

Towards Explainability and Fairness in Swiss Judgement Prediction: Benchmarking on a Multilingual Dataset

Santosh T.Y.S.S¹, Nina Baumgartner², Matthias Stürmer^{2,3},
Matthias Grabmair¹, Joel Niklaus^{2,3,4}

¹Technical University of Munich, ²University of Bern,
³Bern University of Applied Sciences, ⁴Stanford University

Abstract

The assessment of explainability in Legal Judgement Prediction (LJP) systems is of paramount importance in building trustworthy and transparent systems, particularly considering the reliance of these systems on factors that may lack legal relevance or involve sensitive attributes. This study delves into the realm of explainability and fairness in LJP models, utilizing Swiss Judgement Prediction (SJP), the only available multilingual LJP dataset. We curate a comprehensive collection of rationales that ‘support’ and ‘oppose’ judgement from legal experts for 108 cases in German, French, and Italian. By employing an occlusion-based explainability approach, we evaluate the explainability performance of state-of-the-art monolingual and multilingual BERT-based LJP models, as well as models developed with techniques such as data augmentation and cross-lingual transfer, which demonstrated prediction performance improvement. Notably, our findings reveal that improved prediction performance does not necessarily correspond to enhanced explainability performance, underscoring the significance of evaluating models from an explainability perspective. Additionally, we introduce a novel evaluation framework, Lower Court Insertion (LCI), which allows us to quantify the influence of lower court information on model predictions, exposing current models’ biases.

Keywords: Fairness, Explainability, Multilingual, Legal Judgement Prediction

1. Introduction

The task of Legal Judgement Prediction involves analyzing the textual description of case facts to determine various aspects of a case’s outcome, such as the winning party, violated provisions, and motion results. It has garnered substantial attention in the mainstream NLP community (Aletras et al., 2016; Chalkidis et al., 2019; Malik et al., 2021; Niklaus et al., 2021; Semo et al., 2022; Santosh et al., 2023a) and is being considered as a benchmarking task for evaluating the capabilities of legal NLP (Chalkidis et al., 2022b; Niklaus et al., 2023a) and long range models (Condevaux and Harispe, 2022; Niklaus and Giofré, 2022; Chalkidis et al., 2022a; Hua et al., 2022; Niklaus et al., 2023b).

The process of resolving legal cases encompasses evidential reasoning through exchange of arguments between the litigating parties before a decision-making body (Santosh et al., 2022). Earlier methods to deal with outcome prediction task such as IBP (Brüninghaus and Ashley, 2003), SMILE+IBP (Brüninghaus and Ashley, 2005), VJF (Grabmair, 2017) typically involved identification/extraction of the factors from the textual description of the facts, then employing a conceptual schema to relate the factors to legal issues and predicts the outcome by comparing them with the past cases, thus providing the explanations for those predictions in terms that are legally intuitive. However, in the context of modern deep learning-based solutions, the outcome is determined solely from the text of the case facts, effectively bypass-

ing the interpretable legal reasoning process. This poses a significant risk, particularly in high-stakes domains like law, when utilizing such systems that rely on factors that may be predictive but lack legal relevance or involve sensitive attributes (e.g., the race of an accused person). Such reliance can lead to unjust and biased outcomes, undermining the principles of fairness and equal treatment within the legal system. Hence, such systems need to be analyzed from an explainability standpoint, thus making them transparent and thereby enhancing the trust of legal practitioners and stakeholders to comprehend the factors and legal principles that contribute to a particular prediction.

In the line of explainable LJP, Chalkidis et al. 2021 investigated the rationales behind models’ decisions in Legal Judgment Prediction (LJP) for European Court of Human Rights (ECtHR) cases. Subsequent studies by Santosh et al. 2022 extended the above dataset and Malik et al. 2021 created new dataset for Indian Jurisdiction. In contrast to these works in English, our study focuses on assessing the explainability of LJP models trained on the Swiss-Judgment-Prediction (SJP) dataset, which is the only available multilingual LJP dataset. It contains cases from the Federal Supreme Court of Switzerland (FSCS), written in three official Swiss languages (German, French, Italian)¹. To this end, we curate a multilingual set of rationales that ‘support and ‘oppose’ Judgment

¹The dataset consists of non-parallel cases, with each case being unique and decisions being written in a single language.

on 108 cases in German, French and Italian collectively. We employ a perturbation-based explainability approach, namely Occlusion (Zeiler and Fergus, 2014), wherein we remove the factors from the fact statements and measure the change in the prediction confidence in comparison to a non-occluded baseline. This occlusion based method facilitates to identify the contribution of each factor in arriving at the final prediction, which also links to the characteristics of earlier factor based formal methods of LJP which are known for their interpretability. To enable a fair comparison across methods, we release four distinct occlusion test sets². Each test set involves occluding a different number of sentences (1, 2, 3, and 4) per experiment. This comprehensive range of occlusion scenarios allows us to assess the impact of varying levels of factor removal on the prediction outcomes.

Using the occluded datasets, we assess the explainability performance of state-of-the-art models developed for SJP task using both monolingual (Niklaus et al., 2021) and multilingual BERT (Niklaus et al., 2022) architectures, as well as models developed with techniques such as data augmentation and cross-lingual transfer (Niklaus et al., 2022). Our findings highlight the fact that the prediction performance improvement does not translate to explainability improvement.

Furthermore we leverage the peculiar characteristics of the Federal Supreme Court of Switzerland (FSCS), which handles only the most contentious cases that lower courts have struggled to resolve adequately. In their decisions, the FSCS often focuses on specific portions of previous decisions, scrutinizing potential flaws in the lower court’s reasoning. This setup offers an interesting testbed to systematically assess the bias of the lower court in the final predictions generated by our models. This approach is reminiscent of recent works (Chalkidis et al., 2022c; Wang et al., 2021) that have examined the fairness of LJP models by examining group fairness or disparate impact i.e., performance disparities across various attributes, such as gender, age, and region. Our approach, termed Lower Court Insertion (LCI)³, adopts a counterfactual fairness perspective, unlike prior studies examining performance disparities in LJP models. This involves extracting instances of the lower court in each case document and replacing them with other lower courts to measure the resulting changes in prediction confidence

²Our Occlusion dataset is available at https://huggingface.co/datasets/rcds/occlusion_swiss_judgment_prediction

³Our LCI dataset is available at https://huggingface.co/datasets/rcds/lower_court_insertion_swiss_judgment_prediction

scores. Remarkably, despite the lower court’s average length being only 7 words in documents with an average length of 350 words, it has shown the potential to flip the prediction label in some cases.

In sum, our main contributions are as follows:

- We release a new dataset of 108 cases from a trilingual Switzerland Judgment Prediction corpus with rationales annotated by experts to assess the explainability of SJP models.
- We evaluate the state-of-the-art models developed for the SJP task, including monolingual and multilingual models and models trained with several techniques, from an explainability standpoint using the occlusion technique.
- We perform systematic evaluation of lower court bias embodied in these models using the LCI technique, allowing us to quantify the influence of the lower court on the final predictions generated by the models.

2. Related Work

Legal Judgement Prediction: LJP has been studied under various jurisdictions such as the European Court of Human Rights (ECtHR) (Chalkidis et al., 2019; Aletras et al., 2016; Liu and Chen, 2017; Medvedeva et al., 2018, 2021; Santosh et al., 2022, 2023b,a; Chalkidis et al., 2022b, 2021; Kaur and Bozic, 2019) Chinese Criminal Courts (Luo et al., 2017; Yue et al., 2021; Zhong et al., 2020), US Supreme Court (Katz et al., 2017; Kaufman et al., 2019), Indian Supreme Court (Malik et al., 2021; Shaikh et al., 2020) the French Court of Cassation (Şulea et al., 2017b,a), Brazilian courts (Lage-Freitas et al., 2022; Bertalan and Ruiz, 2020), the Turkish Constitutional Court (Sert et al., 2021; Mumcuoğlu et al., 2021) UK courts (Strickson and De La Iglesia, 2020), German courts (Wattl et al., 2017), the Philippine Supreme Court (Virtucio et al., 2018), the Thailand Supreme Court (Kowsrihawat et al., 2018) and the Federal Supreme Court of Switzerland (Niklaus et al., 2021, 2022; Rasiyah et al., 2023) – the only publicly available multi-lingual LJP corpus – which is the main focus of this work.

Swiss Judgement Prediction (SJP): Niklaus et al. 2021 evaluate different methods for the LJP task on the Swiss-Judgment-Prediction (SJP) dataset. They achieve the best performance using a hierarchical variant of BERT that overcomes the token input limitation. Niklaus et al. 2022 further enhance the performance through cross-lingual transfer learning, adapter-based fine-tuning and data augmentation using machine translation. In contrast to previous works, this study examines

the explainability of these models and investigates if improved prediction performance translates into improved explainability performance.

Explainability: Explanations in Explainable Artificial Intelligence (XAI) methods are classified based on two factors: whether the explanation is for an individual prediction or the overall prediction process (local or global), and whether the explanation is derived directly from the prediction process or requires post-processing (self-explaining or post-hoc) (Danilevsky et al., 2020). These methods can be model-agnostic (LIME (Ribeiro et al., 2016), SHAP (Lundberg and Lee, 2017), Occlusion (Li et al., 2016; Zeiler and Fergus, 2014), Anchors (Ribeiro et al., 2018)), applicable to any model, or model-specific (Integrated Gradients (Sundararajan et al., 2017), Gradient Saliency, and Attention-Based Methods), designed for specific models or architectures. In this study, we use occlusion, a model-agnostic, local, and post-hoc explainability technique.

Fairness: Fairness in machine learning has been defined in different ways to address various types of discrimination. These definitions include group fairness, individual fairness, and causality-based fairness. Group fairness ensures equitable predictions across demographic subgroups, avoiding differential treatment based on attributes such as race, gender, or age (Zafar et al., 2017; Hardt et al., 2016). Individual fairness focuses on treating similar individuals similarly, avoiding arbitrary distinctions based on their characteristics (Sharifi-Malvajardi et al., 2019; Yurochkin et al.). Causality-based fairness considers underlying causal mechanisms and aims to identify and mitigate biases caused by confounding variables or indirect discrimination (Wu et al., 2019; Zhang and Bareinboim, 2018). In this study, we examine bias related to the lower court using counterfactual and causal fairness estimation methods.

Explainability and Fairness in LJP: Early works in the field of legal judgment prediction, such as HYPO (Rissland and Ashley, 1987), CATO (Aleven and Ashley, 1997), IBP (Brüninghaus and Ashley, 2003) and IBP+SMILE (Brüninghaus and Ashley, 2005), relied on symbolic AI techniques to incorporate domain knowledge and provide interpretable explanations for the outcomes. However, deep learning models in LJP have prioritized prediction performance over explainability. Nevertheless, recent research emphasizes the significance of explainability in the legal domain for trust and the right to explanation principle. Efforts have been made to investigate explainability in LJP. For instance, Chalkidis et al. 2021 introduced the task

of rationale extraction from facts statements and released a dataset from ECtHR. They used neural models with regularization constraints to select rationales using a learned binary mask. Additionally, Santosh et al. 2022 identified distractor words highly correlated with outcomes but not legally relevant, and proposed an adversarial deconfounding procedure to align model explanations with those chosen by legal experts. Xu et al. 2023 analyzed the token-level alignment with LJP models on the ECtHR corpus. Similarly, Malik et al. 2021 developed a dataset of Indian jurisdiction for explainability assessment using the occlusion method. In this work, we curate a dataset from the trilingual Swiss jurisdiction, and employ the occlusion method to evaluate models for Swiss Judgment Prediction.

Fair machine learning in the legal domain is a relatively new field. Studies such as Angwin et al. 2016 identified racial bias in the COMPAS system, a parole risk assessment tool in the US, where black individuals were more likely to be mislabeled as high risk. Another study by Wang et al. 2021 found significant fairness gaps across gender in LSTM-based models for legal judgment consistency using a dataset of Chinese criminal cases. Recently, Chalkidis et al. 2022c developed the FairLex benchmark to facilitate research on bias mitigation algorithms in the legal domain. It includes four datasets from different jurisdictions and languages, covering various sensitive attributes. While previous works focused on group fairness and quantifying predictions across demographic subgroups, this study examines a specific variable of lower court from a counterfactual perspective.

3. Occlusion & LCI Dataset for SJP

The SJP dataset (Niklaus et al., 2021) comprises 85,000 cases from the Federal Supreme Court of Switzerland (FSCS) spanning the years 2000 to 2020, chronologically split into training (2000-14), validation (2015-16) and test (2017-20) splits and are written in three languages: German, French, and Italian. However, it is important to note that the dataset is not evenly distributed among these languages with Italian having a much smaller number of documents (4K) compared to German (50K) and French (31k). Additionally, this representation disparity is also evident across various legal areas and regions. For more detailed dataset statistics, please refer to the work of Niklaus et al. 2021.

3.1. Rationale and Lower Court Annotation

We sample a total of 108 cases from both the validation and test sets (2015-20). These cases were equally distributed across the three languages.

Within each year of the validation and test sets, we sampled six cases per language, resulting in two cases per legal area. Specifically, each legal area in every year contained one case with the judgment "approved" and one with the judgment "dismissed." It is worth noting that our annotation dataset is balanced in terms of final outcomes and languages, in contrast to the SJP dataset, which contains a majority of dismissed cases (> 3/4). The annotations were conducted by a team of three legal experts, consisting of two law students pursuing their master's degrees and one lawyer, over a period of five months. Two legal experts are native German speakers with intermediate knowledge in French and basic Italian skills. The third expert is a native speaker in German and Italian and fluent in French. The annotation was facilitated using the Prodigy tool ⁴.

The annotation task was to highlight sentences or sub-sentences in the facts section of the judgment that "support" or "oppose" the final outcome of the case. We have chosen sub-sentences as the atomic unit for annotation after consulting with legal experts who expressed that a sentence can contain two sub-sentences opposing each other and hence should be annotated with different labels. The annotators had been given access to the entire case to make their annotation instead of just the facts section, which is the actual input for the models dealing with judgment prediction task. These decisions have been taken to address two points: (i) Experts opined that sentences/sub-sentences may have opposing labels depending on how the court interpreted those facts in its reasoning; hence providing them the entire case would greatly assist them in arriving at explanations leading to higher inter-annotator agreement (ii) Having prior knowledge about a specific case allows an expert to be familiar with its specific legal and factual details, as well as the court's opinions on the matter. As a result, varying levels of prior familiarity with a case can lead to different interpretations and perspectives in understanding it. Hence providing the entire case levels the playing field and eliminates the possibility that some cases are known to only some experts before, possibly leading to different annotations.

The experts are instructed to read through the facts, the considerations, the ruling, and any other needed legal document (such as relevant legislation, analyses or case law) to understand the court case and then annotate the rationale. Unlike the previous works involving explainability annotations in LJP (Chalkidis et al., 2021; Santosh et al., 2022; Malik et al., 2021) which only collect rationales that help to arrive at the final outcome of the case, we introduce and collect rationales at fine-grained level

⁴<https://prodi.gy>

termed as "Supports Judgement" and "Opposes Judgment" which holds significance especially in the task of judgment prediction due to the inherent nature of legal text of often operating within the realm of gray areas rather than clear-cut black-and-white distinctions. Legal cases involve complex issues, conflicting facts leading to alternative legal reasoning, dissenting opinions, alternative interpretations of the law and can serve as potential grounds for challenging the ruling and can serve as a reference point for legal arguments or considerations. Thus, including the fine granularity of labels provides an opportunity to assess more nuanced understanding of the case by the models, acknowledging that legal decisions are not always unanimous and different perspectives may exist within the legal community.

Additionally, we request annotators to label neutral sentences. This is not a label per se, but covers sentences not assigned other labels, as this assists in implementing the occlusion method to partition the facts section into more coherent sentences with minimal effort, as segmenting legal text is a complex task in itself (Read et al., 2012; Savelka et al., 2017; Brugger et al., 2023).

In addition to sentences and sub-sentences indicating towards outcome explanations, we also ask annotators to label the lower court mentions in the fact section as indicated in the rubrum (header including identifiers, and listing judges, lawyers and involved parties) of the ruling.

The annotation task was conducted in two cycles to ensure high quality. The initial cycle involved pilot annotations, highlighting uncertainties regarding guidelines. As a result, we refined the guidelines by providing more precise instructions to address these concerns⁵. Subsequently, a discussion among the legal experts was held to resolve any conflicts and consolidate the annotations in the most effective manner, thereby ensuring the high quality of the annotations.

3.2. Inter Annotator Agreement

We obtained annotations from three annotators only for the German subset. Detailed distribution of labeled tokens per annotator can be found in Fig. 1. Among the three, Annotator 1 annotated the least amount of tokens. Annotator 3 annotated the most comparable to the Annotator 2, especially when using the Supports Judgment label. To measure inter-annotator agreement for explanations, we use the machine translation metrics as suggested by Malik et al. 2021 like ROUGE-1, ROUGE-2, ROUGE-L (Lin, 2004), BLEU (Papineni et al., 2002) (unigram and bigram averag-

⁵Our detailed annotation guidelines and discussions are available [here](#).




Figure 1: Mean number of tokens annotated per label per annotator in German subset

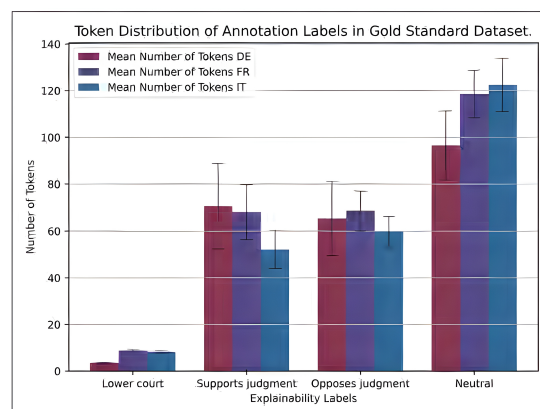


Figure 2: Distribution of the number of tokens per label in the final dataset across each language.

ing), METEOR (Agarwal and Lavie, 2007), Jaccard Similarity, Overlap Maximum, Overlap Minimum and BERTScore (Zhang et al., 2019). We report the inter-annotator agreement scores in the German subset for the first round of annotations of the German dataset in Table 1. These scores are aggregated over all the labels (supports, opposes judgment and lower court). Table 1 demonstrates high agreement across all scores, with values ranging from 0.7 to 0.9. The high BERTScore indicates strong similarities in non-lexical matches, while the indication of OVERLAP Minimum suggests that the annotations frequently overlapped as subsequences. Notably, Experts 2 and 3 exhibit the highest agreement, which can be attributed to their larger number of annotated tokens compared to Expert 1 as observed in Fig. 1. We also notice that the agreement within the categories "Lower Court" and "Supports Judgment" is notably high in comparison to "Opposes Judgment". The experts confirmed that the higher variance in the "Opposes Judgment" label stemmed from the difficulty in identifying these sentences and resolving these conflicts constituted a significant effort in landing with final annotations. Distribution of final number of tokens obtained per label across language is visualised in Fig. 2.

3.3. Occlusion and LCI dataset

To evaluate the explainability of models and enable a fair comparison among them, we derive four distinct occlusion based datasets from the test split⁶ of above annotated rationales data, consisting of 27, 24 and 23 cases in German, French and Italian respectively. For each occlusion test set, we occlude a different number of sentences (1, 2, 3, and 4)

⁶We exclude the instances from the validation split, which is used for hyperparameter tuning during model training, to derive the occluded test set for explainability.

IAA metric	A1-A2	A1-A3	A2-A3
Rouge-1	0.78	0.69	0.87
Rouge-2	0.74	0.64	0.85
Rouge-L	0.77	0.68	0.87
BLEU	0.75	0.69	0.85
METEOR	0.77	0.71	0.88
Jaccard Sim.	0.73	0.64	0.82
Overlap Max.	0.68	0.61	0.74
Overlap Min.	0.83	0.73	0.81
BERTScore	0.91	0.86	0.93

Table 1: IAA score between the annotators in the first cycle for German subset

belonging to same label (Supports/Opposes Judgment/Neutral) per experiment in a case, adding no marker or trace of the occlusion in the fact section to leave it as similar and natural as possible. For every occlusion test instance, we also pair it with a baseline with no text occluded. Thus, we arrive in a total of 28k occluded instances with varying levels of occlusion, across three languages. Using these occluded instances, we analyze the difference in prediction confidence in comparison to the non-occluded baseline.

For LCI, we derive the counterfactual based test set wherein we use the lower court instances annotated by the annotator and replace the lower court instances with other lower court names in every case resulting in a total of 1127 instances. There are a total of 13, 9, and 16 unique lower court instances in German, French, and Italian respectively. Similar to above, each instance is also paired with a baseline representing the case text with the actual lower court name without any replacement, which we use to analyze the change in prediction confidence.

Tables 2 and 3 provide statistics on the total number of instances in both the Occlusion and LCI test sets along with detailed breakdown of the number

		DE	FR	IT
#Documents		27	24	23
Set 1	Opposes	55	34	31
	Neutral	247	164	195
	Supports	98	85	50
Set 2	Opposes	66	22	23
	Neutral	1097	586	827
	Supports	203	246	69
Set 3	Opposes	53	7	8
	Neutral	3158	1260	2429
	Supports	356	659	56
Set 4	Opposes	27	1	1
	Neutral	6622	1801	5704
	Supports	586	1477	28

Table 2: Split of number of instances per label in each occluded test across three languages.

	Occlusion				LCI
	Opp.	Neu.	Sup.	Total	Total
DE	201	11325	1243	12769	351
FR	64	3875	2467	6406	391
IT	63	9218	203	9484	312

Table 3: Total number of instances per label across three languages for occlusion and LCI test. Opp., Neu., Sup. represent ‘Opposes Judgement’, ‘Neutral’ and ‘Supports Judgement’ respectively.

of instances in each occlusion set by language. Across languages, the German subset comprises the largest portion of the test set compared to the rest. This is due to annotation of fewer sentences in Italian and French documents as can be noticed in Fig. 2. Among the labels, the ‘Opposes Judgement’ label has fewer instances, due to the lower number of annotated tokens associated with this label.

4. Experimental Setup

4.1. Models

We assess the following six classes of models, developed on the backbone of hierarchical BERT, developed for the SJP task in previous literature (Niklaus et al., 2021, 2022). We follow the same dataset splits provided by Niklaus et al. 2021 for training and validation. Hierarchical BERT is employed because the SJP dataset includes documents with more than 512 tokens. In this approach, the text is split into 4 consecutive blocks of 512 tokens (90% of cases are less than 2048 tokens) and fed into a shared standard

BERT encoder independently. Then the CLS token of each block is passed through a 2-layer transformer encoder to aggregate the information across blocks, followed by max-pooling and a final classification layer.

MonoLingual: This variant uses monolingually pre-trained BERT models i.e German-BERT (Chan et al., 2019), CamemBERT (Martin et al., 2020) and UmBERTo (Parisi et al., 2020) for German, French and Italian. Each model is fine-tuned and evaluated using that language subset dataset.

MultiLingual: This variant uses the multilingually pre-trained XLM-R model (Conneau et al., 2019) instead of language-specific pre-trained BERT. However, the fine-tuning process is still performed separately for each language, similar to the monolingual approach.

Mono/Multi Lingual with Data Augmentation:

We translate the cases in SJP dataset into other languages from the original language using the EasyNMT2 framework, following the approach proposed by Niklaus et al. 2022. Then these translated instances are then augmented with the original data for a specific language during the fine-tuning process with Mono/Multilingual BERT. This is similar to above experiment in setup, with the main distinction being the additional augmented data.

Joint Training without/with Data Augmentation:

We use a multilingual pre-trained model and fine-tune it across all the three language corpora jointly, which tries to capitalize on the inherited benefit of using larger multilingual corpora during fine-tuning. As discussed above, we translate each document from its original language to the rest of other two and train a data augmented version of models jointly with all the obtained translated data. Unlike the previous approaches where separate models were fine-tuned for each language, this method jointly fine-tunes on all languages, resulting in a single final model instead of multiple models for each language.

4.2. Implementation Details

We use the code repositories from prior work (Niklaus et al., 2021, 2022) to assess the state-of-the-art models on SJP⁷. We employ a learning rate of 1e-5 with early stopping based on macro-F1 on the development set. All models are trained with a batch size of 64 for 10 epochs using AdamW optimizer with mixed precision and gradient ac-

⁷<https://github.com/JoelNiklaus/SwissJudgementPrediction>

Model	German	French	Italian
MonoLingual	69.08	71.78	67.82
MultiLingual	67.92	69.24	65.28
MonoLingual + DA	70.47	71.24	69.21
MultiLingual + DA	68.94	71.06	69.86
Joint Training	68.74	70.82	70.62
Joint Training + DA	70.58	71.62	71.18

Table 4: Prediction Performance on Test set of Niklaus et al. 2021

cumulation using huggingface library (Wolf et al., 2020). We use oversampling to handle class imbalance. We use 4 segments with 512 tokens each in our hierarchical models resulting in a maximum sequence length of 2048.

4.3. Metrics

We report macro-F1 following Niklaus et al. 2021, 2022 for assessing prediction performance. For assessing explainability through occlusion experiments, we calculate the explainability score S_{exp} for every test instance as the difference between the temperature-scaled⁸ confidence of the baseline and the occluded instance. (i.e., baseline - occluded). A negative (positive) S_{exp} score indicates that occluded text is opposing (supporting) its prediction. Then, we assign the label 'Opposes Judgement'/'Neutral'/'Supports Judgement' based on the sign of explainability score. Finally, we report F1-score for each of the labels across all the occluded instances.

In bias estimation using the LCI method, we calculate an explainability score for each instance. As the explainability scores are sign dependent, we separately compute the Mean of Explainability Scores (MES), for positive and negative values, expressed as a percentage. A positive explainability score indicates that the insertion of the lower court decreases the probability, suggesting that the inserted court has a pro-dismissal influence. Conversely, a negative score indicates an increase in the probability, indicating a pro-approval trend of the inserted lower court. For an ideally unbiased model, the presence of the lower court should not affect the probability of the prediction. Therefore, a value of the mean explainability score closer to 0 is desirable. Additionally, we report the percentage of cases where the insertion of the lower court leads to a flip in the label of the prediction, changing it from 0 to 1 or vice versa.

⁸We adopt temperature scaling (Guo et al., 2017) to calibrate the confidence estimates of the model.

4.4. Analysis using Occlusion Test

We present the results of prediction performance and the explainability analysis using occlusion in Tables 4 and 5, respectively. Analyzing Table 5, we observe that the model achieves higher accuracy in classifying instances with Supports Judgment compared to those with Neutral or Opposes Judgment. This could be attributed to the fact that the Opposes Judgment category is underrepresented in the occlusion dataset (due to fewer annotated tokens with this label) and the challenging task to classify Neutral instances. Among the three languages, French exhibits the highest score for the Supports Judgment category, but it also shows lower scores for the other classes.

Despite the MultiLingual model displaying a decrease in predictive performance, it shows some improvement in occlusion performance, particularly for the Supports Judgments class, across all languages. A similar trend is observed in the Joint training model, which consistently demonstrates a significant increase in explainability scores across languages for most classes.

While the inclusion of the DA component in both MonoLingual and MultiLingual models resulted in improved explainability scores for most labels compared to their counterparts, its addition to the Joint training model leads to mixed results. Surprisingly, the addition of the DA component to the Joint training model consistently increases prediction performance but does not maintain consistency in explainability performance. This finding emphasizes the importance of evaluating explainability to develop transparent systems that can make accurate predictions for the right reasons.

Overall, the lower scores across the board indicate the flawed inference about factors predictive for the outcome. Despite the impressive performances of state-of-the-art models on standard LJP prediction performance, there is still much progress to be made to make those models align as closely as possible with the rationales deemed relevant by legal experts. To create practical value for the legal field, the field of LJP should aim for a productive fusion of expert knowledge and data-driven insights, rather than data-driven correlation based learning.

4.5. Analysis using LCI Test

From Table 6, we can observe that the modification of the lower court has a considerable influence on the overall prediction confidence, as indicated by the changes in confidence scores up to 5% in both directions across all languages, despite the lower court name on average spanning around seven tokens in documents of an average length of 350 tokens. However, these small changes in confidence scores did result in label flips.

Model	German			French			Italian		
	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports
MonoLingual	3.02	16.78	15.10	1.95	3.68	40.24	0.49	3.68	11.24
MultiLingual	2.04	11.90	17.46	1.77	3.62	42.77	0.85	5.72	13.48
MonoLingual + DA	3.21	16.26	18.08	1.78	5.98	43.12	0.98	4.39	14.99
MultiLingual + DA	3.64	19.06	20.83	1.43	4.63	45.77	0.83	4.84	15.36
Joint Training	2.62	15.72	26.97	1.67	4.19	48.51	0.54	5.37	18.82
Joint Training + DA	3.75	14.54	21.95	1.93	5.73	45.73	0.63	4.82	19.68

Table 5: Analysis of Explainability using Occlusion - F1-scores across all instances from all four test sets for each label in every language. The higher the scores, the better the explainability.

Model	German				French				Italian			
	+ MES	- MES	Flip 1 → 0	Flip 0 → 1	+ MES	- MES	Flip 1 → 0	Flip 0 → 1	+ MES	- MES	Flip 1 → 0	Flip 0 → 1
MonoLingual	3.48 _{5.12}	-2.3 _{2.82}	2.28	0.43	2.56 _{2.65}	-2.2 _{3.79}	3.84	1.02	1.64 _{3.46}	-2.22 _{7.76}	0.12	1.20
MultiLingual	3.39 _{5.43}	-2.77 _{3.64}	1.71	0.32	4.01 _{4.62}	-3.06 _{3.22}	0.51	0.12	1.72 _{2.18}	-1.78 _{2.52}	1.05	2.88
MonoLingual + DA	3.09 _{5.15}	-2.77 _{5.42}	2.56	0.22	4.12 _{6.73}	-1.83 _{2.34}	1.24	0.82	1.25 _{2.23}	-1.29 _{1.97}	0.32	2.19
MultiLingual + DA	5.32 _{8.27}	-3.35 _{5.24}	4.56	2.56	4.08 _{6.05}	-6.48 _{9.64}	1.53	3.07	2.88 _{3.51}	-6.12 _{8.08}	2.56	3.85
Joint Training	3.32 _{6.18}	-1.86 _{2.89}	3.13	1.99	4.08 _{4.37}	-2.71 _{3.69}	0.51	2.56	6.13 _{6.79}	-2.66 _{2.12}	4.92	4.24
Joint Training + DA	3.23 _{4.46}	-1.84 _{2.45}	2.85	1.99	3.04 _{3.96}	-4.04 _{4.31}	3.07	2.32	6.14 _{7.94}	-3.21 _{3.11}	4.09	4.83

Table 6: Analysis of Lower Court Bias using LCI - Results of Positive and Negative MES Scores, and Label Flips across the three languages. Labels 0 and 1 indicates dismissal and approval respectively. Lower scores indicate a less biased model. Subscript indicates the standard deviation values.

Overall, no consensus exists on which model setting has yielded lower MES scores in both directions consistently across all the languages. MultiLingual model’s prediction performance decreased compared to the MonoLingual model across all three languages and its bias scores increased significantly across all languages, barring -MES for Italian and +MES for German. While the inclusion of the DA (Data Augmentation) component resulted in improved prediction performance compared to the non-DA variants in both MonoLingual and MultiLingual settings, the Multilingual + DA model exhibited a notable increase in bias. This suggests that the model’s reliance on lower court names is more pronounced in the presence of the DA component compared to its non-DA variant.

Joint training models, which aim to generalize across languages, demonstrate improved prediction performance, particularly for Italian, which is underrepresented in the training set. However, this improvement comes at the cost of higher MES scores, indicating potential overfitting to court-specific correlations rather than capturing the actual reasoning behind the predictions. Interestingly, despite training a single model using data from all three languages, the Italian data shows a significant divergence in MES scores compared to German and French. This highlights that representational bias across languages seems to be a crucial part. While adding DA to Joint training models significantly improves prediction performance,

bias scores lack a clear pattern.

Across all the models in the German setting, we can witness an overall pro-dismissal trend (greater + MES scores compared to - MES) echoing with more number of label flips from approval to dismissal. While French notices an overall pro-dismissal trend, Italian shows pro-approval trend, barring the cross-lingual models. These observed biases regarding lower court underscore the need for continuous bias evaluation and mitigation in LJP models.

5. Conclusion

In this work, we present the rationale dataset curated at fine-grained level of both ‘supporting’ and ‘opposing’ factors for Swiss Judgment Prediction (SJP), the only available multilingual LJP dataset. We employ a perturbation-based occlusion approach to assess various state-of-the-art models developed for SJP and also release four distinct occlusion test sets, occluding a different number of sentences in each of the sets. Our lower explainability scores suggest that the current models do not align well with the legal experts which can lead to sub-optimal litigation strategies due to flawed inference about factors responsible for the outcome. Furthermore, we assess the bias of the lower court information in the final predictions generated by the models using LCI test and notice that models learn spurious correlations about court-outcome in the

data. In the future, we would explore deconfounding strategy (Santosh et al., 2022) to improve the alignment between what models and experts deem relevant. One can explore different group robust algorithms such as adversarial removal, IRM, Group DRO and V-REx, as an effective bias-mitigation strategy (Chalkidis et al., 2022c) and investigate its impact on explainability. We hope our data resource will be useful to the research community working on Legal Judgement Prediction.

6. Limitations

In this study, our approach to obtaining rationales involved a consolidation process wherein we aimed to achieve a final set of high-quality annotations through discussions with legal experts. However, it is important to acknowledge that the assumption of a single ground truth may overlook the presence of genuine human variation, which can arise due to factors such as disagreement, subjectivity in annotation, or the existence of multiple plausible answers (Xu et al., 2023). Particularly in the field of law, where complexity and interpretation are inherent, it is well-recognized that lawyers may have differing legal assessments of case facts and how they contribute to the eventual outcome. Instead of attempting to resolve variations in expert labels, it is essential to acknowledge and embrace the inherent variation in human annotations. Moving forward, it is crucial to develop methods that can comprehensively capture and account for variation from data to evaluation, enabling a more comprehensive treatment of this variability in future research.

In the evaluation of our occlusion-based explainability setup, we utilized the F1-score, which focuses solely on the final label obtained from the change in confidence score between the baseline and occluded instances. However, it is important to emphasize the need for a metric that takes into account the magnitude of the difference in confidence scores during aggregation, in order to present a more comprehensive and holistic assessment.

7. Ethics Statement

The dataset used in this work comes from prior work by Niklaus et al. 2021 and these are publicly available on the <https://entscheidsuche.ch> platform and the names of the parties have been redacted by the court to ensure anonymity.

This work does not endorse or advocate for practical use of such systems. Instead our aim in this work is to rather empirically demonstrate that these systems are far from practical use due to their flawed inference about factors leading to outcome prediction. The scope of this work is to

study LJP from an explainability standpoint and to showcase the discrepancy between the prediction performance and explainability performance and emphasize the need to build technology that can help practitioners with reliable insights. Our dataset and findings associated with this work will contribute to advancing the field of explainable legal judgement prediction and provide valuable insights for developing more reliable and unbiased models in the future.

Furthermore, we would like to draw attention to the work by Tsarapatsanis and Aletras 2021 which discusses various normative factors related to ethics in the context of legal natural language processing. These discussions are crucial for fostering ethical thinking within the legal NLP community and ensuring the responsible development of systems that can assist lawyers, judges, and the general public.

Acknowledgements

Thanks to Ilias Chalkidis for providing feedback on initial ideas and experiments and to Lynn Grau, Thomas Lüthi & Angela Stefanelli for providing expert legal annotations.

8. Bibliographical References

- Abhaya Agarwal and Alon Lavie. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. *Proceedings of WMT-08*.
- Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoțiu-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.
- Vincent Aleven and Kevin D Ashley. 1997. Teaching case-based argumentation through a model and examples: Empirical evaluation of an intelligent learning environment. In *Artificial intelligence in education*, volume 39, pages 87–94. Citeseer.
- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias: There’s software used across the country to predict future criminals. *And it’s biased against blacks*. *ProPublica*, 23:77–91.
- Vithor Gomes Ferreira Bertalan and Evandro Eduardo Seron Ruiz. 2020. Predicting judicial outcomes in the brazilian legal system using textual

- features. In *DHandNLP@ PROPOR*, pages 22–32.
- Tobias Brugger, Matthias Stürmer, and Joel Niklaus. 2023. [MultiLegalSBD: A Multilingual Legal Sentence Boundary Detection Dataset](#). ArXiv:2305.01211 [cs].
- Stefanie Brüninghaus and Kevin D Ashley. 2003. Combining case-based and model-based reasoning for predicting the outcome of legal cases. In *Case-Based Reasoning Research and Development: 5th International Conference on Case-Based Reasoning, ICCBR 2003 Trondheim, Norway, June 23–26, 2003 Proceedings 5*, pages 65–79. Springer.
- Stefanie Brüninghaus and Kevin D Ashley. 2005. Generating legal arguments and predictions from case texts. In *Proceedings of ICAIL 2005*, pages 65–74.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural legal judgment prediction in english. In *Proceedings of ACL 2019*, pages 4317–4323.
- Ilias Chalkidis, Xiang Dai, Manos Fergadiotis, Prodromos Malakasiotis, and Desmond Elliott. 2022a. An exploration of hierarchical attention transformers for efficient long document classification. *arXiv preprint arXiv:2210.05529*.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. Paragraph-level rationale extraction through regularization: A case study on european court of human rights cases. In *Proceedings of the NAACL-HLT 2021*, pages 226–241.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022b. Lexglue: A benchmark dataset for legal language understanding in english. In *Proceedings of ACL 2022*, pages 4310–4330.
- Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. 2022c. Fairlex: A multilingual benchmark for evaluating fairness in legal text processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4389–4406.
- Branden Chan, Timo Möller, Malte Pietsch, Tanay Soni, and Chin Man Yeung. 2019. Deepset-open sourcing german bert.
- Charles Condevaux and Sébastien Harispe. 2022. Lsg attention: Extrapolation of pretrained transformers to long sequences. *arXiv preprint arXiv:2210.15497*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. 2020. A survey of the state of explainable ai for natural language processing. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 447–459.
- Matthias Grabmair. 2017. Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs. In *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*, pages 89–98.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- Wenyue Hua, Yuchen Zhang, Zhe Chen, Josie Li, and Melanie Weber. 2022. Legalrelectra: Mixed-domain language modeling for long-range legal text comprehension. *arXiv preprint arXiv:2212.08204*.
- Daniel Martin Katz, Michael J Bommarito, and Josh Blackman. 2017. A general approach for predicting the behavior of the supreme court of the united states. *PloS one*, 12(4):e0174698.
- Aaron Russell Kaufman, Peter Kraft, and Maya Sen. 2019. Improving supreme court forecasting using boosted decision trees. *Political Analysis*, 27(3):381–387.
- Arshdeep Kaur and Bojan Bozic. 2019. Convolutional neural network-based automatic prediction of judgments of the european court of human rights. In *AICS*, pages 458–469.
- Kankawin Kowsrihawat, Peerapon Vateekul, and Prachya Boonkwan. 2018. Predicting judicial decisions of criminal cases from thai supreme

- court using bi-directional gru with attention mechanism. In *2018 5th Asian Conference on Defense Technology (ACDT)*, pages 50–55. IEEE.
- André Lage-Freitas, Héctor Allende-Cid, Orivaldo Santana, and Livia Oliveira-Lage. 2022. Predicting brazilian court decisions. *PeerJ Computer Science*, 8:e904.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. Understanding neural networks through representation erasure. *arXiv preprint arXiv:1612.08220*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Zhenyu Liu and Huanhuan Chen. 2017. A predictive performance comparison of machine learning models for judicial cases. In *2017 IEEE Symposium series on computational intelligence (SSCI)*, pages 1–6. IEEE.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Bingfeng Luo, Yansong Feng, Jianbo Xu, Xiang Zhang, and Dongyan Zhao. 2017. Learning to predict charges for criminal cases with legal basis. In *Proceedings of EMNLP 2017*, pages 2727–2736.
- 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 ACL-IJCNLP 2021*, pages 4046–4062.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villemonte de La Clergerie, Djamé Seddah, and Benoît Sagot. 2020. Camembert: a tasty french language model. In *ACL 2020-58th Annual Meeting of the Association for Computational Linguistics*.
- Masha Medvedeva, Ahmet Üstün, Xiao Xu, Michel Vols, and Martijn Wieling. 2021. Automatic judgement forecasting for pending applications of the european court of human rights. In *ASAIL/LegalAIIA@ ICAIL*.
- 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.
- Emre Mumcuoğlu, Ceyhan E Öztürk, Haldun M Ozaktas, and Aykut Koç. 2021. Natural language processing in law: Prediction of outcomes in the higher courts of turkey. *Information Processing & Management*, 58(5):102684.
- 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.
- Joel Niklaus and Daniele Giofré. 2022. Budget-longformer: Can we cheaply pretrain a sota legal language model from scratch? *arXiv preprint arXiv:2211.17135*.
- 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 preprint arXiv:2301.13126*.
- Joel Niklaus, Veton Matoshi, Matthias Stürmer, Ilias Chalkidis, and Daniel E. Ho. 2023b. [Multilegalpile: A 689gb multilingual legal corpus](#).
- Joel Niklaus, Matthias Stürmer, and Ilias Chalkidis. 2022. An empirical study on cross-x transfer for legal judgment prediction. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*, pages 32–46.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Loreto Parisi, Simone Francia, and Paolo Magrani. 2020. Umberto: an italian language model trained with whole word masking. *Original-date*, 55:31Z.
- Vishvakshen 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].
- Jonathon Read, Rebecca Dridan, Stephan Oepen, and Lars Jørgen Solberg. 2012. Sentence boundary detection: A long solved problem? In *Proceedings of COLING 2012: Posters*, pages 985–994.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international*

- conference on knowledge discovery and data mining, pages 1135–1144.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Edwina L Rissland and Kevin D Ashley. 1987. A case-based system for trade secrets law. In *Proceedings of the 1st international conference on Artificial intelligence and law*, pages 60–66.
- T. Y. S. S Santosh, Marcel Perez San Blas, Phillip Kemper, and Matthias Grabmair. 2023a. Leveraging task dependency and contrastive learning for case outcome classification on european court of human rights cases. *arXiv preprint arXiv:2302.00768*.
- T. Y. S. S Santosh, Oana Ichim, and Matthias Grabmair. 2023b. Zero shot transfer of article-aware legal outcome classification for european court of human rights cases. *arXiv preprint arXiv:2302.00609*.
- T.y.s.s Santosh, Shanshan Xu, Oana Ichim, and Matthias Grabmair. 2022. Deconfounding legal judgment prediction for European court of human rights cases towards better alignment with experts. In *Proceedings of EMNLP 2022*.
- Jaromir Savelka, Vern R Walker, Matthias Grabmair, and Kevin D Ashley. 2017. Sentence boundary detection in adjudicatory decisions in the united states. *Traitement automatique des langues*, 58:21.
- Gil Semo, Dor Bernsohn, Ben Hagag, Gila Hayat, and Joel Niklaus. 2022. Classactionprediction: A challenging benchmark for legal judgment prediction of class action cases in the us. *arXiv preprint arXiv:2211.00582*.
- Mehmet Fatih Sert, Engin Yıldırım, and İrfan Haşlak. 2021. Using artificial intelligence to predict decisions of the turkish constitutional court. *Social Science Computer Review*, page 08944393211010398.
- Rafe Athar Shaikh, Tirath Prasad Sahu, and Veena Anand. 2020. Predicting outcomes of legal cases based on legal factors using classifiers. *Procedia Computer Science*, 167:2393–2402.
- Saeed Sharifi-Malvajerdi, Michael Kearns, and Aaron Roth. 2019. Average individual fairness: Algorithms, generalization and experiments. *Advances in neural information processing systems*, 32.
- Benjamin Strickson and Beatriz De La Iglesia. 2020. Legal judgement prediction for uk courts. In *Proceedings of the 2020 the 3rd international conference on information science and system*, pages 204–209.
- Octavia-Maria Şulea, Marcos Zampieri, Shervin Malmasi, Mihaela Vela, Liviu P Dinu, and Josef van Genabith. 2017a. Exploring the use of text classification in the legal domain.
- Octavia-Maria Şulea, Marcos Zampieri, Mihaela Vela, and Josef van Genabith. 2017b. 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.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *ICML*, pages 3319–3328. PMLR.
- 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.
- Michael Benedict L Virtucio, Jeffrey A Aborot, John Kevin C Abonita, Roxanne S Avinante, Rother Jay B Copino, Michelle P Neverida, Vanesa O Osiana, Elmer C Peramo, Joanna G Syjuco, and Glenn Brian A Tan. 2018. Predicting decisions of the philippine supreme court using natural language processing and machine learning. In *2018 IEEE 42nd annual computer software and applications conference (COMPSAC)*, volume 2, pages 130–135. IEEE.
- Bernhard Walth, Georg Bonczek, Elena Scepankova, Jörg Landthaler, and Florian Matthes. 2017. Predicting the outcome of appeal decisions in germany’s tax law. In *International conference on electronic participation*, pages 89–99. Springer.
- Yuzhong Wang, Chaojun Xiao, Shirong Ma, Haoxi Zhong, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2021. Equality before the law: Legal judgment consistency analysis for fairness. *arXiv preprint arXiv:2103.13868*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 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.

Yongkai Wu, Lu Zhang, and Xintao Wu. 2019. Counterfactual fairness: Unidentification, bound and algorithm. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.

Shanshan Xu, Santosh T.y.s.s, Oana Ichim, Isabella Risini, Barbara Plank, and Matthias Grabmair. 2023. From dissonance to insights: Dissecting disagreements in rationale construction for case outcome classification. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9576, Singapore. Association for Computational Linguistics.

Linan Yue, Qi Liu, Binbin Jin, Han Wu, Kai Zhang, Yanqing An, Mingyue Cheng, Biao Yin, and Dayong Wu. 2021. Neurjudge: a circumstance-aware neural framework for legal judgment prediction. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 973–982.

Mikhail Yurochkin, Amanda Bower, and Yuekai Sun. Training individually fair ml models with sensitive subspace robustness. In *International Conference on Learning Representations*.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*, pages 1171–1180.

Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer.

Junzhe Zhang and Elias Bareinboim. 2018. Fairness in decision-making—the causal explanation formula. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Haoxi Zhong, Yuzhong Wang, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. Iteratively questioning and answering for interpretable legal judgment prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1250–1257.

A. Appendix

A.1. Explainability performance for different levels of occlusion

We report the label wise F1-score for each occlusion test for every language in Tables 7, 8, 9. Overall, "Opposes judgement" and "Neutral" are challenging ones compared to "Supports judgement". In the case of French, there was an improvement in scores as the number of occluded sentences increased. This improvement indicates that the model was able to correctly associate the label "supports judgements" with the occluded sentences, thereby enhancing the model's explainability performance. However, in the case of French and Italian, a different trend was observed. The model did not exhibit the same improvement as the number of occluded sentences increased. It is speculated that the model might have encountered conflicting labels for each occluded sentence, leading to incorrect predictions when multiple occlusions were present.

Model	Set 1			Set 2			Set 3			Set 4		
	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports
MonoLingual	23.08	20.98	33.86	8.80	17.85	27.30	3.28	20.19	16.84	1.15	21.34	11.89
MultiLingual	23.79	18.05	34.46	8.12	13.87	26.81	1.52	12.79	24.25	1.08	14.10	11.17
MonoLingual + DA	24.05	21.77	25.12	10.69	19.29	17.34	3.93	18.24	10.01	0.76	10.80	22.03
MultiLingual + DA	28.80	26.40	36.44	12.45	21.64	27.98	4.49	19.83	17.15	1.04	18.01	11.08
Joint Training	24.52	20.14	31.87	9.86	10.62	28.25	2.69	9.20	27.52	0.98	5.75	26.14
Joint Training + DA	23.85	23.33	38.3	13.40	16.02	25.11	4.71	10.74	12.67	1.24	7.77	8.74

Table 7: Explainability performance for German dataset over different occlusion test sets

Model	Set 1			Set 2			Set 3			Set 4		
	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports
MonoLingual	21.48	10.17	36.67	4.34	4.87	38.17	1.34	4.17	53.64	0.18	3.58	67.99
MultiLingual	26.80	2.40	40.98	6.19	0.68	38.14	1.05	0.79	40.19	0.00	0.22	51.58
MonoLingual + DA	20.05	10.53	28.57	4.72	8.67	43.29	1.28	6.79	58.81	0.14	4.35	69.36
MultiLingual + DA	21.11	4.73	30.41	6.17	2.02	38.24	1.03	1.26	41.39	0.00	1.87	40.86
Joint Training	21.43	3.49	37.37	7.09	0.68	43.02	1.11	1.12	49.77	0.12	1.12	58.29
Joint Training + DA	28.11	14.05	38.78	8.51	9.69	44.26	1.21	10.09	43.65	0.00	8.80	42.04

Table 8: Explainability performance for French dataset over different occlusion test sets

Model	Set 1			Set 2			Set 3			Set 4		
	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports	Opposes	Neutral	Supports
MonoLingual	9.01	24.70	23.53	2.01	26.56	6.72	0.52	31.08	0.22	0.08	32.43	0.00
MultiLingual	12.84	19.18	31.25	2.54	7.66	14.86	0.41	7.22	4.84	0.0	3.92	1.08
MonoLingual + DA	22.11	3.03	27.10	5.61	0.24	13.37	0.84	0.33	6.12	0.05	0.35	1.92
MultiLingual + DA	21.28	11.43	35.82	8.66	2.62	20.77	0.76	1.06	8.51	0.06	0.18	2.46
Joint Training	17.24	11.21	25.61	3.27	5.61	17.55	0.37	1.79	8.07	0.0	0.63	3.01
Joint Training + DA	20.99	17.94	33.53	3.99	6.50	23.02	0.46	2.28	11.3	0.04	0.77	3.91

Table 9: Explainability performance for Italian dataset over different occlusion test sets