Lex Rosetta: Transfer of Predictive Models Across Languages, Jurisdictions, and Legal Domains

Jaromir Savelka
jsavelka@cs.cmu.edu
Carnegie Mellon University
USA

Hannes Westermann
Karim Benyekhlef
Université de Montréal
Canada

Charlotte S. Alexander
Jayla C. Grant
Georgia State University
USA

David Restrepo Amariles
Rajaa El Hamdani
HEC Paris
France

Sébastien Meeùs
Aurore Troussel
HEC Paris
France

Michał Araszkiewicz
Uniwersytet Jagielloński
Poland

Kevin D. Ashley
Alexandra Ashley
University of Pittsburgh
USA

Karl Branting
MITRE Corporation
USA

Mattia Falduti
Libera Università di Bolzano
Italy

Matthias Grabmair
Technische Universität München
Germany

Jakub Harašta
Tereza Novotná
Masarykova univerzita
Czech Republic

Elizabeth Tippet
Shiwanni Johnson
University of Oregon
USA

ABSTRACT

In this paper, we examine the use of multi-lingual sentence embeddings to transfer predictive models for functional segmentation of adjudicatory decisions across jurisdictions, legal systems (common and civil law), languages, and domains (i.e. contexts). Mechanisms for utilizing linguistic resources outside of their original context have significant potential benefits in AI & Law because differences between legal systems, languages, or traditions often block wider adoption of research outcomes. We analyze the use of Language-Agnostic Sentence Representations in sequence labeling models using Gated Recurrent Units (GRUs) that are transferable across languages. To investigate transfer between different contexts we developed an annotation scheme for functional segmentation of adjudicatory decisions. We found that models generalize beyond the contexts on which they were trained (e.g., a model trained on administrative decisions from the US can be applied to criminal law decisions from Italy). Further, we found that training the models on multiple contexts increases robustness and improves overall performance when evaluating on previously unseen contexts. Finally, we found that pooling the training data from all the contexts enhances the models' in-context performance.

CCS CONCEPTS

• **Applied computing** → Law; Annotation; • **Information systems** → Document structure; Structure and multilingual text search; Data mining.

KEYWORDS

multi-lingual sentence embeddings, transfer learning, domain adaptation, adjudicatory decisions, document segmentation, annotation

ACM Reference Format:

Jaromir Savelka, Hannes Westermann, Karim Benyekhlef, Charlotte S. Alexander, Jayla C. Grant, David Restrepo Amariles, Rajaa El Hamdani, Sébastien Meeùs, Aurore Troussel, Michał Araszkiewicz, Kevin D. Ashley, Alexandra Ashley, Karl Branting, Mattia Falduti, Matthias Grabmair, Jakub Harašta, Tereza Novotná, Elizabeth Tippet, and Shiwanni Johnson. 2021. Lex Rosetta: Transfer of Predictive Models Across Languages, Jurisdictions, and Legal Domains. In *Eighteenth International Conference for Artificial Intelligence and Law (ICAIL '21)*, June 21–25, 2021, São Paulo, Brazil. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3462757.3466149>

1 INTRODUCTION

This paper explores the ability of multi-lingual sentence embeddings to enable training of predictive models that generalize beyond individual languages, legal systems, jurisdictions, and domains (i.e., contexts). We propose a new type schema for functional segmentation of adjudicatory decisions (i.e., decisions of trial and appellate court judges, arbitrators, administrative judges and boards) and use it to annotate legal cases across eight different contexts (7 countries, 6 languages). We release the newly created dataset (807 documents

This work is licensed under a Creative Commons Attribution International 4.0 License.
ICAIL '21, June 21–25, 2021, São Paulo, Brazil
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8526-8/21/06.
<https://doi.org/10.1145/3462757.3466149>

with 89,661 annotated sentences) including the annotation schema to the public.¹

In the area of AI & Law, research typically focuses on a single context, such as decisions of a specific court on a specific issue within a specific time range. This is justified by the complexity of legal work and the need for nuanced solutions to particular problems. At the same time, this narrow focus can limit the applicability of the research outcomes, since a proposed solution might not be readily transferable to a different context. In text classification, for example, a model might simply memorize a particular vocabulary characteristic of a given context, rather than acquiring the semantics of a predicted type. Adaptation of such a model to a new context would then require the assembly of a completely new dataset. This may be both time-consuming and expensive, since the annotation of legal documents relies on legal expertise.

Certain tasks appear to be of interest to researchers from multiple countries with different legal traditions (e.g., deontic classification of legal norms embodied in statutory law, argument extraction from case law, summarization/simplification of legal documents, etc.). This suggests that there may be several core tasks in AI & Law that are of general interest in almost any context. One such task is a functional segmentation of adjudicatory decisions, which has been the subject of numerous studies in the past (see Section 2). In this paper, we show that for this particular task it is possible to leverage linguistic resources created in multiple contexts.

This has wide-reaching implications for AI & Law research. Since annotation of training data is expensive, models that are able to use existing data from other contexts might be instrumental in enabling real-world applications that can be applied across contexts. Such approaches may further enable international collaboration of researchers, each annotating their own part of a dataset to contribute to a common pool (as we do in this work) that could be used to train strong models able to generalize across contexts.

1.1 Functional Segmentation

We investigate the task of segmenting adjudicatory decisions based on the functional role played by their parts. While there are significant differences in how decisions are written in different contexts, we hypothesize that selected elements might be universal, such as, for example, sections:

- (1) describing facts that give rise to a dispute;
- (2) applying general legal rules to such facts;
- (3) stating an outcome of the case (i.e., how was it decided).

This conjecture is supported by the results of the comparative project titled *Interpreting Precedents* [23], which aimed to analyze (among other things) the structure in 11 different jurisdictions. The findings of this project suggest that the structure indicated above may be considered a general model followed in the investigated jurisdictions, although variations exist that are characteristic of particular legal systems and types of courts and their decisions.

The ability to segment cases automatically could be beneficial for many tasks. It could support reading and understanding of legal decisions by students, legal practitioners, researchers, and the public. It could facilitate empirical analyses of the discourse structure of decisions. It could enhance the performance, as well

as the user experience, of legal search tools. For example, if a user searches for an application of a legal rule, they might restrict the search to the section where a judge applies the rule to a factual situation. Judges themselves might find the technique useful, within their own jurisdictions but also in transnational disputes involving the application of different legal standards. The same benefits apply to non-court settings, e.g., international arbitration, where many jurisdictions' laws and their interpretation matter. Further, the segmentation of decisions into meaningful sections could serve as an important step in many legal document processing pipelines.

1.2 Hypotheses

To investigate how well predictive models, based on multi-lingual sentence embeddings, learn to segment cases into functional parts across different contexts, we evaluated the following hypotheses:

- (H1) A model trained on a single context can generalize when transferred to other, previously unseen, contexts.
- (H2) A model trained on data pooled from multiple contexts is more robust and generalizes better to unseen contexts than a model trained on a single context.
- (H3) A context-specific model benefits from pooling the in-domain data with data from other contexts.

1.3 Contributions

By carrying out this work, we provide the following contributions to the AI & Law research community:

- Detailed definition and analysis of a functional segmentation task that is widely applicable across different contexts.
- A new labeled dataset consisting of 807 documents (89,661 sentences) from seven countries in six different languages.
- Evidence of the effectiveness of multi-lingual embeddings on processing legal documents.
- Release of the code used for data preparation, analysis, and the experiments in this work.

2 RELATED WORK

Segmenting court decisions into smaller elements according to their function or role is an important task in legal text processing. Prior research utilizing supervised machine learning (ML) approaches or expert crafted rules can roughly be distinguished into two categories. First, the task could be to segment the text into a small number of contiguous parts typically comprising multiple paragraphs (this work). Different variations of this task were applied to several legal domains from countries, such as Canada [15], the Czech Republic [17], France [8], or the U.S. [27]. Second, the task could instead be labeling smaller textual units, often sentences, according to some predefined type system (e.g., rhetorical roles, such as evidence, reasoning, conclusion). Examples from several domains and countries include administrative decisions from the U.S. [33, 41], multi-domain court decisions from India [6], international arbitration decisions [9], or even multi-{domain,country} adjudicatory decisions in English [28]. Identifying a section that states an outcome of the case has also received considerable attention separately [25, 38]. To the best of our knowledge, existing work on functional segmentation of court decisions is limited to a single

¹https://github.com/lexrosetta/caselaw_functional_segmentation_multilingual

language—ours being the first paper exploring the task jointly on legal documents in multiple languages.

In NLP, the success of word embeddings was followed by an increasing interest in learning continuous vector representations of longer linguistic units, such as sentences (a trend that has been reflected in AI & Law research as well [34, 41]). Multi-lingual representations recently attracted ample attention. While most of the earlier work was limited to a few close languages or pairwise joint embeddings for English and one foreign language, several approaches to obtain general-purpose massively multi-lingual sentence representations were proposed [5, 11, 13]. Such representations were utilized in many downstream applications, such as document classification [21], machine translation [2], question answering [22], hate speech detection [4], or information retrieval (IR) in the legal domain [40]. Our work is one of the first such applications in the legal domain and to the best of our knowledge the first dealing with more than two languages.

Approaches other than language-agnostic sentence embeddings (this work) were used in AI & Law research focused on texts in multiple languages. A recent line of work mapped recitals to articles in EU directives and normative provisions in Member states' legislation [24]. There, mono-lingual models were used (i.e., one model per language). Other published applications in multi-lingual legal IR were based on thesauri [14, 29]. A common technique to bridge the language gap was the use of ontologies and knowledge graphs [1, 3, 7, 16]. The multi-lingual environments, such as EU or systems established by international treaties, attracted work on machine translation [20], meaning equivalence verification [32], and building of parallel corpora [30, 31].

3 DATASET

In creating the dataset, the first goal was to identify a task that would be useful across different contexts. After extensive literature review, we identified the task of functional segmentation of adjudicatory decisions as a viable candidate. To make the task generalizable, we decided to include only a small number of core types.

- (1) *Out of Scope* – Parts outside of the main document body (e.g., metadata, editorial content, dissents, end notes, appendices).
- (2) *Heading* – Typically an incomplete sentence or marker starting a section (e.g., “Discussion,” “Analysis,” “II.”).
- (3) *Background* – The part where the court describes procedural history, relevant facts, or the parties' claims.
- (4) *Analysis* – The section containing reasoning of the court, issues, and application of law to the facts of the case.
- (5) *Introductory Summary* – A brief summary of the case at the beginning of the decision.
- (6) *Outcome* – A few sentences stating how the case was decided (i.e., the overall outcome of the case).

We created detailed annotation guidelines defining the individual types as well as describing the annotation workflow (tooling, steps taken during annotation). Eight teams of researchers from six different countries (14 persons) were trained in the annotation process through online meetings. After this, each annotator conducted a dry-run annotation on 10 cases and received detailed feedback. Then, each team was tasked with assembling approximately 100

adjudicatory decisions. Each team developed specifications for the decisions to be included in their part of the dataset.

Four of the contexts were double-annotated by two annotators (Canada, Czech R., France, U.S.A. I); the remaining four by just one. Each team had at least one member with a completed law degree. When a team had more than one member, law students were allowed to be included.

A high-level description of the resulting dataset is provided in Table 1. It consists of eight contexts from seven different countries (two parts are from the U.S.) with 807 documents in six languages (three parts are in English). Most of the contexts include judicial decisions, while U.S.A. II was the only context that consisted solely of administrative decisions. There are considerable variations in the length of the documents. While an average document in the U.S.A. I context comprises of 530.6 sentences, an average document in the France context is about ten times shorter (59.0 sentences).

The four double-annotated parts enabled us to examine the inter-annotator agreement. Table 2 shows the raw agreement on a character level. While it appears that recognizing the *Outcome* was rather straightforward in the France and U.S.A. I contexts, it was more complicated in case of Canada and the Czech R. This might be due to a presence/absence of some structural clue. We also observe that in the Czech R. context it was presumably much easier to distinguish between the *Background* and *Analysis* than in case of the other three contexts.

In this paper, we focus on prediction of the *Background*, *Analysis*, and *Outcome* types. We decided to exclude the *Introductory Summary* type, since it is mainly present in the data from the United States. For the double-annotated datasets, we picked the annotations that appeared to be of higher quality (either by consensus between the annotators themselves or by a decision of a third unbiased expert).

We first removed all the spans of text annotated with either of the *Out of Scope* or *Heading* types. The removal of *Out of Scope* leaves the main body of a decision stripped of potential metadata or editorial content at the beginning of the document as well as dissents, or end notes at the end. The removal of the text spans annotated with the *Heading* type might appear counter-intuitive since headings often provide a clear clue as to the content of the following section (e.g., “Outcome”, “Analysis” etc.). We remove these (potentially valuable) headings because we want to focus on the more interesting task of recognizing the sections purely by the semantics of their constitutive sentences. This task is more challenging, and more closely emulates generalization to domains where headings are not used or not present in all cases, or are not reliable indicators.

The transformed documents are separated into several segments based on the annotations of the three remaining types. Each segment is then split into sentences.² A resulting document is a sequence of sentences labeled with one of the *Background*, *Analysis*, or *Outcome* types. The highlighted (green) part of Table 1 provides

²We used the processing pipeline from <https://spacy.io/> (large models). For the Czech language we used <https://github.com/TakeLab/spacy-udpipe> with the Czech model (PDT) from <https://universaldependencies.org/>. The output was further processed with several regular expressions. A different method was used for the French dataset, which consists of a few very long sections, internally separated by a semicolon. After consultation with an expert we decided to split the cases by the semicolon as well.

Table 1: Descriptive statistics of the created dataset. Each entry provides information about the country, the language of the decisions (Lang), and the number of documents (Docs) in a specific context. The Sentence-Level Statistics subsection reports basic descriptive statistics focused on sentences as well as the number of sentences labeled with each type (OoS - Out of Scope, Head - Heading, Int.S. - Introductory Summary, Back - Background, Anl - Analysis, Out - Outcome). The part highlighted in green contains the counts of sentences labeled with the types we focus on in this work.

Country	Lang	Docs	Sentence-Level Statistics										Description
			Count	Avg	Min	Max	OoS	Head	Int.S.	Back	Anl	Out	
Canada	EN	100	12168	121.7	8	888	873	438	20	3319	7190	328	Random selection of cases retrieved from www.canlii.org from multiple provinces. The selection is not limited to any specific topic or court.
Czech R.	CS	100	11283	112.8	10	701	945	1257	2	3379	5422	278	A random selection of cases from Constitutional Court (30), Supreme Court (40), and Supreme Administrative Court (30). Temporal distribution was taken into account.
France	FR	100	5507	55.1	8	583	3811	220	0	485	631	360	A selection of cases decided by Cour de cassation between 2011 and 2019. A stratified sampling based on the year of publication of the decision was used to select the cases.
Germany	DE	104	10724	103.1	12	806	406	333	38	2960	6697	290	A stratified sample from the federal jurisprudence database spanning all federal courts (civil, criminal, labor, finance, patent, social, constitutional, and administrative).
Italy	IT	100	4534	45.3	10	207	417	1098	0	986	1903	130	The top 100 cases of the criminal courts stored between 2015 and 2020 mentioning "stalking" and keyed to the Article 612 bis of the Criminal Code.
Poland	PL	101	9791	96.9	4	1232	796	303	0	2736	5820	136	A stratified sample from trial-level, appellate, administrative courts, the Supreme Court, and the Constitutional tribunal. The cases mention "democratic country ruled by law."
U.S.A. I	EN	102	24898	244.1	34	1121	574	1235	475	6042	16098	474	Federal district court decisions in employment law mentioning "motion for summary judgment," "employee," and "independent contractor."
U.S.A. II	EN	100	10756	107.6	24	397	1766	650	639	3075	4402	224	Administrative decisions from the U.S. Department of Labor. Top 100 ordered in reverse chronological rulings order, starting in October 2020, were selected.
Overall	6	807	89661	105.6	4	1232	9588	5534	1174	22982	48163	2220	

Table 2: Raw agreement on a character level for the four datasets with two human annotators. The agreement is computed as a percentage of characters where both the annotators agree on a specific type over all the characters annotated by that type by any of the annotators. (NM=Not Marked)

	OoS	Head	Int.S.	Back	Anl	Out	NM
Canada	97.2	68.2	44.0	83.3	92.2	79.9	43.4
Czech R.	80.3	54.6	0.0	92.6	94.5	46.9	10.0
France	93.5	92.5	N/A	43.0	72.2	99.1	1.0
U.S.A. I	90.8	71.0	74.2	78.4	93.7	91.1	18.4
Overall	91.8	70.4	72.1	82.1	92.6	77.3	3.8

basic descriptive statistics of the resulting dataset per the individual contexts. Our final dataset for analysis consists of 807 cases split into 74,539 annotated sentences.

4 MODELS

In our experiments we use the Language-Agnostic Sentence Representations (LASER) model [5] to encode sentences from different languages into a shared semantic space. Each document becomes a series of vectors which represent the semantic content of a single sentence. We use these vectors to train a bidirectional Gated Recurrent Unit (GRU) model [10] for predicting sentence labels.

The LASER model is a language-agnostic bidirectional LSTM encoder coupled with an auxiliary decoder and trained on parallel corpora. The sentence embeddings are obtained by applying a max-pooling operation over the output of the encoder. The resulting sentence representations (after concatenating both directions) are 1024-dimensional. The released trained model,³ which we use in this work, supports 93 languages (including the six in our dataset) belonging to 30 different families and written in 28 different scripts. The model was trained on 223 million parallel sentences. The joint

encoder itself has no information on the language or writing script of the tokenized text, while the tokenizer is language specific. It is even possible to mix multiple languages in one sentence. The focus of the LASER model is to produce vector representations of sentences that are general with respect to two dimensions: the input language and the NLP task [5]. An interesting property of such universal multi-lingual sentence embeddings is the increased focus on the sentence semantics, as the syntax or other surface properties are unlikely to be shared among languages.

The GRU neural network [10] is an architecture based on a recurrent neural network (RNN) that is able to learn the mapping from a sequence of an arbitrary length to another sequence. GRUs are able to either score a pair of sequences or to generate a target sequence given a source sequence (this work). In a bidirectional GRU, two separate sequences are considered (one from right to left and the other from left to right). Traditional RNNs work well for shorter sequences but cannot be successfully applied to long sequences due to the well-known problem of vanishing gradients. Long Short-Term Memory (LSTM) networks [18] have been used as an effective solution to this problem (the forget gate, along with the additive property of the cell state gradients). GRUs have been proposed as an alternative to LSTMs with a reduced number of parameters. In GRUs there is no explicit memory unit, and the forget gate and the update gate are combined. The performance of GRUs was shown to be superior to that of LSTMs in the scenario of long texts and small datasets [39], which is the situation in this work. For these reasons, we chose to use GRUs over LSTMs.

The overall structure of the employed model is shown in Figure 1.⁴ Each case is transformed into a 1080×1024 matrix. The number 1080 represents the maximum length (in sentences) of any case in our dataset. Shorter cases are padded to be of uniform length. The vectors are passed to the model in batches of size 32. They first go through a masking layer, which masks the sentences used for

³<https://github.com/facebookresearch/LASER>

⁴The model was implemented using the Keras framework (<https://keras.io/>).

Figure 1: The structure of the sequential model used for prediction. Each case m is split into n sentences, which are converted to language-independent LASER vector embeddings. These are fed through a bidirectional recurrent GRU-model, which predicts one of the three labels per input sentence.

padding. The data is then passed to a bidirectional GRU model with 256 units and a dropout of 0.2. Finally, the model contains a time-distributed dense output layer with softmax activation, which outputs the predicted label (i.e., *Background*, *Analysis*, *Outcome*) for each sentence. As a loss function, we use categorical cross-entropy. As optimizer we use the Adam algorithm [19] with initial learning rate set to $4e^{-3}$ (reduced to $4e^{-4}$ once validation loss has stopped decreasing, with a patience of 50 epochs). We train the model for up to 1000 epochs. We halt the training once the validation accuracy has not increased in 80 epochs. For prediction, we use the best model as determined by validation accuracy.

5 EXPERIMENTAL DESIGN

We performed three experiments to test the hypotheses (Section 1.2). The first experiment (H1) focused on model generalization across different contexts (Out-Context). The second experiment (H2) assessed model robustness when trained on multiple pooled contexts different from the one where the model was applied (Pooled Out-Context). Finally, the third experiment (H3) analyzed the effects of pooling the target context data with data from other contexts (Pooled with In-Context). The different pools of training data, and the baselines we compare them against, are summarized in Table 3 and described in Sections 5.1-5.3.

Since our dataset is limited, we performed a 10-fold cross-validation. The folds were kept consistent across all the experiments. The experiments were conducted in the following manner:

- (1) Index $i = 1$ is set.
- (2) Pool of training context(s) is selected (see Sections 5.1-5.3).
- (3) A single test context is selected (see Sections 5.1-5.3).
- (4) Eight of the folds from the training context(s) are used as training data (index different from i and $(i + 1) \bmod 10$).
- (5) The folds with index $(i + 1) \bmod 10$ from the training context(s) are designated as validation data.

Figure 2: A visualization of the three experimental setups. The figure shows how training contexts are selected for H1, H2 and H3, and examples of how the folds are assigned to Train, Val and Test.

Table 3: Description of different training data selection and baselines used for H1, H2 and H3. The random (H1), the best single Out-Context models (H2), and the In-Context models (H3) baselines are highlighted.

Name	Trained on	Hyp.	Baseline
Random	Target Context	-	-
In-Context	Target Context	-	-
Out-Context	Non-Target Context	H1	Random
Pooled Out-Context	Pooled Non-Target Contexts	H2	Best performing Out-Context model per context
Pooled with In-Context	Non-Target and Target Pooled Contexts	H3	In-Context

- (6) The i fold from the test context is designated as test data.
- (7) The models are trained and evaluated.
- (8) Index $i = i + 1$ is set.
- (9) If $i \leq 10$ go to (4) else finish.

Note that from here on we highlight the baselines using the colors as shown in Table 3 to improve clarity.

5.1 Out-Context Experiment (H1)

In this experiment, we investigated the ability of the models to generalize beyond the training context. The pool of training data consisted of a single context at a time (see Figure 2). As a baseline, we used a random classifier that learns the distribution of the target dataset (train and validation folds) and generates predictions based on that distribution. The baseline thus had a slight advantage in having access to the signal from the target dataset. The Out-Context models performing better than the random baseline would show

that the models are able to transfer (at least some) knowledge from one context to another.

5.2 Pooled Out-Context Experiment (H2)

The second experiment focused on the ability of the models to gain robustness as they learn from more than one context, when applied to an unseen context. Therefore, the training pool for the experiments consisted of data from all the contexts *except* the current target context (see Figure 2). As a baseline for each context, we used the performance measurement taken from the **best Out-Context model for that context** (see Section 5.1). This was a very competitive baseline since the best Out-Context model likely stems from the context that is the most similar to the target context. If the Pooled Out-Context model performed better than the best Out-Context model, this would indicate that pooling datasets increases the robustness of the model, allowing it to perform better on previously unseen contexts.

5.3 Pooled With In-Context Experiment (H3)

The third experiment focused on pooling the target context’s data with data from the other contexts. The training pool therefore incorporates all of the contexts, including the target one (see Figure 2). As a baseline, we used the **In-Context models**, trained on the target context. Again, this is a very strong baseline since the In-Context model should be able to learn the most accurate signal from the target context (including context specific peculiarities). If the model trained on the pooled contexts (including target) is able to outperform the In-Context model, this would indicate that pooling contexts is beneficial in terms of both robustness and absolute performance on the individual contexts.

6 EVALUATION METHOD

To evaluate the performance of the systems trained on data from different contexts, we used Precision (P), Recall (R), and F_1 -measure. We compute the evaluation metrics for each class (i.e., *Background*, *Analysis*, *Outcome*) per fold. We then report per class F_1 -scores as well as an overall F_1 -score (micro) including its standard deviation.

For statistical evaluation we used the Wilcoxon signed-rank test [37] to perform pair comparisons between methods and the baselines as suggested by [12]. We used the overall (micro) F_1 -score as the evaluation metric. The null-hypothesis states that the mean performance of the two compared methods (i.e., an assessed model and a baseline) is equivalent.

Since the number of samples is rather small (7 contexts in testing H1, and 8 contexts in testing H2 and H3), we are not able to reject the null hypotheses at $\alpha = 0.05$ for any of H1–H3. This is as expected given the size of the dataset, and it is not unacceptable given the exploratory nature of our work at this stage. We will be in a position to formally evaluate the hypotheses once we extend the dataset (see Section 10). Instead, we report raw p-values produced by the testing as an evidence of the assessed methods’ effectiveness. The p-value is the probability of seeing the results at least as extreme as those we observed given the null hypothesis is correct (i.e., their mean performance is the same).

Figure 3: A description of performance metrics reported in Table 4.

7 RESULTS

The results of all three experiments are shown in Table 4. Each row reports a performance of a specific model across different contexts (columns). Each cell shows a bold number on the first row which correspond to the (micro) average F_1 -score over the three predicted classes, with standard deviation across the 10 folds. The F_1 -scores of the three classes are reported in the second row of each cell ordered as *Background*, *Analysis*, *Outcome*. A visual explanation of cell contents can be found in Figure 3.

7.1 Out-Context Experiment (H1)

The performance of the models trained during the Out-Context experiment is reported by the eight rows of Table 4 starting with Canada and ending with U.S.A. II. Here, the models are trained on a single context (row), and then evaluated on all of the contexts (columns). It appears the models perform reasonably well under the Out-Context condition. The application of the trained models outperforms the **random baseline** across the board (in 54 out of 56 instances). This is further corroborated by the low p-values obtained for these methods when compared to the random baseline (reported in Table 4).

Several interesting patterns emerge from the results. First, it appears that models trained on contexts with the same language or a language from the same family perform better. For example, the models trained on Canada and the two U.S.A. contexts perform well among each other. A similar observation applies to the models trained on the Poland and the Czech R. contexts. Quite surprisingly the models succeeded in identifying the *Outcome* sentences to some extent, despite a heavy under-representation of these sentences in our dataset. For example, a model trained on Canada predicts the *Outcome* sentences of the both U.S.A. contexts with an average F_1 -score close to 0.7. At the same time the *Outcome* sentences are by far the most challenging ones. On multiple occasions the models completely fail to predict these (e.g., Canada→Germany, or Poland→France).

Overall, the results demonstrate the ability of the models to effectively transfer the knowledge learned on one context to another. At the same time, it is clear that the models trained on one context and applied to a different one do not perform as well as the **In-Context models**.

Table 4: Results of the Out-Context, Pooled Out-Context, and Pooled with In-Context experiments. Each row reports a performance of a model across contexts. A bold number in each cell reports the (micro) average F₁-score over the predicted classes and the standard deviation across the 10 folds. The F₁-scores of the three classes are reported below ordered as Background, Analysis, Outcome. A visual explanation of cell contents can be found in Figure 3. The random (H1), the best single Out-Context models (H2), and the In-Context (H3) baselines are highlighted. The p-values are color coded to match the baselines.

	Canada	Czech R.	France	Germany	Italy	Poland	U.S.A. I	U.S.A. II	Avg (-test)	Avg (+test)
Random	.54 ± .04	.49 ± .02	.33 ± .05	.56 ± .04	.50 ± .04	.55 ± .04	.58 ± .03	.51 ± .02	.51 ± .07	
(dist.)	.31 .66 .06	.36 .59 .04	.32 .37 .25	.29 .68 .03	.31 .62 .04	.32 .67 .00	.26 .72 .02	.37 .61 .02	.32 .62 .10	
Canada	.82 ± .09	.68 ± .08	.64 ± .09	.81 ± .06	.73 ± .08	.73 ± .07	.87 ± .05	.88 ± .03	.76 ± .09	.77 ± .08
<i>p</i> = .016	.75 .87 .70	.53 .80 .39	.64 .66 .57	.71 .89 .00	.55 .82 .66	.50 .85 .00	.75 .92 .69	.87 .90 .70	.65 .83 .43	.66 .84 .46
Czech R.	.76 ± .09	.91 ± .04	.47 ± .08	.82 ± .08	.82 ± .05	.83 ± .07	.84 ± .06	.86 ± .05	.77 ± .13	.79 ± .13
<i>p</i> = .016	.71 .80 .31	.90 .92 .64	.52 .36 .48	.75 .89 .01	.81 .84 .40	.73 .90 .00	.73 .90 .28	.86 .88 .41	.73 .80 .27	.75 .81 .32
France	.71 ± .07	.61 ± .06	.86 ± .08	.66 ± .10	.73 ± .08	.65 ± .07	.72 ± .07	.68 ± .08	.68 ± .04	.70 ± .07
<i>p</i> = .016	.45 .83 .69	.37 .78 .45	.81 .83 .98	.37 .82 .00	.59 .81 .69	.33 .82 .00	.32 .86 .69	.48 .81 .73	.42 .82 .46	.47 .82 .53
Germany	.72 ± .10	.64 ± .09	.29 ± .12	.88 ± .11	.69 ± .09	.77 ± .09	.73 ± .10	.83 ± .07	.67 ± .16	.69 ± .17
<i>p</i> = .031	.68 .76 .01	.42 .81 .01	.42 .32 .00	.82 .93 .66	.50 .84 .00	.54 .88 .54	.47 .85 .00	.82 .88 .00	.55 .76 .08	.58 .78 .15
Italy	.55 ± .12	.76 ± .09	.63 ± .08	.78 ± .09	.95 ± .02	.73 ± .10	.53 ± .13	.74 ± .08	.67 ± .10	.71 ± .13
<i>p</i> = .047	.57 .55 .49	.74 .81 .12	.69 .63 .52	.73 .83 .00	.92 .96 .94	.66 .78 .00	.50 .54 .42	.74 .74 .63	.66 .70 .31	.69 .73 .39
Poland	.76 ± .08	.83 ± .05	.38 ± .11	.85 ± .08	.83 ± .05	.93 ± .05	.73 ± .08	.83 ± .07	.74 ± .15	.77 ± .16
<i>p</i> = .016	.66 .84 .00	.82 .89 .01	.48 .52 .00	.73 .91 .44	.80 .91 .00	.89 .95 .88	.44 .86 .00	.82 .87 .00	.68 .83 .06	.71 .84 .17
U.S.A. I	.83 ± .06	.65 ± .08	.47 ± .14	.81 ± .07	.65 ± .15	.67 ± .09	.91 ± .03	.89 ± .03	.71 ± .13	.74 ± .14
<i>p</i> = .016	.76 .87 .59	.45 .79 .49	.35 .61 .35	.71 .89 .00	.40 .80 .58	.38 .83 .00	.84 .94 .73	.87 .91 .68	.56 .81 .38	.60 .83 .43
U.S.A. II	.81 ± .08	.67 ± .06	.53 ± .15	.84 ± .10	.75 ± .10	.70 ± .08	.86 ± .05	.94 ± .02	.74 ± .11	.76 ± .12
<i>p</i> = .016	.74 .86 .65	.49 .80 .31	.50 .63 .41	.76 .91 .00	.57 .84 .75	.43 .84 .00	.72 .92 .59	.93 .96 .82	.60 .83 .39	.64 .85 .44
Pooled	.83 ± .06	.87 ± .03	.66 ± .08	.90 ± .04	.85 ± .04	.88 ± .05	.81 ± .10	.92 ± .03	.84 ± .08	
<i>p</i> = .148	.77 .86 .66	.87 .91 .03	.59 .67 .71	.87 .95 .01	.81 .89 .65	.83 .94 .01	.64 .88 .73	.91 .93 .65	.79 .88 .43	
Pooled+	.88 ± .05	.94 ± .03	.82 ± .09	.96 ± .02	.94 ± .04	.94 ± .04	.92 ± .03	.96 ± .02	.92 ± .04	
<i>p</i> = .195	.83 .91 .77	.94 .95 .70	.76 .78 .96	.94 .98 .64	.91 .95 .90	.92 .97 .65	.86 .95 .84	.95 .96 .80	.89 .93 .78	

7.2 Pooled Out-Context Experiment (H2)

The performance of Pooled Out-Context models is reported in the Pooled row of Table 4. The experiment concerns the resulting models' robustness, i.e., if a model trained on multiple contexts adapts well to unseen contexts. We are especially interested if such a model adapts better than the models trained on single contexts.

The results suggest that training on multiple contexts leads to models that are robust and perform better than the models trained on a single context. The multi-context models outperform the best single extra-context model baseline in 7 out of 8 cases. The $p = 0.148$ needs to be understood in terms of the small number of samples (contexts) and competitiveness of the baseline.

Interestingly, the Pooled Out-Context models even appear to be competitive with several In-Context models (Canada, Czech R., Germany, U.S.A. II). The overall average F₁-scores are often quite high (over 0.80 or 0.90). This is a surprising outcome considering the fact that no data from the context on which a model is evaluated is used during training.

7.3 Pooled With In-Context Experiment (H3)

The performance of the Pooled with In-Context model is reported in the Pooled+ row of Table 4. This experiment models a scenario where a sample of labeled data from the target context is available. The question is whether combining In-Context data with data from other contexts leads to improved performance.

The results appear to suggest that pooling the target context with data from other contexts does lead to improved performance. In case of 3 out of the 8 contexts (Canada, Czech R., and U.S.A. I) the improvement is clear and substantial across all the three classes (*Background*, *Analysis*, *Outcome*). For three additional contexts (Germany, Poland, and U.S.A. II), the performance also improved in terms of overall (micro) F₁-score, but took a slight hit for the challenging *Outcome* class. With respect to the two remaining contexts (France, Italy) the overall performance of the pooled models is lower than that of the In-Context models. As in the previous experiment, the $p = 0.195$ needs to be understood in terms of the small number of samples (contexts) and high competitiveness of the In-Context baseline.

Figure 4: Average LASER embeddings for each case document, projected to 2-dimensional space using Principal Component Analysis.

8 DISCUSSION

It appears that the multilingual sentence embeddings generated by the LASER model excel at capturing the semantics of a sentence. A model trained on a single context could in theory capture the specific vocabulary used in that context. This would almost certainly lead to poor generalization across different contexts. However, the performance statistics we observed when transferring the Out-Context models to other contexts suggest that the sentence embeddings provide a representation of the sentences that enable the model to learn aspects of the meaning of a sentence, rather than mere surface linguistic features.

The results clearly point to certain relationships where contexts within the same or related languages appear to work well together, e.g., {Canada, U.S.A. I, and U.S.A. II} or {Czech R. and Poland}. This could indicate that the multi-lingual embeddings work better when the language is the same or similar. It could also point to similarities in legal traditions, e.g. the use of somewhat similar sentences to indicate a transition from one type of section to the next. Note that we removed headings, which means that explicit cues could not be relied on by the models we trained on such transformed documents. Finally, the cause could also be topical (domain) similarity of the contexts (e.g., both the U.S.A. contexts deal with employment law). Also note that the above are just possible explanations. We did not perform feature importance analysis on the models.

To gain insight into this phenomenon, we visualized the relationships among the contexts on a document level. We first calculated the average sentence embedding for each document. This yielded 1024-dimensional vectors for 807 documents representing their semantics. We arranged the resulting vectors in a matrix (1024×807) and performed a Principal Component Analysis (PCA) reducing the dimensionality of the document vectors to 2. This operation enabled a convenient visualization shown in Figure 4.

Figure 5: Distribution of labels across datasets. The X-axis corresponds to unique cases, whereas the Y-axis corresponds to the normalized length of the cases. The colors correspond to the different labels of the sentences (Background, Analysis, Outcome).

Overall, the cases from the same contexts appear to cluster together. This is expected as the documents written in the same language, having similar topics, or sharing some similarities due to legal traditions, are likely to map to vectors that are closer. The documents from both the U.S.A. contexts occupy the same region in the embedding space. This is not surprising as they come from the same jurisdiction, are written in the same language, and deal with similar topics (employment law). The Canadian cases, which are also in English, occupy a nearby space. This could be linked to the language as well as likely similarities in legal traditions. French, German, and Italian documents occupy the middle space; they are closer to the English documents than those from the Czech R. and Poland. Interestingly, the Czech and Polish documents occupy almost the same space. The Polish context focuses on the rule of law while the Czech one is supposed to be more general. As the latter deals with the decisions of the top-tier courts (one of them Constitutional), it is possible that the topics substantially overlap. Moreover, Poland and the Czech R. share similar legal traditions and languages from the same family (Slavic), so the close proximity of the documents in the embedding space might not be unexpected. Finally, German, French, and Canadian cases occupy wider areas than documents from other contexts. This could be due to their lack of focus on specific legal domains.

We observed a peculiar phenomenon where the Out-Context models trained on the German and Polish contexts failed to detect *Outcome* sentences on the six remaining contexts and vice-versa. The cause is readily apparent from the visualization shown in Figure 5. Each segment of the figure depicts a spatial distribution of sentences color coded with their labels across documents for a particular context. As can be seen, the cases typically follow a pattern of a contiguous *Background* segment, followed by a long *Analysis* section. The several *Outcome* sentences are placed at the very end of the documents. In the Polish and German decisions, however, the *Outcome* sentences come first. The GRU models we use rely on the structure as well as semantics in making their predictions. As we can see, a model trained exclusively on cases that begin with a *Background* might therefore have difficulties correctly identifying outcome sections at the beginning, and vice-versa. However, as we

will see below, a model trained with data featuring both structures can learn to correctly identify the correct structure based on the semantics of the sentences.

The model trained on the French context appears to perform better on detecting *Outcome* sentences than models trained on other contexts. This is somewhat surprising as the French model's overall performance is among the weakest (e.g. compare the Czech model's average $F_1 = 0.77$ to the $F_1 = 0.68$ of the French model). Again, Figure 5 provides an insight into why this happens. The French context is the only one where the count of *Outcome* sentences is comparable to those of the other two categories. For all the other contexts, the *Outcome* sentences are heavily underrepresented. This reveals an interesting direction for future work where the use of re-sampling may yield models with better sensitivity for identifying *Outcome* sentences.

In two instances models trained on a single context under-performed the **random baseline**. The model trained on the Germany context achieved the average $F_1 = 0.29$ when applied to the French context (Random $F_1 = 0.33$). As the model trained on Polish data also performed poorly on the French context ($F_1 = 0.38$) the cause appears to be the inability of the two models to detect the *Outcome* sentences at the end (discussed above). As the *Outcome* sentences are heavily present in the French context, this problem manifests in a score lower than the random baseline. The second instance is the model trained on the Italy context applied to the U.S.A. I data ($F_1 = 0.53$ versus $F_1 = 0.58$ on Random). Here, the cause appears to be different. Note that the Italian context likely has a very specific notion of *Outcome* sentences ($F_1 = 0.94$ on Italy→Italy). It appears that many *Analysis* sentences from the U.S.A. I context were labeled as *Outcome* by the model. Summary judgments often address multiple legal issues with their own conclusions which could have triggered the model to label such sentences as *Outcome*.

An important finding is the performance of the Pooled Out-Context model (H2) shown in the Pooled row of Table 4. The experiment simulates training of a model on several contexts, and then applying it to an unseen context. The Pooled Out-Context models, having no access to the data from a target context, reliably outperform the **best single Out-Context models**. They appear to be competitive with several **In-Context models**. These results are achieved with a fairly small dataset of 807 cases. We expect that expanding the dataset would lead to further improved performance.

The Pooled with In-Context experiment (H3) models the situation where data from a target context is available in addition to labeled data from other contexts. Our experiments indicate that the use of data from other contexts (if available) in addition to data from the target context is preferable to the use of the data from the target context only. This is evidenced by the improved performance of the models trained on the pooled contexts over the single **In-Context models**. The models have an interesting property of being able to identify the *Outcome* sentences with effectiveness comparable to (or higher than) the models trained on the same context only. This holds for all the contexts, except Poland where the *Outcome* performance is a bit lower (0.65 vs. 0.88). This indicates that the model is able to learn the two possible modes of the *Outcome* section placement. It successfully distinguishes cases where the section is at the beginning from the cases where the *Outcome* sentences are found toward the end.

The inclusion of the In-Context data in the pooled data leads to a remarkable improvement over only using the pooled Out-Context data. The magnitude of the improvement highlights the importance of including such data in the training. We envision that the models trained on different contexts used in combination with high-speed similarity annotation frameworks [35, 36] could enable highly cost efficient annotation in situations where resources are scarce. Perhaps, adapting a model to an entirely new context could be as simple as starting with a model trained on other contexts, and spending a few hours correcting the misconceptions of the model to teach it the particularities of the new context.

9 CONCLUSIONS

We analyzed the use of multi-lingual sentence embeddings in sequence labeling models to enable transfer across languages, jurisdictions, legal systems (common and civil law), and domains. We created a new type schema for functional segmentation of adjudicatory decisions and used it to annotate legal cases across eight different contexts. We found that models generalize beyond the contexts they were trained on and that training the models on multiple contexts increases their robustness and improves the overall performance when evaluating on previously unseen contexts. We also found that pooling the training data of a model with data from additional contexts enhances its performance on the target context. The results are promising in enabling re-use of annotated data across contexts and creating generalizable and robust models. We release the newly created dataset (807 documents with 89,661 annotated sentences), including the annotation schema and the code used in our experiments, to the public.

This work suggests a promising path for the future of international collaboration in the field of AI & Law. While previous annotation efforts have typically been limited to a single context, the experiments presented here suggest that researchers can work together by annotating cases from many different contexts at the same time. Such a combined effort could aid researchers in creating models that perform well on the data from the context they care about, while at the same time helping other groups train even better models for other contexts. We encourage these research directions and hope to form such collaborations under the Lex Rosetta project.

10 FUTURE WORK

The application of multi-lingual sentence embeddings to functional segmentation of case law across different contexts yielded promising results. At the same time, the work is subject to limitations and leaves much room for improvement. Hence, we suggest several directions for future work:

- Extension of the datasets from different contexts used in this work beyond ~100 documents per context.
- Annotation of data from contexts beyond the eight used here (multi-lingual models support close to 100 languages).
- Analysis of automatic detection of *Introductory Summary*, *Headings*, and *Out of Scope*.
- Identification and investigation of other tasks applicable across different contexts.
- Evaluation of the application of other multilingual models (e.g., those mentioned in Section 2).

- Exploring other transfer learning strategies beyond simple data pooling, such as the framework proposed in [26].
- Using multi-lingual models for annotation tasks with high-speed annotation framework, such as [35, 36].
- Performing the transfer across contexts with related (but different) tasks, such as in [28].
- Further exploring the differences in the distribution of the multilingual embeddings for purposes of comparing and analyzing domains, languages, and legal traditions.

ACKNOWLEDGMENTS

Hannes Westermann, Karim Benyeklef, and Kevin D. Ashley would like to thank the Cyberjustice Laboratory at Université de Montréal, the LexUM Chair on Legal Information and the Autonomy through Cyberjustice Technologies (ACT) project for their support of this research. Kevin D. Ashley also thanks the Canadian Legal Information Institute for providing the corpus of legal cases. Matthias Grabmair thanks the SINC GmbH for supporting this research. Jakub Harašta and Tereza Novotná acknowledge the support of the ERDF project “Internal grant agency of Masaryk University” (No. CZ.02.2.69/0.0/0.0/19_073/0016943).

REFERENCES

- [1] Tommaso Agnoloni, Lorenzo Bacci, Enrico Francesconi, P Spinoso, Daniela Tiscornia, Simonetta Montemagni, and Giulia Venturi. 2007. Building an ontological support for multilingual legislative drafting. *Frontiers in Artificial Intelligence and Applications* 165 (2007), 9.
- [2] Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. Massively Multilingual Neural Machine Translation. In *NAACL-HLT, Vol. 1 (Long and Short Papers)*. 3874–3884.
- [3] Gianmaria Ajani, Guido Boella, Luigi Di Caro, Livio Robaldo, Llio Humphreys, Sabrina Praduroux, Piercarlo Rossi, and Andrea Violato. 2016. The European Legal Taxonomy Syllabus: A multi-lingual, multi-level ontology framework to untangle the web of European legal terminology. *Applied Ontology* 11, 4 (2016).
- [4] Sai Saket Aluru, Binny Mathew, Punyajoy Saha, and Animesh Mukherjee. 2020. Deep learning models for multilingual hate speech detection. *arXiv preprint arXiv:2004.06465* (2020).
- [5] Mikel Artetxe and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics* 7 (2019), 597–610.
- [6] Paheli Bhattacharya, Shounak Paul, Kripabandhu Ghosh, Saptarshi Ghosh, and Adam Wyner. 2019. Identification of rhetorical roles of sentences in Indian legal judgments. In *JURIX 2019*, Vol. 322. IOS Press, 3.
- [7] Guido Boella, Luigi Di Caro, Michele Graziadei, Loredana Cupi, Carlo Emilio Salargio, Llio Humphreys, Hristo Konstantinov, Kornel Marko, Livio Robaldo, Claudio Ruffini, et al. 2015. Linking legal open data: breaking the accessibility and language barrier in european legislation and case law. In *ICAIL 2015*. 171–175.
- [8] Paul Boniol, George Panagopoulos, Christos Xypolopoulos, Rajaa El Hamdani, David Restrepo Amariles, and Michalis Vazirgiannis. 2020. Performance in the Courtroom: Automated Processing and Visualization of Appeal Court Decisions in France. In *Proceedings of the Natural Legal Language Processing Workshop 2020*.
- [9] Karl Branting, Brandy Weiss, Bradford Brown, Craig Pfeifer, A Chakraborty, Lisa Ferro, M Pfaff, and A Yeh. 2019. Semi-supervised methods for explainable legal prediction. In *ICAIL 2019*. 22–31.
- [10] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP 2014*.
- [11] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *ACL 2020*. 8440–8451.
- [12] Janez Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research* 7, Jan (2006), 1–30.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL 2019, Volume 1 (Long and Short Papers)*. 4171–4186.
- [14] Luca Dini, Wim Peters, Doris Liebwald, Erich Schweighofer, Laurens Mommers, and Wim Voermans. 2005. Cross-lingual legal information retrieval using a WordNet architecture. In *ICAIL 2005*. 163–167.
- [15] Atefeh Farzindar and Guy Lapalme. 2004. LetSum, an Automatic Text Summarization system in Law field. *JURIX 2004*.
- [16] Jorge González-Conejero, Pompeu Casanovas, and Emma Teodoro. 2018. Business Requirements for Legal Knowledge Graph: the LYNX Platform.. In *TERECOM@JURIX 2018*. 31–38.
- [17] Jakub Harašta, Jaromír Šavelka, František Kasl, and Jakub Míšek. 2019. Automatic Segmentation of Czech Court Decisions into Multi-Paragraph Parts. *Jusletter IT* 4, M (2019).
- [18] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [19] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [20] Philipp Koehn, Alexandra Birch, and Ralf Steinberger. 2009. 462 Machine Translation Systems for Europe. In *Proceedings of the Twelfth Machine Translation Summit*. Association for Machine Translation in the Americas, 65–72.
- [21] Guokun Lai, Barlas Oguz, Yiming Yang, and Veselin Stoyanov. 2019. Bridging the domain gap in cross-lingual document classification. *arXiv preprint arXiv:1909.07009* (2019).
- [22] Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. MLQA: Evaluating Cross-lingual Extractive Question Answering. In *ACL 2020*. 7315–7330.
- [23] D.N. MacCormick, R.S. Summers, and A.L. Goodhart. 2016. *Interpreting Precedents: A Comparative Study*. Taylor & Francis.
- [24] Rohan Nanda, Llio Humphreys, Lorenzo Grossio, and Adebayo Kolawole John. 2020. Multilingual Legal Information Retrieval System for Mapping Recitals and Normative Provisions. In *Proceedings of Jurix 2020*. IOS Press, 123–132.
- [25] Alina Petrova, John Armour, and Thomas Lukasiewicz. 2020. Extracting Outcomes from Appellate Decisions in US State Courts. In *JURIX 2020*. 133.
- [26] Jaromír Šavelka and Kevin D Ashley. 2015. Transfer of predictive models for classification of statutory texts in multi-jurisdictional settings. In *ICAIL 2015*. 216–220.
- [27] Jaromír Šavelka and Kevin D Ashley. 2018. Segmenting US Court Decisions into Functional and Issue Specific Parts.. In *JURIX 2018*. 111–120.
- [28] Jaromír Šavelka, Hannes Westermann, and Karim Benyeklef. 2020. Cross-Domain Generalization and Knowledge Transfer in Transformers Trained on Legal Data. In *ASAIL@ JURIX 2020*.
- [29] Párai Sheridan, Martin Bräschler, and Peter Schauble. 1997. Cross-language information retrieval in a Multilingual Legal Domain. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 253–268.
- [30] Ralf Steinberger, Mohamed Ebrahim, Alexandros Poulis, Manual Carrasco-Benitez, Patrick Schluter, Marek Przybyszewski, and Signe Gilbro. 2014. An overview of the European Union’s highly multilingual parallel corpora. *Language Resources and Evaluation* 48, 4 (2014), 679–707.
- [31] Kyoko Sugisaki, Martin Volk, Rodrigo Polanco, Wolfgang Alschner, and Dmitriy Skougarevskiy. 2016. Building a Corpus of Multi-lingual and Multi-format International Investment Agreements. In *JURIX 2016*.
- [32] Linyuan Tang and Kyo Kageura. 2019. Verifying Meaning Equivalence in Bilingual International Treaties. In *JURIX 2019*. 103–112.
- [33] Vern R Walker, Krishnan Pillaipakkamatt, Alexandra M Davidson, Marysa Linares, and Domenick J Pesce. 2019. Automatic Classification of Rhetorical Roles for Sentences: Comparing Rule-Based Scripts with Machine Learning.. In *ASAIL@ ICAIL 2019*.
- [34] Hannes Westermann, Jaromír Šavelka, and Karim Benyeklef. 2021. Paragraph Similarity Scoring and Fine-Tuned BERT for Legal Information Retrieval and Entailment. In *New Frontiers in Artificial Intelligence (Lecture Notes in Computer Science)*. Springer International Publishing.
- [35] Hannes Westermann, Jaromír Šavelka, Vern R Walker, Kevin D Ashley, and Karim Benyeklef. 2019. Computer-Assisted Creation of Boolean Search Rules for Text Classification in the Legal Domain. In *JURIX 2019*, Vol. 322. IOS Press, 123.
- [36] Hannes Westermann, Jaromír Šavelka, Vern R Walker, Kevin D Ashley, and Karim Benyeklef. 2020. Sentence Embeddings and High-Speed Similarity Search for Fast Computer Assisted Annotation of Legal Documents. In *JURIX 2020*, Vol. 334. IOS Press, 164.
- [37] Frank Wilcoxon. 1992. Individual comparisons by ranking methods. In *Breakthroughs in statistics*. Springer, 196–202.
- [38] Huihui Xu, Jaromír Šavelka, and Kevin D Ashley. 2020. Using Argument Mining for Legal Text Summarization. In *JURIX 2020*, Vol. 334. IOS Press.
- [39] Shudong Yang, Xueying Yu, and Ying Zhou. 2020. LSTM and GRU Neural Network Performance Comparison Study: Taking Yelp Review Dataset as an Example. In *IWECAI 2020*. IEEE, 98–101.
- [40] Vladimir Zhebel, Denis Zubarev, and Ilya Sochenkov. 2020. Different Approaches in Cross-Language Similar Documents Retrieval in the Legal Domain. In *International Conference on Speech and Computer*. Springer, 679–686.
- [41] Linwu Zhong, Ziyi Zhong, Zinian Zhao, Siyuan Wang, Kevin D Ashley, and Matthias Grabmair. 2019. Automatic summarization of legal decisions using iterative masking of predictive sentences. In *ICAIL 2019*. 163–172.