Unlocking Legal Knowledge: A Multilingual Dataset for Judicial Summarization in Switzerland

Anonymous ACL submission

Abstract

Legal research is a time-consuming task that most lawyers face on a daily basis. A large part of legal research entails looking up relevant caselaw and bringing it in relation to the case at hand. Lawyers heavily rely on summaries (also called headnotes) to find the right cases quickly. However, not all decisions are annotated with headnotes and writing them is time-consuming. Automated headnote creation has the potential to make hundreds of thousands of decisions more accessible for legal research in Switzerland alone. To kickstart this, we introduce the Swiss Leading Decision Summarization (SLDS) dataset, a novel cross-lingual resource featuring 18K court rulings from the Swiss Federal Supreme Court (SFSC), in German, French, and Italian, along with German headnotes. We fine-tune and evaluate three mT5 variants, along with proprietary models. Our analysis highlights that while proprietary models perform well in zero-shot and one-shot settings, fine-tuned smaller models still provide a strong competitive edge. We publicly release the dataset to facilitate further research in multilingual legal summarization and the development of assistive technologies for legal professionals.

1 Introduction

A significant part of legal work involves research, where lawyers must find similar cases and navigate numerous judicial decisions, especially when interpreting laws with room for debate. Due to the time-intensive nature of this task, they usually rely on judgment summaries. However, creating these summaries is labor-intensive and requires expertise from judges and clerks, who are already burdened with heavy caseloads (Bieri, 2015) and time-pressure (Ludewig and Lallave, 2013).

To alleviate this increasing need for efficient ways to navigate vast amounts of legal documents,

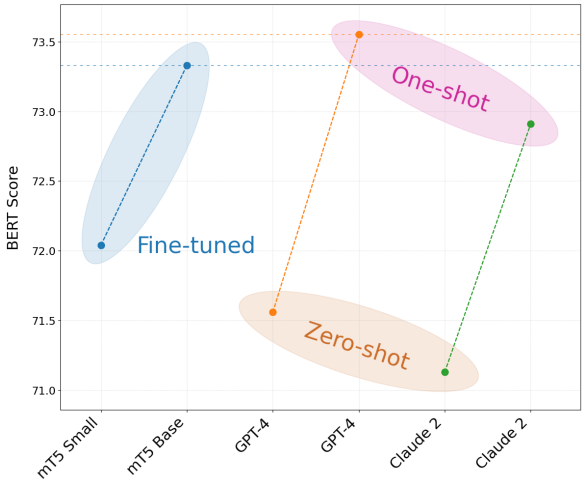

Figure 1: Small fine-tuned Language Models (LMs) are still competitive when compared to large proprietary models with few-shot prompting when summarizing judicial decisions.

legal document summarization has become a critical area of interest in NLP (Jain et al., 2021). Over the years, researchers have made significant strides in both extractive and abstractive summarization of legal texts. Earlier works focused on extracting key sentences to create concise summaries (Grover et al., 2004; Hachey and Grover, 2006; Kim et al., 2013; Bhattacharya et al., 2021), while recent advancements have turned towards abstractive methods, which generate condensed paraphrases of the most important information in a document (Shukla et al., 2022; Niklaus and Giofré, 2022; Moro et al., 2023; Jain et al., 2024; Niklaus et al., 2024).

Datasets with legal documents and their corresponding summaries have been instrumental in enabling these advancements, yet they primarily focus on monolingual corpora. Therefore, existing datasets do not adequately address the unique challenges posed by multilingual jurisdictions, such as Switzerland, where legal decisions are written in multiple languages and need to be summarized consistently. This gap is significant because most legal NLP tools and models are trained on English-centric datasets, which often do not generalize well

to cross-lingual environments.

To fill this gap, we introduce the Swiss Leading Decision Summarization (SLDS) dataset, a large collection of Swiss legal cases in German, French, and Italian, created for legal summarization. In Switzerland, lawyers often rely on leading decisions Leading Decisions (LDs) from the Swiss Federal Supreme Court (SFSC), which come with headnotes – short summaries that highlight key legal points and important laws. The SLDS dataset focuses on these LDs, which explain how Swiss laws are interpreted and applied. It is designed to help develop tools that assist legal professionals working in multilingual environments. It also enables training language models to handle the complexities of legal text in different languages, making it a valuable resource for research in cross-lingual legal summarization. The dataset is available on Hugging Face under a CC-BY-4.0 license[1].

Contributions Our contributions are two-fold:

1. **SLDS dataset release:** We introduce and publicly release the SLDS dataset, a cross-lingual legal resource comprising 18K rulings from the SFSC in German, French, and Italian, along with corresponding German summaries.
2. **Comprehensive benchmarking:** We evaluate three fine-tuned mT5 variants and four proprietary Large Language Models (LLMs), establishing detailed performance baselines for multilingual legal summarization. Our results show that small mT5 models can outperform proprietary models in summarizing legal texts, emphasizing the effectiveness of domain-specific fine-tuning over generic, zero-shot models in specialized legal tasks.

2 Related Work

Recently, research on legal text summarization has increasingly shifted towards abstractive summarization, with a particular emphasis on creating datasets that facilitate fine-tuning of pre-trained language models for this task. In the following section, we will discuss the most relevant contributions.

2.1 Monolingual Datasets

English Two English legal text summarization datasets with a focus on U.S. legislative particularly stand out. The first one is BillSum (Kornilova and Eidelman, 2019) comprising 22K bills from the U.S. Congress and the State of California. The authors evaluated different extractive algorithms on their dataset. They trained two binary classifiers, a Random Forest (Breiman, 2001) based on hand-crafted features and TF-IDF (Luhn, 1957; Sparck Jones, 1972) scores, and a BERT (Kenton and Toutanova, 2019) encoder model, whose purpose is to predict whether a sentence should be included in the summary. To create the summary, they use the Maximal Marginal Relevance (Goldstein et al., 2000) algorithm. The authors evaluate both classifiers, an ensemble of both classifiers, along with two other unsupervised baselines on their dataset using ROUGE scores (Lin, 2004a). They find that the ensemble performs best, with most of the performance coming from BERT. Moreover, they applied transfer learning in summarization from federal to state laws.

The second dataset is Multi-LexSum (Shen et al., 2022), which targets long civil rights lawsuits, with an average length of over 75K words. The corpus consists of 9K documents and allows for in-depth study at different summary lengths: short (25 words), medium (130 words), and long (650 words), a unique feature of the Multi-LexSum dataset. Models based on BART (Lewis et al., 2020) and PEGASUS (Zhang et al., 2020a) were evaluated on this dataset. There are also larger datasets. A significant contribution in the recent past comes from Bauer et al. (2023) who focused on extracting key passages from 430K U.S. court opinions to create concise summaries. Interestingly, they find that their reinforcement-learning-based MemSum model out-performs two transformer-based models that were also evaluated. However, due to licensing issues, the authors were not able to release the dataset their models were trained on.

Portuguese RulingBR (de Vargas Feijó and Moreira, 2018) is another monolingual dataset for legal document summarization. It consists of over 10K rulings from the Brazilian Federal Surpreme Court, along with the corresponding summaries in Portuguese. The summaries are divided into four parts that represent different types of information, with one of them being the reference summary that can be used to train a summarization model. The authors also published evaluations on extractive summarization algorithms that serve as a baseline.

[1] Link available upon acceptance

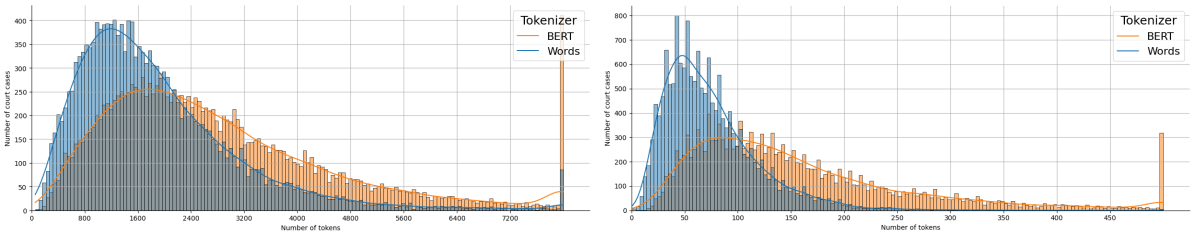

(a) Decisions (b) Summaries

Figure 2: Distributions for the text length of the decisions (input) and summaries (output) in the SLDS

2.2 Multilingual Datasets

There exists also multilingual corpora. Below, we briefly describe two of these datasets.

EUR-Lex-Sum Derived from EUR-Lex, the official online repository of European Union law, EUR-Lex-Sum (Aumiller et al., 2022) encompasses the 24 official languages of the European Union and features more than 1,500 document/summary pairs per language. A subset of 375 legal acts is cross-lingually aligned and available in all 24 languages. Although legal acts tend to be dense and technical, they generally adhere to a more structured format compared to court decisions, which are characterized by complex reasoning, factual narratives, and often involve multiple legal perspectives.

A significant distinction between our work and EUR-Lex-Sum lies in our focus on court cases within a single jurisdiction, as opposed to their focus on legal acts from a supranational entity. While summarizing international laws is valuable, understanding their interaction with national laws and their application to specific cases is just as important, if not more so. Our dataset prioritizes understanding how laws are applied in a national context, as opposed to simply detailing what the laws are. Additionally, SLDS provides over 13 times more samples for French-to-German cross-lingual summarization and more than double the number of Italian-to-German samples. This enables a more comprehensive analysis of cross-lingual summarization for these specific languages, which hold particular relevance in the Swiss context.

MILDSum The dataset, introduced by Datta et al. (2023), was designed to address language barriers in the Indian legal system, focusing on cross-lingual translation of English legal documents into Hindi. MILDSum comprises 3,000 Indian legal case judgments in English, along with their corresponding summaries in both English and Hindi. A key finding by the authors is that the *Summarize-then-Translate* approach outperforms direct cross-lingual summarization. The average lengths of the case judgments and summaries are 4,696 tokens for the judgments, and 724 and 695 tokens for the English and Hindi summaries, respectively.

While MILDSum and our dataset share a focus on summarizing court decisions and cross-linguality, there are notable differences. Our dataset includes three instead of two languages, and none of them is English. This likely makes the summarization task in SLDS more challenging, given that English dominates most pre-training datasets. Another distinction lies in the structure of the summaries: our dataset contains headnotes, which include citations, keywords, and summarized decisions, whereas MILDSum does not mention any similar structure for their summaries.

3 Data

We present SLDS, a novel dataset for cross-lingual summarization in the legal domain. The dataset consists of 18,175 leading decisions in German, French or Italian published by the SFSC, along with their corresponding summaries in German written by clerks and judges. It serves as an important resource to study cross-lingual summarization, which is relatively underexplored in the legal domain. In contrast to datasets like EUR-Lex-Sum that focus on legislation, our dataset contains judicial decisions, making it especially relevant for developing tools aimed at supporting legal practitioners and researchers working with court rulings.

3.1 Data Collection

The original data was published on the official website[2] of the SFSC. These leading decisions are then scraped by the platform Entscheidsuche[3], from where we downloaded them.

3.2 Fields

The dataset contains the following fields:

[2]Available at https://www.bger.ch/
[3]Available at https://entscheidsuche.ch/

- `id`: identifier for a specific decision
- `header`: a short header containing metadata about the case such as the case citation, the division that handled the case, date of judgment, parties involved, and the type of complaint
- `text`: the decision
- `summary`: headnote (*Regeste* in German) that summarizes the decision. Consist of three components: i) important citations to laws and prior cases, ii) thematic keywords derived from a legal thesaurus, and iii) a free-form summary of key considerations.
- `language`: the language of the decision, i.e. the column `text`
- `chamber`: an identifier for one out of seven chambers of the SFSC

3.3 Descriptive Statistics

Dataset Splits The dataset is divided into *train*, *validation*, and *test* subsets. The training set contains 16.5K samples, while the validation and test sets have 500 and 1000 samples, respectively. We partitioned the data by year: training includes decisions from 1954 to 2015, validation from 2016 to 2017, and testing from 2018 to 2022. This was done to prevent data leakage and ensures consistency with current summarization styles.

Languages The decisions, which are written by clerks and judges in the language of the proceedings, are either in German, French or Italian. There are roughly 12K German decisions, almost 5K in French, and 835 in Italian. The decisions are stored under the column *text*. The headnotes are in the column *summary*, all written in German.

Text Length The distributions of the number of words and tokens, when using a pre-trained BERT tokenizer, for the decisions and their headnotes are shown in Figure 2. To save space, we removed the tail of the distribution by truncating all texts to a maximum token and word length of 8000, resulting in a peak at the end of the x-axis.

Decision Year The judgments in this dataset were made between 1954 and 2022, with only little data available for the years 2001 to 2006. The number of decisions per year ranges from 2 to 420, with a median of 278 decisions per year. The exact distribution can be seen in Appendix C.

```
Objective: Generate a headnote based on a Swiss court
↪  judgment.
Background: A Swiss court judgment consists of facts,
↪  recitals and dispositive. The headnote serves as a
↪  brief summary and contains the guiding principles of
↪  the judgment. Only leading decisions have a headnote.
Instruction:
1. facts: Read and understand the given facts.
2. considerations: Analyze the recitals to identify the
↪  main arguments and causes.
3. dispositive: Pay attention to the dispositive, as it
↪  contains the final judgment.
4. compile the headnote: The headnote should consist of
↪  three very short parts: a. Cite the most important
↪  relevant article numbers (without the article title).
↪  b. List short, relevant, descriptive keywords about
↪  the subject matter of the case. c. Formulate a very
↪  short body text that cites and briefly summarizes the
↪  most important considerations.
Output: The headnote should provide a clear and
↪  structured summary of the judgment, consisting of
↪  cited articles, keywords and a very short text.

Here is an example of a court judgment with a headnote:

[Given facts, recitals and dispositive]:
{ Specific decision in German, French or Italian }

[Headnote in German]: { Headnote for that specific
↪  decision }

[Given facts, recitals and dispositive]:
{ Decision from the validation set }

[Headnote in German]:
```

Prompt 1: A translated version of the prompt that was used for the zero-shot and one-shot experiments. The blue part is only added in case of a one-shot setting.

3.4 Licensing

We release the data under the CC-BY-4.0 license, which complies with the SFSC licensing[4].

3.5 Ethical Considerations

Due to the sensitive nature of court cases and their corresponding rulings, the SFSC anonymizes personal or sensitive information according to their guidelines[5] before publishing them online.

4 Experimental Setup

To establish baselines, we fine-tuned the two variants *small* and *base* of mT5 (Xue, 2020), a multilingual variant of the T5 model (Raffel et al., 2020), on the training subset of our dataset. We then evaluated the models using **BERT**Score (Zhang et al., 2020b), **BLEU** (Papineni et al., 2002), **METEOR** (Banerjee and Lavie, 2005), and **ROUGE** (Lin, 2004b). Since each individual metric has inherent weaknesses (Zhang et al., 2020b), it is necessary to employ multiple metrics for a more comprehen-

[4]For more information, see https://www.bger.ch/files/live/sites/bger/files/pdf/de/urteilsveroeffentlichung_d.pdf

[5]Anonymization guidelines at https://www.bger.ch/home/juridiction/anonymisierungsregeln.html

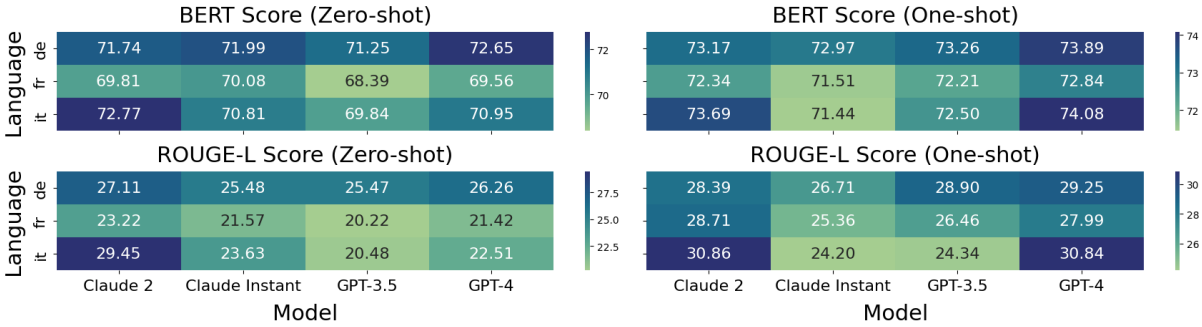

Figure 3: Few-shot evaluation of the proprietary models used in our experiments on decisions in different languages.

Model	Setting	BERT ↑	BLEU ↑	MET ↑	R1 / R2 / RL ↑
$mT5_{Base}$	fine-tuned	**73.33**	**30.81**	**23.50**	**32.43 / 12.78 / 30.87**
$mT5_{Small}$	fine-tuned	72.04	28.68	21.29	29.61 / 10.31 / 28.12
GPT-4	one-shot	**73.55**	**47.75**	**34.72**	**30.82 / 9.68 / 28.89**
GPT-3.5-Turbo-16K	one-shot	72.89	45.21	32.76	29.69 / 9.25 / 27.94
Claude 2	one-shot	72.91	47.55	33.57	30.28 / 9.12 / 28.58
Claude Instant	one-shot	72.44	44.80	30.29	27.89 / 8.56 / 26.18
GPT-4	zero-shot	**71.56**	48.35	**32.97**	26.52 / **8.93** / 24.51
GPT-3.5-Turbo-16K	zero-shot	70.28	46.08	30.60	25.18 / 7.58 / 23.59
Claude 2	zero-shot	71.13	**49.20**	32.54	**27.70** / 8.39 / **25.90**
Claude Instant	zero-shot	71.33	45.65	29.22	26.13 / 8.16 / 24.15

Table 1: Results of the experiments with a maximum input length of 4096 tokens. The input is truncated to 4096 tokens. **Bold**: best within setup; underlined: best overall.

sive assessment. Besides this quantitative evaluation, we suggest evaluating predictions with trained lawyers as future work.

We also compared the two fine-tuned models against GPT-3.5-Turbo, GPT-4, Claude 2, and Claude Instant in a zero-shot and one-shot setting. The original prompt and a translated version can be found in Appendix F. We used a German prompt because the headnotes are always in German in our dataset, and we found that models tend to generate in the prompt's language, regardless of the input's language. The decisions in our dataset consist of 3081 tokens on average, while the headnotes average 168 tokens. The maximum sequence length for the decisions is 38K, and the 90th percentile is 5646 tokens. For our main experiment, we truncated the inputs to 4096 tokens and the output to 256 tokens, preserving the full output in 82% of cases. Additionally, we performed an extended version of the experiment with input sequence lengths ranging from 512 to 16K tokens. In this extended version, we also evaluated the *large* variant of mT5[6].

[6]Unfortunately, we were not able to go up to a maximum sequence length of 4096 tokens during fine-tuning due to memory limitations.

5 Results

5.1 Main Experiment

Overall Results The proprietary models perform well on SLDS in the zero-shot and one-shot setting, at least in terms of BLEU and METEOR (see Table 1). According to the ROUGE scores, the fine-tuned mT5 models outperform the pre-trained proprietary models, while BERT-Score does not discriminate clearly. We assume that summarization is a large portion of internal instruction tuning datasets used for optimizing these models.

In general, the generated summaries are of high quality. They demonstrate a good stylistic imitation of legal language and consistent logical coherence. To see examples of three generated summaries by our fine-tuned $mT5_{Base}$ model, refer to Appendix G. Summaries generated by the proprietary models can be found in Appendix I. GPT-3.5-Turbo (zero-shot) offered a narrative-style summary, while others adhered to the traditional format of headnotes in Switzerland. Notably, GPT-3.5-Turbo made a factual error by negating a crucial element, and Claude 2 referenced an outdated legal provision. We conduct a detailed error analysis in Section 6.

Model	Setting	In Len	BERT ↑	BLEU ↑	MET ↑	R1 / R2 / RL ↑
$mT5_{Large}$	fine-tuned	2048	73.10	27.21	21.88	31.47 / 12.22 / 29.94
$mT5_{Large}$	fine-tuned	512	70.67	26.89	18.31	24.76 / 6.15 / 23.48
$mT5_{Base}$	fine-tuned	4096	73.33	30.81	23.50	**32.43** / **12.78** / **30.87**
$mT5_{Base}$	fine-tuned	2048	72.45	30.13	21.94	30.09 / 10.79 / 28.71
$mT5_{Base}$	fine-tuned	512	70.60	27.10	18.31	24.72 / 6.15 / 23.55
$mT5_{Small}$	fine-tuned	4096	72.04	28.68	21.29	29.61 / 10.31 / 28.12
$mT5_{Small}$	fine-tuned	2048	71.38	24.64	19.28	27.88 / 9.19 / 26.54
$mT5_{Small}$	fine-tuned	512	69.66	20.73	15.95	22.91 / 5.36 / 21.85
GPT-4	one-shot	4096	**73.55**	47.75	34.72	30.82 / 9.68 / 28.89
GPT-3.5-Turbo-16K	one-shot	8192	72.92	46.15	33.68	29.69 / 9.47 / 27.97
GPT-3.5-Turbo-16K	one-shot	4096	72.89	45.21	32.76	29.69 / 9.25 / 27.94
Claude 2	one-shot	16384	73.21	**49.79**	**35.58**	31.19 / 9.76 / 29.44
Claude 2	one-shot	4096	72.91	47.55	33.57	30.28 / 9.12 / 28.58
Claude Instant	one-shot	16384	72.23	46.29	31.05	28.26 / 8.78 / 26.42
Claude Instant	one-shot	4096	72.44	44.80	30.29	27.89 / 8.56 / 26.18
GPT-4	zero-shot	4096	71.56	48.35	32.97	26.52 / 8.93 / 24.51
GPT-3.5-Turbo-16K	zero-shot	4096	70.28	46.08	30.60	25.18 / 7.58 / 23.59
Claude 2	zero-shot	16384	71.45	49.29	33.51	28.61 / 8.84 / 26.59
Claude 2	zero-shot	4096	71.13	49.20	32.54	27.70 / 8.39 / 25.90
Claude Instant	zero-shot	16384	70.64	45.19	28.92	26.23 / 7.68 / 24.28
Claude Instant	zero-shot	4096	71.33	45.65	29.22	26.13 / 8.16 / 24.15

Table 2: Results of the extended experiments with different maximum sequence lengths. *In Len* denotes input length in tokens. Underlined: best within setting; **bold**: best overall. All models, except for the mT5 models, were evaluated on the validation set.

Language-specific Results Figure 3 shows the evaluation of proprietary models on German, French, and Italian decisions (detailed in Appendix H). GPT-4 (one-shot) leads in BERTScore across all languages. In the zero-shot setting, the best model varies: GPT-4 for German, Claude Instant for French, and Claude 2 for Italian. Claude 2 competes closely with GPT-4, often surpassing it in zero-shot and achieving more exact matches in one-shot, especially in French and Italian, as reflected by higher ROUGE and BLEU scores.

5.2 Extended Experiment

Table 2 presents the extended experiment results, revealing two key trends: First, longer inputs generally improve scores across models, as truncated inputs miss information, leading to lower evaluation scores. Claude Instant was an exception, showing worse performance with longer inputs in both zero-shot and one-shot settings. Second, larger mT5 models outperform smaller ones, though the difference between *base* and *large* is subtle. For proprietary models, one-shot settings yielded slightly better scores than zero-shot, as anticipated.

6 Error Analysis

Automatic metrics have limits in evaluating text generation tasks (Schluter, 2017; Zhang et al., 2020b). To address this and gain clearer insights, a legal expert on this project conducted an error analysis. This analysis, based on a reference sample (Figure 4) and generated summaries from proprietary models like GPT-4 (Figure 5), and others (see Appendix I), looked at a case where the SFSC ruled against a wife's appeal on her husband's reduced maintenance, as she couldn't prove his unemployment was meant to harm her.

All models demonstrated an authentic stylistic representation. GPT-3.5-Turbo provides a more narrative-style summary, while the other models stick more closely to a classic Swiss headnote format. Claude 2 cited an outdated legal provision, Article 137 of the ZGB, inactive since 2011, even though the case decision was from 2017 and no such reference was in the input. Despite this, the article cited remains relevant to the context. Meanwhile, GPT-4's reference to BGE 137 III 118 and GPT-3.5's mention of BGE 143 III 233 S. 235 were accurate and existed in the input.

Reference Sample

Decision (Input)

Original: Sachverhalt ab Seite 234 BGE 143 III 233 S. 234 A. Mit Eingabe vom 22. März 2013 begehrte B. (Ehemann) beim Zivilgericht Basel-Stadt die Scheidung der Ehe mit A. (Ehefrau). Mit Entscheid vom 1. Mai 2013 wurde er verpflichtet, der Ehefrau für die Dauer der Aufhebung des gemeinsamen Haushalts einen monatlichen Unterhaltsbeitrag von Fr. 1'520.- bis 31. August 2013 und von Fr. 3'000.- ab September 2013 zu entrichten. B. B.a Am 2. August 2015 ersuchte der Ehemann um Reduktion des Unterhaltsbeitrages für die Ehefrau auf monatlich Fr. 1'500.- resp. Fr. 500.-. Die Ehefrau widersetzte sich diesem Begehren. Mit Entscheid vom 10. November 2015 hob der Instruktionsrichter des Zivilgerichts Basel-Stadt den monatlichen Unterhaltsbeitrag für die Ehefrau für die Dauer des Scheidungsverfahrens rückwirkend per 1. Juli 2015 auf. B.b Am 1. März 2016 hiess der Ausschuss des Appellationsgerichts des Kantons Basel-Stadt (nachfolgend: Appellationsgericht) die Berufung der Ehefrau teilweise gut und änderte die erstinstanzliche Unterhaltsregelung für die Ehefrau ab. Der vom Ehemann an den Unterhalt der Ehefrau zu bezahlende Unterhaltsbeitrag von Fr. 3'000.- wurde per 1. August 2015 auf Fr. 500.- reduziert und mit Wirkung per 1. September 2015 aufgehoben. C. Die Ehefrau (nachfolgend: Beschwerdeführerin) hat am 22. April 2016 (Postaufgabe) gegen den vorgenannten Entscheid des Appellationsgerichts beim Bundesgericht Beschwerde erhoben. Sie beantragt, den angefochtenen Entscheid aufzuheben und das Begehren des Ehemannes (nachfolgend: Beschwerdegegner) um Aufhebung des Ehegattenunterhalts für die Dauer des Scheidungsverfahrens abzuweisen. Der Beschwerdegegner hat sich am 16. Februar 2017 vernehmen lassen. Er schliesst auf Abweisung der Beschwerde. Das Bundesgericht hebt die kantonalen Entscheide auf und weist das Gesuch des Beschwerdegegners um Abänderung des Ehegattenunterhalts für die Dauer des Scheidungsverfahrens ab. (Zusammenfassung) BGE 143 III 233 S. 235 Erwägungen Aus den Erwägungen: 3. 3.1 Abänderungsgrund bildet die veränderte Einkommenslage beim Beschwerdegegner. Die erste Instanz hat dessen seit Januar 2015 bestehende Arbeitslosigkeit berücksichtigt und dabei angenommen, der Beschwerdegegner habe seine Arbeitsstelle nicht freiwillig, sondern einzig deshalb gekündigt, weil ihm die Auflösung des Arbeitsverhältnisses von seiner Arbeitgeberin nahegelegt worden sei. Das Appellationsgericht hat geprüft, ob die Kündigung als eigenmächtiges, mithin rechtsmissbräuchliches Verhalten zu werten sei, das einer Abänderung der Eheschutzmassnahmen entgegensteht. Als Ergebnis seiner Würdigung der tatsächlichen Gegebenheiten ist es zum Schluss gelangt, die Beschwerdeführerin habe nicht glaubhaft gemacht, dass die Kündigung des Arbeitsverhältnisses durch den Beschwerdegegner ausschliesslich zu ihrer Schädigung erfolgt sei. Das Appellationsgericht hat daher - wie die erste Instanz - dem Beschwerdegegner kein hypothetisches Einkommen angerechnet. 3.2 Bei der Bemessung des Unterhaltsbeitrages ist grundsätzlich vom tatsächlich erzielten Einkommen des Unterhaltspflichtigen auszugehen. Soweit dieses Einkommen allerdings nicht ausreicht, um den ausgewiesenen Bedarf zu decken, kann ein hypothetisches Einkommen angerechnet werden, sofern dieses zu erreichen zumutbar und möglich ist (BGE 137 III 118 E. 2.3). [...]

Translated: Facts from page 234 BGE 143 III 233 p. 234 A. By petition dated 22 March 2013, B. (husband) applied to the Civil Court of Basel-Stadt for a divorce from his marriage to A. (wife). By decision of 1 May 2013, he was ordered to pay the wife a monthly maintenance contribution of CHF 1,520 until 31 August 2013 and CHF 3,000 from September 2013 for the duration of the dissolution of the joint household. B. B.a On 2 August 2015, the husband requested that the maintenance contribution for the wife be reduced to CHF 1,500 per month and CHF 500 per month respectively. The wife opposed this request. In a decision dated 10 November 2015, the instructing judge of the Civil Court of Basel-Stadt cancelled the monthly maintenance contribution for the wife for the duration of the divorce proceedings with retroactive effect from 1 July 2015. B.b On 1 March 2016, the Committee of the Court of Appeal of the Canton of Basel-Stadt (hereinafter: Court of Appeal) partially upheld the wife's appeal and modified the first-instance maintenance arrangement for the wife. The maintenance contribution of CHF 3,000 to be paid by the husband to the wife was reduced to CHF 500 with effect from 1 August 2015 and cancelled with effect from 1 September 2015. C. On 22 April 2016 (date of posting), the wife (hereinafter: complainant) lodged an appeal with the Federal Supreme Court against the aforementioned decision of the Court of Appeal. She requested that the contested decision be set aside and that the husband's (hereinafter: respondent) request for cancellation of spousal maintenance for the duration of the divorce proceedings be dismissed. The respondent was heard on 16 February 2017. He concludes that the appeal should be dismissed. The Federal Supreme Court overturns the cantonal decisions and rejects the respondent's request for a modification of spousal maintenance for the duration of the divorce proceedings. (Summary) BGE 143 III 233 p. 235 Considerations From the considerations: 3. 3.1 The reason for the modification is the change in the respondent's income situation. The first instance took into account his unemployment since January 2015 and assumed that the respondent had not resigned from his job voluntarily, but only because he had been advised by his employer to terminate the employment relationship. The Court of Appeal examined whether the termination was to be regarded as unauthorised, and therefore abusive, behaviour that precluded a modification of the matrimonial protection measures. As a result of its assessment of the factual circumstances, it concluded that the complainant had not credibly demonstrated that the respondent had terminated the employment relationship solely to her detriment. The Court of Appeal therefore - like the first instance - did not impute any hypothetical income to the respondent. 3.2 When calculating the maintenance contribution, the actual income of the maintenance debtor must be taken as a basis. However, if this income is not sufficient to cover the stated needs, a hypothetical income can be taken into account, provided that it is reasonable and possible to achieve this (BGE 137 III 118 E. 2.3). [...]

Headnote (Ground Truth)

Original: Regeste Abänderung des Unterhaltsbeitrages für die Dauer des Scheidungsverfahrens infolge veränderter Verhältnisse; Rechtsmissbrauch (Art. 179 Abs. 1, Art. 2 Abs. 2 ZGB). Vermindert der Unterhaltspflichtige sein Einkommen in Schädigungsabsicht, ist eine Abänderung der Unterhaltsleistung selbst dann ausgeschlossen, wenn die Einkommensreduktion nicht mehr rückgängig gemacht werden kann (Änderung der Rechtsprechung; E. 3).

Translated: Headnote Modification of the maintenance contribution for the duration of the divorce proceedings due to changed circumstances; abuse of rights (Art. 179 para. 1, Art. 2 para. 2 ZGB). If the maintenance debtor reduces his or her income with the intention of causing damage, a modification of the maintenance payment is excluded even if the reduction in income can no longer be reversed (change in case law; E. 3).

Metadata

Decision ID: 3e70603c-bbcd-47de-a066-23124945fcc1, Year: 2017, Language: German, Court: CH_BGE, Source: BGE 143 III 233

Figure 4: A reference sample from the dataset used for a qualitative comparison of the generated headnotes from different LLMs.

GPT-4 (one-shot)

Headnote (Output)

Original: Regeste Abänderung von Unterhaltsbeiträgen bei Reduktion des Einkommens (BGE 137 III 118). Bei der Bemessung des Unterhaltsbeitrages ist grundsätzlich vom tatsächlich erzielten Einkommen des Unterhaltspflichtigen auszugehen. Eine Abänderung der Unterhaltsleistung ist jedoch auszuschliessen, wenn der Unterhaltspflichtige sein Einkommen in Schädigungsabsicht vermindert hat, selbst dann wenn die Einkommensverminderung nicht mehr rückgängig gemacht werden kann (E. 3.4).

Translated: Rulings Modification of maintenance contributions in the event of a reduction in income (BGE 137 III 118). When calculating the maintenance contribution, the actual income of the maintenance debtor must be taken as a basis. However, a reduction in the maintenance payment must be ruled out if the maintenance debtor has reduced his or her income with the intention to pay maintenance, even if the reduction in income can no longer be attributed to him or her (E. 3.4).

Evaluation

BERT: 79.82, **BLEU:** 68.93, **METEOR:** 55.47, **ROUGE-1/2/L:** 45.16 / 28.04 / 43.01

Figure 5: Generated output for the sample shown in Figure 4 with GPT-4 in a one-shot setting. On the left is the original German output, on the right the translation. We provide BERT, BLEU, METEOR and ROUGE scores for the generated headnote.

GPT-4 (one-shot) Among the four generated outputs, this model was able to express the key legal consideration very precisely right from the beginning. Unlike the other outputs, it made an attempt to identify the involved party, specifically the "maintenance debtor." While GPT-4 did not correctly identify this party, the effort to do so is considered a positive practice in summarization, as it enhances clarity. Despite the error in identifying the party, GPT-4 still successfully grasped and explained the main legal principle.

Claude 2 (one-shot) The model referred to the wrong article of the Swiss Civil Procedure Code (SCPC). Instead of referencing article 127, which is relevant to the current dispute, it cited article 137, likely confusing the article number with the cited judgment (BGE 137 III 118 E.2.3). The output lacks precision and only loosely addresses the legal principle. Overall, the summary is more vague and less of a clear, concise deduction.

GPT-3.5-Turbo (zero-shot) The model produced a summary based on an incorrect premise and misinterpreted the key considerations. For example, it incorrectly suggested the husband had displayed abusive behavior, although this was never recorded. The Court may have discussed what could constitute abusive behavior in general, and the model seems to have mistakenly linked those conditions to the husband's actions. When addressing the modification application, the model failed to reach a correct conclusion and instead focused on when such modifications cannot be reversed.

Claude Instant (zero-shot) The model mistakenly referenced article 107 of the Swiss Civil Procedure Code (SCPC), which is not applicable. However, it correctly cited a highly relevant SFSC judgment fitting the context. Despite this, it failed to clearly state the case's outcome, offering only a vague explanation of the legal principles involved.

Summary When comparing the summaries from the four models, GPT-4 (one-shot) produced the most accurate summary of the actual outcome and demonstrated a strong understanding of the legal reasoning behind it.

7 Conclusions and Future Work

In this article, we introduce SLDS, a cross-lingual dataset for summarizing judicial decisions to advance cross-lingual summarization and support legal practitioners in Switzerland. Our experiments demonstrate that fine-tuned mT5 models can compete with proprietary models like GPT-4 in summarizing judicial decisions, achieving similar or better performance while offering significant computational efficiency. Our results highlight the effectiveness of domain-specific fine-tuning for multilingual legal tasks.

Future research may address two key areas. First, finetuning larger models or exploring advanced summarization techniques, such as multi-step summarization. Second, expanding the error analysis with additional input from more legal professionals or collecting quantitative preference ratings at scale could help evaluate models in more detail by uncovering deeper issues.

Limitations

We tried as carefully as possible to design the experiment in a way that the evaluation is meaningful, which is why we included BERTScore to measure semantic similarity and multiple other metrics that measure lexicographic similarity. We also performed a qualitative error analysis to further improve the evaluation. Due to resource constraints, we were not able to involve additional legal experts into the evaluation process. We leave this up to future work, along with the incorporation of a LLM-based evaluation such as LLM-as-a-Judge (Zheng et al., 2023).

Although the mT5 model family might be considered somewhat outdated given today's wide range of LLMs, we use this model to show that fine-tuning is still competitive compared to prompting larger proprietary models. Subsequent work should investigate fine-tuning larger LMs that are currently considered state-of-the-art.

References

Dennis Aumiller, Ashish Chouhan, and Michael Gertz. 2022. EUR-Lex-Sum: A Multi- and Cross-lingual Dataset for Long-form Summarization in the Legal Domain. *arXiv preprint*. ArXiv:2210.13448 [cs].

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

Emmanuel Bauer, Dominik Stammbach, Nianlong Gu, and Elliott Ash. 2023. Legal extractive summarization of u.s. court opinions. *Preprint*, arXiv:2305.08428.

Paheli Bhattacharya, Soham Poddar, Koustav Rudra, Kripabandhu Ghosh, and Saptarshi Ghosh. 2021. Incorporating domain knowledge for extractive summarization of legal case documents. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 22–31.

Peter Bieri. 2015. Law clerks in switzerland-a solution to cope with the caseload? In *IJCA*, volume 7, page 29. HeinOnline.

Leo Breiman. 2001. Random forests. *Machine learning*, 45:5–32.

Debtanu Datta, Shubham Soni, Rajdeep Mukherjee, and Saptarshi Ghosh. 2023. MILDSum: A novel benchmark dataset for multilingual summarization of Indian legal case judgments. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5291–5302, Singapore. Association for Computational Linguistics.

Diego de Vargas Feijó and Viviane Pereira Moreira. 2018. Rulingbr: A summarization dataset for legal texts. In *Computational Processing of the Portuguese Language: 13th International Conference, PROPOR 2018, Canela, Brazil, September 24–26, 2018, Proceedings 13*, pages 255–264. Springer.

Jade Goldstein, Vibhu O Mittal, Jaime G Carbonell, and Mark Kantrowitz. 2000. Multi-document summarization by sentence extraction. In *NAACL-ANLP 2000 workshop: automatic summarization*.

Claire Grover, Ben Hachey, and Ian Hughson. 2004. The HOLJ Corpus. Supporting Summarisation of Legal Texts. In *Proceedings of the 5th International Workshop on Linguistically Interpreted Corpora*, pages 47–54, Geneva, Switzerland. COLING.

Ben Hachey and Claire Grover. 2006. Extractive summarisation of legal texts. *Artificial Intelligence and Law*, 14(4):305–345.

Deepali Jain, Malaya Dutta Borah, and Anupam Biswas. 2021. Summarization of legal documents: Where are we now and the way forward. *Computer Science Review*, 40:100388.

Deepali Jain, Malaya Dutta Borah, and Anupam Biswas. 2024. Summarization of lengthy legal documents via abstractive dataset building: An extract-then-assign approach. *Expert Systems with Applications*, 237:121571.

Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Mi-Young Kim, Ying Xu, and Randy Goebel. 2013. Summarization of legal texts with high cohesion and automatic compression rate. In *New Frontiers in Artificial Intelligence*, pages 190–204, Berlin, Heidelberg. Springer Berlin Heidelberg.

Anastassia Kornilova and Vladimir Eidelman. 2019. Billsum: A corpus for automatic summarization of us legislation. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.

Chin-Yew Lin. 2004a. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

Chin-Yew Lin. 2004b. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Revital Ludewig and Juan Lallave. 2013. Professional stress, discrimination and coping strategies: Similarities and differences between female and male judges in switzerland.

Hans Peter Luhn. 1957. A statistical approach to mechanized encoding and searching of literary information. *IBM Journal of research and development*, 1(4):309–317.

Gianluca Moro, Nicola Piscaglia, Luca Ragazzi, and Paolo Italiani. 2023. Multi-language transfer learning for low-resource legal case summarization. *Artificial Intelligence and Law*, pages 1–29.

Joel Niklaus and Daniele Giofré. 2022. BudgetLongformer: Can we Cheaply Pretrain a SotA Legal Language Model From Scratch? *arXiv preprint*. ArXiv:2211.17135 [cs].

Joel Niklaus, Lucia Zheng, Arya D. McCarthy, Christopher Hahn, Brian M. Rosen, Peter Henderson, Daniel E. Ho, Garrett Honke, Percy Liang, and Christopher Manning. 2024. FLawN-T5: An Empirical Examination of Effective Instruction-Tuning Data Mixtures for Legal Reasoning. *arXiv preprint*. ArXiv:2404.02127 [cs].

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 on Association for Computational Linguistics*, ACL '02, page 311–318, USA. Association for Computational Linguistics.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.

Natalie Schluter. 2017. The limits of automatic summarisation according to ROUGE. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 41–45, Valencia, Spain. Association for Computational Linguistics.

Zejiang Shen, Kyle Lo, Lauren Yu, Nathan Dahlberg, Margo Schlanger, and Doug Downey. 2022. MultiLexSum: Real-World Summaries of Civil Rights Lawsuits at Multiple Granularities. *arXiv preprint*. ArXiv:2206.10883 [cs].

Abhay Shukla, Paheli Bhattacharya, Soham Poddar, Rajdeep Mukherjee, Kripabandhu Ghosh, Pawan Goyal, and Saptarshi Ghosh. 2022. Legal case document summarization: Extractive and abstractive methods and their evaluation. 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 1048–1064.

Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.

L Xue. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. *arXiv:1912.08777 [cs]*. ArXiv: 1912.08777.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. BERTScore: Evaluating Text Generation with BERT. *arXiv:1904.09675 [cs]*. ArXiv: 1904.09675.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

A Potential Risks

We believe the release of SLDS poses minimal risk. On the contrary, we expect our dataset to foster further research and encourage the development of assistive technologies that can make the work of lawyers, judges, and clerks more efficient. However, it is crucial not to rely on these summaries blindly. We recommend using such systems as tools to enhance efficiency, rather than as substitutes for human oversight. Users must ensure that the generated summaries accurately reflect the decisions and do not introduce any misleading content, since lawyers will rely on these summaries to find relevant cases faster.

B Use of AI Assistants

We used ChatGPT to improve the content of this article. It was used to rephrase certain passages, as well as condense them to make the text less redundant and easier to understand. We carefully checked that the generated paraphrases corresponded to our own ideas and that no errors were introduced during this process.

C Number of Leading Decisions by Year

In Figure 6, we provide a distribution of LDs over the years.

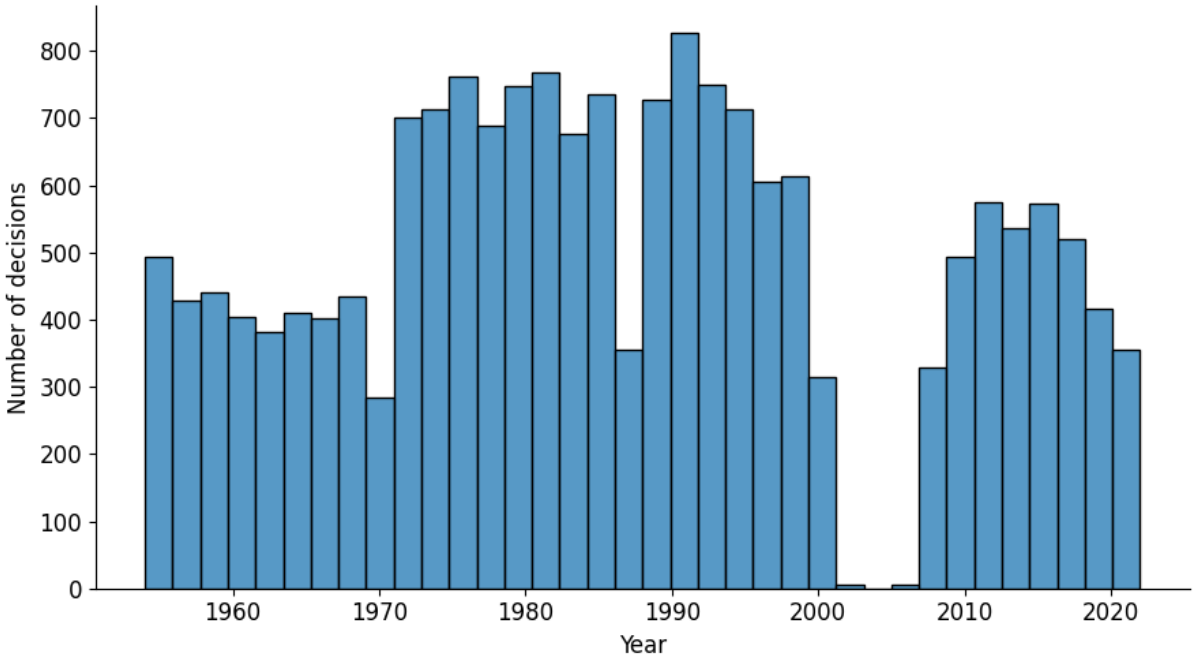

Figure 6: Number of leading decisions published per year.

D Resources Used

For the experiments, we used two NVIDIA A100 80GB GPUs. We used 10 GPU days for our experiments.

E Hyperparameters

We used gradient accumulation in our fine-tuning experiments to achieve an effective batch size of 16. The mT5 models were trained for 10 epochs and evaluations were performed every 1000 steps.

F Prompts

We used Prompt 2 for the few-shot experiments. To avoid language barriers, we included Prompt 1, the English translation of the original German prompt, in the main body of this article.

```
Ziel: Generiere eine Regeste basierend auf einem Schweizer Gerichtsurteil.
Hintergrund: Ein Schweizer Gerichtsurteil setzt sich aus Sachverhalt, Erwägungen und Dispositiv zusammen. Die Regeste dient
↪  als Kurzzusammenfassung und beinhaltet Leitsätze des Urteils. Nur Leitentscheide haben eine Regeste.
Anweisung:
1. Sachverhalt: Lies und verstehe den gegebenen Sachverhalt.
2. Erwägungen: Analysiere die Erwägungen, um die Hauptargumente und Gründe zu identifizieren.
3. Dispositiv: Beachte das Dispositiv, da es das endgültige Urteil enthält.
4. Erstelle die Regeste: Die Regeste sollte aus drei sehr kurzen Teilen bestehen: a. Zitiere die wichtigsten relevanten
↪  Artikelziffern (ohne den Artikeltitel). b. Nenne kurze, relevante, deskriptive Keywords, über die Thematik des Falls. c.
↪  Formuliere einen sehr kurzen Fliesstext, der die wichtigsten Erwägungen zitiert und kurz zusammenfasst.
Output: Die Regeste sollte eine klare und strukturierte Kurzzusammenfassung des Urteils bieten, die aus zitierten Artikeln,
↪  Keywords und einem sehr kurzen Fliesstext besteht.

Hier ein Beispiel für ein Gerichtsurteil mit Regeste:

[Gegebener Sachverhalt, Erwägungen und Dispositiv]:
{ Specific decision in German, French or Italian }

[Regeste auf Deutsch]: { Headnote for that specific decision }

[Gegebener Sachverhalt, Erwägungen und Dispositiv]:
{ Decision from the validation set }

[Regeste auf Deutsch]:
```

Prompt 2: The original prompt that was used for the zero-shot and one-shot experiments. The blue part is only added in case of a one-shot setting.

G Generated Summaries

Figure 7 shows three different samples, along with generated summaries, ground truths, and corresponding evaluations. The headnotes were generated with our fine-tuned mT5_{Base} model, and they are ordered according to high, average, and low evaluation scores.

H Few-shot results per language

We provide the detailed results for the cross-lingual evaluations in our experiment in Table 3. The results are grouped by few-shot setting, either zero-shot or one-shot, and language.

Model	Setting	BERT ↑	BLEU ↑	MET ↑	R1 ↑	R2 ↑	RL ↑
GPT-4	one-shot	**73.89** / **72.84** / **74.08**	49.32 / 45.85 / 38.08	**36.50** / 32.01 / 28.75	**31.14** / 30.03 / 32.43	9.91 / **9.13** / 10.82	**29.25** / 27.99 / 30.84
GPT-3.5	one-shot	73.26 / 72.21 / 72.50	47.44 / 41.50 / 39.08	35.23 / 28.48 / 27.37	30.64 / 28.22 / 26.11	9.46 / 8.87 / 8.96	28.90 / 26.46 / 24.34
Claude 2	one-shot	73.17 / 72.34 / 73.69	48.62 / 45.50 / 47.68	35.07 / **30.69** / 33.91	30.13 / **30.30** / **32.57**	9.29 / 8.52 / 11.50	28.39 / **28.71** / **30.86**
Claude Instant	one-shot	72.97 / 71.51 / 71.44	45.66 / 43.40 / 42.06	32.02 / 27.20 / 27.40	28.40 / 27.17 / 25.58	8.80 / 8.12 / 8.28	26.71 / 25.36 / 24.20
GPT-4	zero-shot	72.65 / 69.56 / 70.95	49.17 / **46.96** / 46.89	36.26 / 27.39 / 27.10	28.27 / 23.45 / 24.25	**10.88** / 5.36 / 7.74	26.26 / 21.42 / 22.51
GPT-3.5	zero-shot	71.25 / 68.39 / 69.84	47.39 / 43.68 / 44.17	33.31 / 25.51 / 27.91	27.11 / 21.73 / 21.86	9.22 / 4.64 / 4.75	25.47 / 20.22 / 20.48
Claude 2	zero-shot	71.74 / 69.81 / 72.77	**50.21** / 46.94 / **52.35**	34.72 / 28.09 / **35.60**	28.93 / 25.02 / 31.04	9.27 / 6.38 / **11.55**	27.11 / 23.22 / 29.45
Claude Instant	zero-shot	71.99 / 70.08 / 70.81	46.82 / 43.59 / 43.33	31.41 / 25.06 / 27.55	27.43 / 23.60 / 25.63	9.17 / 6.11 / 8.48	25.48 / 21.57 / 23.63

Table 3: Results of the experiments grouped by language. Scores are presented in the order: German, French, and Italian. Underlined: best within setting; **bold**: best overall.

I Qualitative Analysis

The reference sample that was used in our qualitative analysis can be seen in Figure 4. The generated summaries along with their evaluation scores are shown in Figures 5 to 10.

Example: Summary with high evaluation scores

Target

regeste: Regeste Teilklage und negative Feststellungswiderklage (Art. 86 und 224 ZPO). Die Ausnahme vom Erfordernis der gleichen Verfahrensart gemäss Art. 224 Abs. 1 ZPO für negative Feststellungswiderklagen ist nicht auf den Fall beschränkt, dass es sich bei der Hauptklage um eine sogenannte echte Teilklage handelt, sondern gilt allgemein dann, wenn die Teilklage eine Ungewissheit zur Folge hat, die es rechtfertigt, die Feststellung des Nichtbestands einer Forderung oder eines Rechtsverhältnisses zu verlangen (E. 2).

Predicted

regeste: Regeste Voraussetzungen, unter denen eine Widerklage erhoben werden darf, wenn eine echte Teilklage beurteilt wird (E. 2.1). Zur Anwendung von Art. 224 Abs. 1 ZPO ist die Ausnahme vom Erfordernis der gleichen Verfahrensart gemäss Art. 224 Abs. 1 ZPO nicht auf diesen Fall beschränkt, sondern gilt allgemein dann, wenn die Teilklage eine Ungewissheit zur Folge hat, die es rechtfertigt, im Sinne von Art. 88 ZPO die Feststellung des Nichtbestands einer Forderung oder eines Rechtsverhältnisses zu verlangen (E. 2.2).

Metrics

BERT: 88.03, **BLEU:** 30.45, **METEOR:** 43.78, **ROUGE-1/2/L:** 72.44 / 54.05 / 70.87

Example: Summary with average evaluation scores

Target

regeste: Regeste Uferschutzplanung Wohlensee; Gesamtinteressenabwägung. Die Behörden nehmen bei der Genehmigung einer Sondernutzungsplanung eine umfassende Interessenabwägung vor (vgl. Art. 3 RPV ; E. 3.2). Würdigung des Gutachtens der Vogelwarte Sempach und der Stellungnahme des BAFU als Fachbehörde (E. 5.5, 6.3 und 6.4). Im zu beurteilenden Fall wird das öffentliche Interesse an einer ufernahen Wegführung (vgl. Art. 3 Abs. 2 lit. c RPG) relativiert, weil ein durchgehender Sichtschutz zum Schutz der im Gebiet vorkommenden Vogelarten erforderlich wäre, was Spaziergängern das Erleben der Uferlandschaft nur sehr beschränkt ermöglichen würde (E. 6.5 sowie 3.3). Das Gebiet des Wohlensees stellt eines von insgesamt 25 Wasser- und Zugvogelreservaten von nationaler Bedeutung dar (vgl. Art. 11 Abs. 2 JSG, Art. 2 Abs. 1 i.V.m. Anhang 1 WZVV).

Predicted

regeste: Regeste Art. 3 Abs. 2 lit. c RPG ; schweizerische Vogelwarte Sempach; Umweltschutz. Gutachten im Bereich des Umweltrechts, insbesondere von dem UVG/USG ergangenen Streitwert (E. 3.1). Ein öffentlicher Weg, der über Privatgrundstücke führt und mit dem Enteignungsrecht zugunsten des Gemeinwesens verbunden ist, stellt eine öffentlich-rechtliche Eigentumsbeschränkung dar (E. 3.3). Interessenabwägungen sind grundsätzlich frei (E. 3.4 und 3.5). Umfang des geplanten Uferweges; ein ausweisenes öffentliches Interesse an einem direkt dem Ufer entlangführenden Weg sowie an einer zumindest ufernahen Wegführung (E. 3.5).

Metrics

BERT: 73.91, **BLEU:** 47.21, **METEOR:** 30.91, **ROUGE-1/2/L:** 32.75 / 10.63 / 32.75

Example: Summary with low evaluation scores

Target

regeste: Regeste a Schwere Bedrängnis, grosse seelische Belastung und heftige Gemütsbewegung; Art. 48 lit. a Ziff. 2 und Art. 48 lit. c StGB. Abgrenzung dieser Strafmilderungsgründe untereinander (E. 2.1-2.3). Regeste b Mord, heftige Gemütsbewegung und grosse seelische Belastung; Art. 48 lit. c und Art. 112 StGB. Die für den Mord typische Skrupellosigkeit schliesst die Bejahung der Strafmilderungsgründe der heftigen Gemütsbewegung und der grossen seelischen Belastung aus (E. 2.5).

Predicted

regeste: Regeste Art. 112 und 49 Abs. 2 StGB ; tiefe Gewalt, Schwere Gewalt, Mord. Die Beeinträchtigung oder der schwere Gewalt ist einen psychischen Klumpen im Sinne des Art. 48 StGB, wenn beide Personen einen Beziehungskonflikt bilden und sich ein bestimmtes Verhältnis zwischen einseitigen Umständen des Tatvergehens und der Begründung der Tatbegehren haben (E. 2 und 3).

Metrics

BERT: 69.3, **BLEU:** 12.14, **METEOR:** 14.72, **ROUGE-1/2/L:** 30.43 / 10.26 / 30.43

Figure 7: Three examples of generated summaries by $mT5_{Base}$ with input length 4096, showcasing high, average, and low scored summaries

Claude 2 (one-shot)

Headnote (Output)

Original: Regeste Art. 137 ZGB (Unterhaltsabänderung). Abänderung des Ehegattenunterhalts bei veränderter Einkommenslage. Bei freiwilliger Schädigung ist eine Abänderung selbst dann ausgeschlossen, wenn die Einkommensverminderung nicht mehr rückgängig gemacht werden kann (E. 3).

Translated: Regulations Art. 137 CC (modification of maintenance). Modification of spousal maintenance in the event of a change in income. In the case of voluntary termination, a modification is excluded even if the reduction in income can no longer be reversed (E. 3).

Evaluation

BERT: 83.88, **BLEU:** 47.53, **METEOR:** 42.44, **ROUGE-1/2/L:** 53.33 / 29.27 / 53.33

Figure 8: Generated output for the sample shown in Figure 4 with Claude 2 in a one-shot setting. On the left is the original German output, on the right the translation. We provide BERT, BLEU, METEOR and ROUGE scores for the generated headnote.

GPT-3.5-Turbo-16K (zero-shot)

Headnote (Output)

Original: BGE 143 III 233 S. 235 - Ehegattenunterhalt: Abänderung aufgrund veränderter Einkommenslage des Unterhaltspflichtigen; Anrechnung eines hypothetischen Einkommens; Rechtsmissbräuchliches Verhalten des Unterhaltspflichtigen. Der unterhaltspflichtige Ehemann beantragte eine Abänderung des Ehegattenunterhalts aufgrund seiner veränderten Einkommenslage. Das Gericht stellte fest, dass eine Abänderung ausgeschlossen ist, wenn die veränderte Einkommenslage durch eigenmächtiges, widerrechtliches oder rechtsmissbräuchliches Verhalten des Unterhaltspflichtigen herbeigeführt wurde. Es wurde entschieden, dass in diesem Fall kein hypothetisches Einkommen angerechnet werden kann, selbst wenn die Einkommensverminderung nicht mehr rückgängig gemacht werden kann.

Translated: BGE 143 III 233 p. 235 - Spousal maintenance: Modification due to a change in the debtor's income situation; imputation of a hypothetical income; abuse of law by the debtor. The husband, who was obliged to pay maintenance, applied for a modification of the spousal maintenance on the basis of his changed income situation. The court found that a modification is excluded if the changed income situation was brought about by unlawful, wrongful or abusive behaviour on the part of the maintenance debtor. It was decided that no hypothetical income can be taken into account in this case, even if the reduction in income can no longer be attributed.

Evaluation

BERT: 72.21, **BLEU:** 43.71, **METEOR:** 44.39, **ROUGE-1/2/L:** 27.52 / 12.60 / 23.85

Figure 9: Generated output for the sample shown in Figure 4 with GPT-3.5-Turbo-16K in a zero-shot setting. On the left is the original German output, on the right the translation. We provide BERT, BLEU, METEOR and ROUGE scores for the generated headnote.

Claude Instant (zero-shot)

Headnote (Output)

Original: Regeste: a) Art. 107ff ZGB, 128 III 4 E. 4 b) Ehegattenunterhalt, Abänderungsverfahren, hypothetisches Einkommen c) Das Bundesgericht hielt fest, dass dem Unterhaltspflichtigen, der sein Einkommen in Schädigungsabsicht vermindert hat, selbst dann ein hypothetisches Einkommen anzurechnen ist, wenn die Einkommensverminderung nicht mehr rückgängig gemacht werden kann. Somit ist eine Abänderung der Unterhaltsleistung in einem solchen Fall ausgeschlossen.

Translated: Regeste: a) Art. 107ff ZGB, 128 III 4 E. 4 b) Spousal maintenance, modification procedure, hypothetical income c) The Federal Supreme Court held that a hypothetical income must be credited to the maintenance debtor who has reduced his or her income with the intention to pay maintenance, even if the reduction in income can no longer be attributed to him or her. A modification of the maintenance payment is therefore excluded in such a case.

Evaluation

BERT: 76.91, **BLEU:** 67.20, **METEOR:** 49.20, **ROUGE-1/2/L:** 41.67 / 24.30 / 41.67

Figure 10: Generated output for the sample shown in Figure 4 with Claude Instant in a zero-shot setting. On the left is the original German output, on the right the translation. We provide BERT, BLEU, METEOR and ROUGE scores for the generated headnote.