation process of the models, and Phillip Rust and Desmond Elliott for providing valuable feedback on the original draft of the manuscript.

References

- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lampos. 2016. [Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective](#). *PeerJ Computer Science*, 2:e93. Publisher: PeerJ Inc.
- Ashish Bhan and Mohit Rohatgi. 2021. [Legal systems in India: Overview](#). *Thomsons Reuters - Practical Law*.
- Rich Caruana, Steve Lawrence, and C. Giles. 2001. [Overfitting in Neural Nets: Backpropagation, Conjugate Gradient, and Early Stopping](#). In *Advances in Neural Information Processing Systems*, volume 13. MIT Press.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in English](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, and Ion Androutsopoulos. 2021. [MultiEURLEX - a multi-lingual 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, Dublin, Ireland. Association for Computational Linguistics.
- Branden Chan, Timo Möller, Malte Pietsch, Tanay Soni, and Chin Man Yeung. 2019. [deepset - Open Sourcing German BERT](#).
- 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.
- Xiang Dai, Ilias Chalkidis, Sune Darkner, and Desmond Elliott. 2022. [Revisiting transformer-based models for long document classification](#).
- 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.
- Kasper Drawzeski, Andrea Galassi, Agnieszka Jablonowska, Francesca Lagioia, Marco Lippi, Hans Wolfgang Micklitz, Giovanni Sartor, Giacomo Tagiuri, and Paolo Torrioni. 2021. [A corpus for multilingual analysis of online terms of service](#). In *Proceedings of the Natural Legal Language Processing Workshop 2021*, pages 1–8, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rotem Dror, Segev Shlomov, and Roi Reichart. 2019. [Deep dominance - how to properly compare deep neural models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2785, Florence, Italy. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. [Beyond english-centric multilingual machine translation](#). *CoRR*, abs/2010.11125.
- Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edouard Hovy. 2021. [A survey of data augmentation approaches for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 968–988, Online. Association for Computational Linguistics.
- Andrea Galassi, Kasper Drazewski, Marco Lippi, and Paolo Torrioni. 2020. [Cross-lingual annotation projection in legal texts](#). In *Proceedings of the 28th*

¹⁰<https://www.nfp77.ch/en/>

¹¹<https://innovationsfonden.dk/en>

- International Conference on Computational Linguistics*, pages 915–926, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Maurizio Gotti. 2014. [Linguistic Features of Legal Texts: Translation Issues](#). *Statute Law Review*, 37(2):144–155.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for nlp](#).
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Edward Ma. 2019. [Nlp augmentation](https://github.com/makcedward/nlpaug). <https://github.com/makcedward/nlpaug>.
- Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, and Ashutosh Modi. 2021. [ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4046–4062, Online. Association for Computational Linguistics.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Karen McAuliffe. 2014. [Translating Ambiguity](#). *Journal of Comparative Law*, 9(2).
- Masha Medvedeva, Michel Vols, and Martijn Wieling. 2018. [Judicial decisions of the European Court of Human Rights: Looking into the crystal ball](#). In *Proceedings of the Conference on Empirical Legal Studies*, page 24.
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2020. [On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines](#).
- 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.
- Loreto Parisi, Simone Francia, and Paolo Magnani. 2020. [UmBERTo: an Italian Language Model trained with Whole Word Masking](#). Original-date: 2020-01-10T09:55:31Z.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Fernando Prieto Ramos. 2021. [Translating legal terminology and phraseology: between inter-systemic incongruity and multilingual harmonization](#). *Perspectives*, 29(2):175–183.
- C.D. Robertson. 2016. [Multilingual Law: A Framework for Analysis and Understanding](#). Law, language and communication. Routledge.
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Dan Garrette, Graham Neubig, and Melvin Johnson. 2021. [XTREME-R: Towards more challenging and nuanced multilingual evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10215–10245, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mahendra Pal Singh and Niraj Kumar. 2019. [Tracing the History of the Legal System in India](#). In *The Indian Legal System: An Enquiry*. Oxford University Press.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. [Rethinking the inception architecture for computer vision](#). *CoRR*, abs/1512.00567.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Dimitrios Tsarapatsanis and Nikolaos Aletras. 2021. [On the ethical limits of natural language processing on legal text](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3590–3599, Online. Association for Computational Linguistics.

Dennis Ulmer. 2021. [deep-significance: Easy and Better Significance Testing for Deep Neural Networks](https://github.com/Kaleidophon/deep-significance). <https://github.com/Kaleidophon/deep-significance>.

Fridolin M.R. Walther. 2001. [The swiss legal system a guide for foreign researchers](#). *International Journal of Legal Information*, 29(1):1–24.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Stratos Xenouleas, Alexia Tsoukara, Giannis Panagiotakis, Ilias Chalkidis, and Ion Androutsopoulos. 2022. [Realistic zero-shot cross-lingual transfer in legal topic classification](#). In *Proceedings of the 12th Hellenic Conference on Artificial Intelligence, SETN '22*, New York, NY, USA. Association for Computing Machinery.

Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. [Lawformer: A pre-trained language model for chinese legal long documents](#). *CoRR*, abs/2105.03887.

Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xi-anpei Han, Zhen Hu, Heng Wang, and Jianfeng Xu. 2018. [CAIL2018: A Large-Scale Legal Dataset for Judgment Prediction](#). *arXiv:1807.02478 [cs]*. ArXiv: 1807.02478.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Lucia Zheng, Neel Guha, Brandon R. Anderson, Peter Henderson, and Daniel E. Ho. 2021. [When does pretraining help? assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings](#). In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law, ICAIL '21*, page 159–168, New York, NY, USA. Association for Computing Machinery.

Octavia-Maria Şulea, Marcos Zampieri, Mihaela Vela, and Josef van Genabith. 2017. [Predicting the Law Area and Decisions of French Supreme Court](#)

[Cases](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 716–722, Varna, Bulgaria. INCOMA Ltd.

A Hyperparameter Tuning

We experimented with learning rates in $\{1e-5, 2e-5, 3e-5, 4e-5, 5e-5\}$ as suggested by [Devlin et al. \(2019\)](#). However, like reported by [Mosbach et al. \(2020\)](#), we also found RoBERTa-based models to exhibit large training instability with learning rate $3e-5$, although this learning rate worked well for BERT-based models. $1e-5$ worked well enough for all models. To avoid either over- or under-fitting, we use Early Stopping ([Caruana et al., 2001](#)) on development data. To combat the high class imbalance, we use oversampling, following ([Niklaus et al., 2021](#)).

We opted to use the standard Adapters of [Houlsby et al. \(2019\)](#), as the language Adapters introduced by [Pfeiffer et al. \(2020\)](#) are more resource-intensive and require further pre-training per language. We tuned the adapter reduction factor in $\{2\times, 4\times, 8\times, 16\times\}$ and got the best results with $2\times$ and $4\times$; we chose $4\times$ for the final experiments to favor less additional parameters. We tuned the learning rate in $\{1e-5, 5e-5, 1e-4, 5e-4, 1e-3\}$ and achieved the best results with $5e-5$.

We additionally applied label smoothing ([Szegedy et al., 2015](#)) on cross-entropy loss. We achieved the best results with a label smoothing factor of 0.1 after tuning with $\{0, 0.1, 0.2, 0.3\}$.

Model Type	M1	M2	M3	M4
M1: NativeBERTs	1.0	1.0	1.0	1.0
M2: NativeBERTs + MT CH	0.0	1.0	1.0	1.0
M3: XLM-R + MT CH	0.0	0.0	1.0	1.0
M4: XLM-R + MT CH + IN	0.0	0.0	0.0	1.0

Table 5: Almost stochastic dominance ($\epsilon_{\min} < 0.5$) with ASO. + *MT CH* stands for augmentation with machine translation inside the Swiss dataset and + *MT CH+IN* is the code for augmentation with machine-translations with the Swiss **and** Indian dataset.

B Statistical Significance Testing

Since our experiments present several incremental improvements, we assessed the stability of the performance improvements with statistical significance testing by comparing the most crucial settings. Using Almost Stochastic Order (ASO) ([Dror et al., 2019](#)) with a confidence level $\alpha = 0.05$, we

find the score distributions of the core models (NativeBERTs, w/ and w/o MT Swiss, XLM-R w/ and w/o MT Indian and/or Swiss) stochastically dominant ($\epsilon_{\min} = 0$) over each other in order. We compared all pairs of models based on three random seeds each using ASO with a confidence level of $\alpha = 0.05$ (before adjusting for all pair-wise comparisons using the Bonferroni correction). Almost stochastic dominance ($\epsilon_{\min} < 0.5$) is indicated in Table 5 in Appendix A. We use the deep-significance Python library of Ulmer (2021).

C Distances Between Legal Area Distributions per Origin Regions

	ZH	ES	CS	NWS	EM	RL	TI	FED
ZH	<u>.02</u>	.02	.03	.02	.01	.02	.05	.12
ES	.03	<u>.03</u>	.04	.03	.02	.01	.06	.11
CS	.02	.01	<u>.01</u>	.02	.01	.04	.06	.13
NWS	.05	.04	.06	<u>.04</u>	.04	.03	.04	.09
EM	.03	.03	.04	.02	<u>.03</u>	.03	.04	.10
RL	.06	.05	.07	.05	.05	<u>.05</u>	.04	.07
TI	.07	.07	.08	.05	.07	.08	<u>.02</u>	.06
FED	.10	.10	.12	.09	.10	.10	.06	<u>.02</u>

Table 6: Wasserstein distances between the legal area distributions of the training and the test set per origin region across languages. The training sets are in the columns and the test sets in the rows.

In Table 6 we show the Wasserstein distances between the legal area distributions of the training and the test sets per origin region across languages. Unfortunately, this analysis does not explain why the NWS model (zero-shot) outperforms the ZH model (in-domain) on the ZH test set, as found in Table 2.

D Additional Results

In Tables 7, 8, 9 and 10 we present detailed results for all experiments. All tables include both the average score across repetitions, as reported in the original tables in the main article, but also the standard deviations across repetitions.

E Responsible NLP Research

We include information on limitations, licensing of resources, and computing foot-print, as suggested by the newly introduced Responsible NLP Research checklist.

E.1 Limitations

In this appendix, we discuss core limitations that we identify in our work and should be considered in future work.

Data size fluctuations We did not control for the sizes of the training datasets, which is why we reported them in the Tables 2, 3 and 4. This mimics a more realistic setting, where the training set size differs based on data availability. Although we discussed representational bias in RQ4, we cannot completely rule out different performance based on simply more training data.

Mismatch in in/out of region model performance As described in Section 3.4.1, certain zero-shot evaluations outperform in-domain evaluations. Although we try to find an explanation for this in Section 3.4, and Appendix C, it remains an open question since there are many confounding factors.

Re-use of Indian cases Although we have empirical results confirming the statistically significant positive effect of training with additional translated Indian cases, we do not have a profound legal justification or even a hypothesis for this finding at the moment.

E.2 Licensing

The SJP dataset (Niklaus et al., 2021) we mainly use in this work is available under a CC-BY-4 license. The second dataset, ILDC (Malik et al., 2021), comprising Indian cases is available upon request. The authors kindly provided their dataset. All used software and libraries (EasyNMT, Hugging Face Transformers, deep-significance, and several other typical scientific Python libraries) are publicly available and free to use, while we always cite the original work and creators. The artifacts (i.e., the translations and the code) we created, target academic research and are available under a CC-BY-4 license.

E.3 Computing Infrastructure

We used an NVIDIA GeForce RTX 3090 GPU with 24 GB memory for our experiments. In total, the experiments took approx. 80 GPU days, excluding the translations. The translations took approx. 7 GPU days per language from Indian to German, French, and Italian. The translation within the Swiss corpus took approx. 4 GPU days in total.

Legal Area	#D	Public Law	Civil Law	Penal Law	Social Law	All
Public Law	45.6K	<u>56.4</u> \pm 2.2	52.2 \pm 2.0	59.7 \pm 4.9	60.1 \pm 5.8	57.1 \pm 3.2
Civil Law	34.5K	44.4 \pm 7.9	<u>64.2</u> \pm 0.6	45.5 \pm 13.1	43.6 \pm 5.2	49.4 \pm 8.6
Penal Law	35.4K	40.8 \pm 10.1	55.8 \pm 2.9	84.5 \pm 1.3	61.1 \pm 7.5	60.6 \pm 15.7
Social Law	29.1K	52.6 \pm 4.2	56.6 \pm 2.0	69.0 \pm 5.5	<u>70.2</u> \pm 2.0	62.1 \pm 7.6
<i>All</i>	60K	58.0 \pm 3.0	67.2 \pm 1.6	84.4 \pm 0.2	70.2 \pm 1.3	70.0 \pm 9.5
<i>All (w/o MT)</i>	60K	57.4 \pm 2.0	66.1 \pm 3.1	81.4 \pm 1.4	70.8 \pm 2.0	68.9 \pm 8.7
<i>All (Native)</i>	60K	58.1 \pm 3.0	64.5 \pm 3.7	83.0 \pm 1.3	71.1 \pm 4.3	69.2 \pm 9.2

Table 7: Test results for models (XLM-R with MT unless otherwise specified) **fine-tuned** per legal area (domain) or across all legal areas (domains). Best overall results are in **bold**, and in-domain are underlined. **Cross-domain transfer is beneficial for 3 out of 4 legal areas and has the best overall results.** The shared multilingual model trained across all languages and legal areas outperforms the baseline (monolingual BERT models). The mean and standard deviations are computed across languages per legal area and across legal areas for the right-most column. #D is the number of training examples per legal area.

Legal Area	#D	Public Law	Civil Law	Penal Law	Social Law	All
Public Law	45.6K	<u>57.2</u> \pm 1.8	53.8 \pm 2.1	58.9 \pm 5.2	61.7 \pm 4.1	57.9 \pm 2.9
Civil Law	34.5K	41.4 \pm 6.6	<u>57.6</u> \pm 1.1	42.8 \pm 9.1	43.0 \pm 4.1	46.2 \pm 6.6
Penal Law	35.4K	37.4 \pm 12.8	56.4 \pm 2.0	86.3 \pm 0.1	61.6 \pm 6.7	60.4 \pm 17.4
Social Law	29.1K	51.4 \pm 5.8	54.8 \pm 2.8	73.9 \pm 1.9	<u>70.3</u> \pm 2.2	62.6 \pm 9.7
<i>All</i>	60K	58.6 \pm 2.7	66.8 \pm 2.8	83.1 \pm 1.3	71.3 \pm 2.4	69.9 \pm 8.8
<i>All (w/o MT)</i>	60K	58.4 \pm 2.5	66.1 \pm 2.4	83.1 \pm 1.2	71.1 \pm 1.4	69.7 \pm 9.0

Table 8: Test results for models (XLM-R with MT unless otherwise specified) **adapted** per legal area (domain) or across all legal areas (domains). Best overall results are in **bold**, and in-domain are underlined. The mean and standard deviations are computed across languages per legal area and across legal areas for the right-most column. #D is the number of training examples per legal area.

Region	#D	#L	ZH	ES	CS	NWS	EM	RL	TI	FED	All
ZH	26.4K	de	<u>65.5 ± 0.0</u>	65.6 ± 0.0	63.7 ± 0.0	68.2 ± 0.0	62.0 ± 2.9	57.9 ± 6.7	63.2 ± 0.0	54.8 ± 5.1	62.6 ± 4.1
ES	17.1K	de	62.9 ± 0.0	<u>66.9 ± 0.0</u>	62.8 ± 0.0	65.2 ± 0.0	62.2 ± 1.1	60.2 ± 5.3	57.8 ± 0.0	55.1 ± 6.3	61.6 ± 3.6
CS	14.4K	de	62.5 ± 0.0	65.5 ± 0.0	<u>63.2 ± 0.0</u>	65.1 ± 0.0	60.7 ± 1.6	57.8 ± 3.7	60.5 ± 0.0	55.9 ± 0.5	61.4 ± 3.1
NWS	17.1K	de	66.0 ± 0.0	68.6 ± 0.0	65.2 ± 0.0	<u>67.9 ± 0.0</u>	61.6 ± 1.7	57.0 ± 4.9	57.1 ± 0.0	55.5 ± 5.7	62.4 ± 4.9
EM	24.9K	de,fr	64.1 ± 0.0	66.6 ± 0.0	63.3 ± 0.0	66.7 ± 0.0	<u>64.0 ± 0.7</u>	66.8 ± 2.9	63.2 ± 0.0	58.4 ± 0.3	64.1 ± 2.6
RL	40.2K	fr,de	61.0 ± 0.0	64.7 ± 0.0	60.2 ± 0.0	63.7 ± 0.0	63.4 ± 3.3	<u>69.8 ± 2.7</u>	67.6 ± 0.0	54.3 ± 7.2	63.1 ± 4.4
TI	6.9K	it	55.0 ± 0.0	56.3 ± 0.0	53.2 ± 0.0	54.5 ± 0.0	56.0 ± 0.4	54.7 ± 0.9	<u>66.0 ± 0.0</u>	53.1 ± 6.4	56.1 ± 3.9
FED	3.9K	de,fr,it	57.5 ± 0.0	59.6 ± 0.0	56.8 ± 0.0	58.9 ± 0.0	55.0 ± 1.0	56.5 ± 1.1	53.5 ± 0.0	<u>54.9 ± 2.9</u>	56.6 ± 1.9
<i>All</i>	60K	de,fr,it	69.2 ± 0.0	72.9 ± 0.0	68.3 ± 0.0	73.3 ± 0.0	69.9 ± 1.6	71.7 ± 2.8	70.4 ± 0.0	65.0 ± 3.9	70.1 ± 2.5
<i>All (w/o MT)</i>	60K	de,fr,it	68.5 ± 0.0	71.3 ± 0.0	67.7 ± 0.0	71.2 ± 0.0	69.0 ± 1.5	71.4 ± 0.3	67.4 ± 0.0	64.6 ± 5.2	68.9 ± 2.2
<i>All (Native)</i>	60K	de,fr,it	69.0 ± 0.0	72.1 ± 0.0	68.6 ± 0.0	72.0 ± 0.0	69.9 ± 1.6	71.9 ± 0.7	68.8 ± 0.0	64.8 ± 7.0	69.6 ± 2.3

Table 9: Test results for models (XLM-R with MT unless otherwise specified) **fine-tuned** per region (domain) or across all regions (domains). Best overall results are in **bold**, and in-domain are underlined. The mean and standard deviations are computed across languages per origin region and across origin regions for the right-most column. The regions where only one language is spoken thus show std 0. #D is the number of training examples per origin region. #L are the languages covered.

Region	#D	#L	ZH	ES	CS	NWS	EM	RL	TI	FED	All
ZH	26.4K	de	65.4 ± 0.0	68.7 ± 0.0	63.9 ± 0.0	68.2 ± 0.0	63.6 ± 3.5	61.0 ± 2.8	66.4 ± 0.0	56.3 ± 1.8	64.2 ± 3.8
ES	17.1K	de	64.2 ± 0.0	69.4 ± 0.0	63.9 ± 0.0	66.0 ± 0.0	61.7 ± 2.3	59.4 ± 4.6	61.2 ± 0.0	56.5 ± 6.1	62.8 ± 3.7
CS	14.4K	de	63.1 ± 0.0	66.5 ± 0.0	64.1 ± 0.0	65.0 ± 0.0	61.0 ± 2.6	57.5 ± 2.1	62.2 ± 0.0	56.7 ± 2.5	62.0 ± 3.2
NWS	17.1K	de	65.8 ± 0.0	69.0 ± 0.0	63.8 ± 0.0	67.4 ± 0.0	59.9 ± 3.3	58.6 ± 1.1	58.9 ± 0.0	54.2 ± 2.7	62.2 ± 4.8
EM	24.9K	de,fr	63.9 ± 0.0	67.5 ± 0.0	64.4 ± 0.0	66.8 ± 0.0	64.7 ± 0.5	69.1 ± 1.7	66.4 ± 0.0	59.5 ± 1.0	65.3 ± 2.7
RL	40.2K	fr,de	62.3 ± 0.0	66.2 ± 0.0	62.0 ± 0.0	64.7 ± 0.0	65.2 ± 4.2	70.8 ± 6.8	65.5 ± 0.0	56.9 ± 6.0	64.2 ± 3.7
TI	6.9K	it	56.4 ± 0.0	62.1 ± 0.0	53.7 ± 0.0	56.3 ± 0.0	55.1 ± 0.2	57.4 ± 1.1	68.3 ± 0.0	50.5 ± 2.3	57.5 ± 5.1
FED	3.9K	de,fr,it	52.7 ± 0.0	52.7 ± 0.0	51.3 ± 0.0	53.1 ± 0.0	52.8 ± 0.7	52.0 ± 2.3	52.8 ± 0.0	50.0 ± 4.0	52.2 ± 1.0
<i>All</i>	60K	de,fr,it	69.2 ± 0.0	73.3 ± 0.0	69.9 ± 0.0	73.0 ± 0.0	70.3 ± 1.9	72.1 ± 0.7	70.9 ± 0.0	63.8 ± 6.1	70.3 ± 2.8
<i>All (w/o MT)</i>	60K	de,fr,it	69.2 ± 0.0	73.9 ± 0.0	67.9 ± 0.0	72.6 ± 0.0	69.0 ± 2.1	72.1 ± 0.3	70.1 ± 0.0	64.2 ± 4.6	69.9 ± 2.9

Table 10: Test results for models (XLM-R with MT unless otherwise specified) **adapted** per region (domain) or across all regions (domains). Best overall results are in **bold**, and in-domain are underlined. The mean and standard deviations are computed across languages per origin region and across origin regions for the right-most column. The regions where only one language is spoken thus show std 0. #D is the number of training examples per origin region. #L are the languages covered.