

# Structural Text Segmentation of Legal Documents

Dennis Aumiller\*

Institute of Computer Science, Heidelberg University  
Heidelberg, Germany

aumiller@informatik.uni-heidelberg.de

Sebastian Lackner

Institute of Computer Science, Heidelberg University  
Heidelberg, Germany

lackner@informatik.uni-heidelberg.de

Satya Almasian\*

Institute of Computer Science, Heidelberg University  
Heidelberg, Germany

almasian@informatik.uni-heidelberg.de

Michael Gertz

Institute of Computer Science, Heidelberg University  
Heidelberg, Germany

gertz@informatik.uni-heidelberg.de

## ABSTRACT

The growing complexity of legal cases has led to an increasing interest in legal information retrieval systems that can effectively satisfy user-specific information needs. However, such downstream systems typically require documents to be properly formatted and segmented, which is often done with relatively simple pre-processing steps, disregarding topical coherence of segments. Systems generally rely on representations of individual sentences or paragraphs, which may lack crucial context, or document-level representations, which are too long for meaningful search results. To address this issue, we propose a segmentation system that can predict topical coherence of sequential text segments spanning several paragraphs, effectively segmenting a document and providing a more balanced representation for downstream applications. We build our model on top of popular transformer networks and formulate structural text segmentation as topical change detection, by performing a series of independent classifications that allow for efficient fine-tuning on task-specific data. We crawl a novel dataset consisting of roughly 74,000 online Terms-of-Service documents, including hierarchical topic annotations, which we use for training. Results show that our proposed system significantly outperforms baselines, and adapts well to structural peculiarities of legal documents. We release both data and trained models to the research community for future work.<sup>1</sup>

## CCS CONCEPTS

• **Applied computing** → Law; **Annotation**; • **Information systems** → **Document structure**.

## KEYWORDS

Document Understanding, Outline Generation, Text Segmentation

### ACM Reference Format:

Dennis Aumiller, Satya Almasian, Sebastian Lackner, and Michael Gertz. 2021. Structural Text Segmentation of Legal Documents. In *Eighteenth International Conference for Artificial Intelligence and Law (ICAIL'21)*,

\*Both authors contributed equally to this research.

<sup>1</sup><https://github.com/dennlinger/TopicalChange>

ICAIL'21, June 21–25, 2021, São Paulo, Brazil


© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Eighteenth International Conference for Artificial Intelligence and Law (ICAIL'21)*, June 21–25, 2021, São Paulo, Brazil, <https://doi.org/10.1145/3462757.3466085>.

June 21–25, 2021, São Paulo, Brazil. ACM, New York, NY, USA, 10 pages.  
<https://doi.org/10.1145/3462757.3466085>

## 1 INTRODUCTION

Written texts are often a sequence of semantically coherent segments, designed to create a smooth transition between various subtopics discussed in a single document. Usually, the information needs of a user are satisfied by retrieving only the relevant subtopic, and retrieving the whole document is unwieldy and may result in information overload [29, 43]. However, the context of a single subtopic frequently spans multiple sentences and contains localized context, which is crucial for proper understanding. Despite the clear relevance of segmentation to downstream performance, many (legal) retrieval systems choose structurally rigid representations of only a single text element (generally either the full document [8, 28], or a single paragraph/sentence [34, 42]), disregarding the semantic coherence. Especially in legal documents, which can be extremely lengthy and contain domain-specific complexities in their topics, it is important to suitably represent entire topics in a single cohesive unit. Furthermore, aside from semi-structured legal texts, such as laws, other documents do not necessarily contain uniform and easily separable segments, to begin with. Especially input formats such as PDF or scans frequently lack any sort of meta descriptors for hierarchical information about the document contents, which makes this a challenging task. To find a fitting representation that captures the precise topical context in the text, a robust and flexible framework to obtain such a structural segmentation is required. We therefore propose a new approach for the estimation of topic boundaries to generate more suitable document representations for the mentioned downstream applications, by considering the topical coherence of paragraphs. We define a coherent section in a document as a unit consisting of potentially one or multiple paragraphs, which together share a common topic. Section boundaries often coincide with a change in topic and can thus be assumed to generate candidates for the later segmentation.

Despite their importance, many previous works for structural text segmentation ignore the notion of paragraphs and focus only on the granularity of sentences. This is contrary to the nature of written text, where paragraphs represent a semantically cohesive unit, which is already available and represents a coarser and more meaningful structure than sentences. In this work, we assume that topic boundaries generally do not appear in the middle of a paragraph, and, consequently, operating on paragraph level can reduce the risk of false-positive segmentations and lower the computation cost of per



**Figure 1: Visual cues such as paragraphs often give away a notion of semantic coherence, which is disregarded in sentence-level models.**

sentence prediction. Figure 1 shows how paragraphs group sentences and divide a text into coherent parts and how by overlooking this valuable information the structure in the text is lost to the model.

Focusing on the field of text segmentation in the context of the Natural Language Processing (NLP), we find an already large body of existing work. Because no large labeled dataset existed, early approaches to text segmentation were mainly unsupervised, using heuristics to identify whether two sentences belong to the same topic or not. Such approaches either exploit the fact that topically related words tend to appear in semantically coherent segments [6, 18, 22, 26, 39], or focus on the representation of text in terms of latent-topic vectors using topic modeling methods, such as Latent Dirichlet Allocation [2, 29, 30, 37]. Recently, with the availability of annotated data, text segmentation has also been formulated as a supervised learning problem. Most methods utilize a hierarchical neural model, where the lower-level network creates sentence representations, and a secondary network models the dependencies between the embedded sentences [14, 21]. These models use sentence dependencies to predict the potential segment boundaries. One drawback of these approaches is their sentence-level granularity, which disregards the paragraph coherence previously mentioned. This problem is partially solved by hierarchical neural models, where the dependency between sentences is modeled in a hierarchical structure, to combine sentence into bigger units. However, training such models are, due to the different document lengths, computationally expensive. Moreover, these models fail to take advantage of larger pre-trained language representations, such as BERT [9] and RoBERTa [23], which have recently proven to be valuable feature generators with a low cost of fine-tuning, advancing the state-of-the-art in several disciplines.

In this paper, we aim to tackle the task of structural text segmentation using transformer-based models, and introduce a novel dataset of Terms-of-Service documents, containing annotated paragraphs belonging to the same topic. We focus on the text segmentation alone and leave the use of segmentation for retrieval enhancement to future work. We formulate topical coherence as a special case of binary classification of Same Topic Prediction (STP) and fine-tune our transformer-based models to detect paragraphs belonging to sections with similar topics. We consider sections as the top-level hierarchy in our dataset and do not consider subsections. We assume topical independence between consecutive paragraphs and show that it does not lead to a deterioration of performance while avoiding the costly computation of hierarchical models. The hypothesis is that by fine-tuning the paragraph embedding for topic similarity we can generate segment features that detect coherent topical structures in

a document. We evaluate our models against the traditional embedding baselines and compare them to supervised and unsupervised approaches for text segmentation and find significant improvements by our method.

**Contributions.** The contributions of this paper are as follows: (i) We present the task of structural text segmentation on coarser cohesive text units (paragraphs/sections). (ii) We investigate the performance of transformer-based models for topical change detection, and (iii) frame the task as a collection of independent binary predictions, reducing overhead for hierarchical training and simplified training sample generation. (iv) We present a new dataset consisting of online Terms-of-Service documents partitioned into hierarchical sections, and make the data available for future research. (v) We evaluate our model against classical baselines for text segmentation, and (vi) show the effectiveness of our generated embeddings for structural segmentations to obtain superior performance to other text segmentation techniques.

## 2 RELATED WORK

Our work is closely tied to the broader fields of legal document understanding, topic analysis, text segmentation, and transformer language models, and we briefly review related work in each of these areas in this section.

### 2.1 Legal Document Understanding

The area of legal document understanding and legal information retrieval has a long-standing history. A great overview is presented by Moens [31], who details more of the applications that were mentioned in the introduction. While there is existing work on the topic of legal document segmentation [24, 25, 27], they are generally concerned with a specific information extraction interest. In the case of Mencía, they are concerned with metadata of French law documents [27]. Specifically, they also make use of an existing HTML/XML structure in their input documents, but do not generalize to arbitrary text inputs without structural features. Lu et al., on the other hand, are utilizing clustering techniques to identify sub-topics in legal documents. However, these topics are irrespective of actual section boundaries within the document [24]. They further include additional metadata, such as citations, headnotes, and key numbers, in their task setup, which are suitable to their specific application in case law. Similarly, Conrad et al. [8] have previously attempted to employ clustering for heterogeneous document collections, again focusing on hierarchical representations in the form of short topic descriptors and not focusing on the actual textual content of the documents themselves. Lyte and Branting [25] classify metadata labels based on a CRF (Conditional Random Field) model, building on prior work by Branting [3] in the same direction, but focus mostly on element classification. Slightly longer segments in the form of entire sentences are both used by Poudyal et al. [34], who mine arguments from European case-law decision, and Westermann et al. [42], where a system for efficient similarity search based on sentence embeddings is presented.

## 2.2 Topic Analysis

Detection and analysis of topical change are grounded in topic modeling approaches. Earlier work such as LDA [2] treat documents as bag-of-words, where each document is assigned to a topic distribution, and each topic is a distribution over all words. More recent work has adopted a more sophisticated representation than bag-of-words and generally models Markovian topic or state transitions to capture dependencies between words in a document [15, 41]. With the rise of distributed word representation, the focus has shifted to the combination of LDA and word embeddings [11, 32]. Since we are interested in a primary segmentation without necessarily predicting topics, we put a stronger focus on the related work of segmentation methods, as discussed in the following section.

## 2.3 Text Segmentation

Text Segmentation is the task of dividing a document into a multi-paragraph discourse unit that is topically coherent, with the cut-off point usually indicating a change in topic [17, 39]. Although the task itself dates back to 1994 [17], most existing text segmentation datasets are small and limit their scope to sentences (predicting whether two sentences discuss the same topic or not). The most common dataset is by Choi [7], containing only 920 synthesized passages from the Brown corpus. Choi’s method (C99) is a probabilistic algorithm measuring similarity via term overlap. GraphSeg [13] is an unsupervised graph method that segments documents using a semantic relatedness graph of a document. GraphSeg is also evaluated on a small set of 5 manually-segmented political manifestos from the Manifesto project<sup>2</sup>. Another class of methods are topic-based document segmentations, which are statistical models that find latent topic assignments reflecting the underlying structure of the document [1, 4, 5, 10, 29, 38]. TopicTiling [38] performs best among this family of methods and uses LDA to detect topic shifts, with computing similarities between adjacent blocks based on their term frequency vectors. Brants et al. [4] follow a similar approach but employ PLSA [19] to compute the estimated word distributions. Another noteworthy approach based on Bayesian topic models is by Chen et al. [5], where they constrain latent topic assignments to reflect the underlying organization of document topics. They also publish a test dataset with 218 Wikipedia articles about cities and chemical elements.

All mentioned methods are unsupervised learning approaches, and small annotated datasets are only used for evaluation and, hence, are not directly comparable to our approach. Instead, we focus on supervised learning of topics and introduce a new dataset with 43,056 automatically labeled documents.

The only two comparable supervised approaches are from Koshorek et al. [21] and Glavas et al. [14]. Koshorek et al. [21] propose a hierarchical LSTM architecture for learning sentence representation and their dependencies. They train their hierarchical model on a dataset of cleaned Wikipedia articles, called Wiki-727k. Glavas et al. [14] introduce Coherence-Aware Text Segmentation, which encodes a sentence sequence using two hierarchically connected transformer networks. The two latter models are closest to our work in terms of data size and problem formulation. However, they rely solely on per sentence predictions, which is incomparable to our paragraph-based

method. The model by Glavas et al. is similar to our approach in that it is based on a transformer architecture, yet, they do not take advantage of transfer learning from pre-trained language models and learn all the features from scratch. Finally, Zhang et al. [45] extend text segmentation by outline generation and trained an end-to-end LSTM-model for identifying sections and generating corresponding headings for Wikipedia documents.

## 2.4 Transformer Language Models

The transformer architecture, much like recurrent neural networks, aims to solve sequence-to-sequence tasks, relying entirely on self-attention to compute representations of its input and output [40]. Transformers have made a significant step in bringing transfer learning to the NLP community, which allows the easy adaptation of a generically pre-trained model for specific tasks. Pre-trained models such as BERT, GPT-2, and RoBERTa [9, 23, 35] use language modeling for pre-training on large corpora. These models are powerful feature generators, which with minimal task-specific fine-tuning achieve state-of-the-art performance on a wide variety of tasks. Although at the core of all these models lies the idea of transformers and attention mechanisms, many have been modified and optimized to fit various downstream applications. One variation based on BERT is Sentence-BERT [36], which combines two BERT-based models in Siamese fashion to derive semantically meaningful sentence embeddings. By its design, Sentence-BERT also allows for longer input sequences for pairwise training tasks and outperforms BERT on semantic textual similarity tasks, making it a suitable choice for embedding paragraphs. Another notable variant of BERT is RoBERTa, a retraining of BERT with improved training methodology and more training data, it achieves slightly better results than BERT on some natural understanding tasks. Due to the advantages of RoBERTa, we chose RoBERTa and Sentence-RoBERTa from the Sentence-BERT variant for the setup in our approach.

## 3 SAME TOPIC PREDICTION

We formulate structural text segmentation as a supervised learning task of the same topic prediction. Our model consists of two steps: (i) Independent and Identically Distributed Same Topic Prediction (IID STP) and (ii) Sequential inference over a full document. As mentioned previously, sections are the considered level of hierarchy in our model and the structure of sub-sections is ignored in this study. However, the model is easily adaptable to any granularity, and our dataset contains information for all the levels. In the first step, we fine-tune transformer-based models to detect topical change for both paragraphs and entire sections. Given two paragraphs or sections, the classifier should correctly identify if they discuss the same subject or not. We assume that the topic of each paragraph or section is independent of the text before and after, meaning that the topic of one paragraph does not affect the likelihood of the next paragraph belonging to the same topic. We later prove that this assumption yields good performance without a costly training of hierarchical models. In the second step, we use the fine-tuned transformer-based classifiers for sequential inference on entire documents, where the segment boundaries are defined by topical change. In the following, we discuss these steps in more detail.

<sup>2</sup><https://manifestoproject.wzb.eu>


### 3.1 IID Same Topic Prediction (STP)

A document  $d \in D$  is represented as a sequence of  $N$  sections  $S_d = (s_1, \dots, s_N)$ , where each section is assigned one of  $M$  topics  $T = (t_1, \dots, t_M)$ , and each section contains up to  $K$  paragraphs  $P_n = (p_1, \dots, p_K)$ . We assume topical consistency within a paragraph and argue that the results for classification do not change based on the position of the paragraph in the document, since the most relevant part for our inference is the intra-section information. Therefore, all paragraphs in a section belong to the same topic. If the topic assignment is defined by the function  $Topic$ , we have:

$$\begin{aligned} s_n = (p_1, \dots, p_k) \wedge Topic(s_n) = t_1 &\implies \\ Topic(p_1) = t_1 = \dots = Topic(p_k) = t_1 &\quad (1) \end{aligned}$$

If we define  $C$  as a chunk of text corresponding to either a section or paragraph, the topic prediction task is defined for section and paragraph granularity as follows: Given two chunks of text of the same type (both paragraphs or both sections)  $(c_1, c_2)$  and labels  $y \in \{0, 1\}$ , indicating whether the two chunks belong to the same topic, topical change detection can be formulated as a binary classification problem. The positive class indicates that both chunks have the same topic, whereas the negative class indicates a change in topic and potentially the beginning of a new segment in text. Note that we only consider chunks of the same type, namely, either *only* sections or *only* paragraphs, in each model. By formulating the problem as a binary classification, detecting the topic consistency between two chunks of text can now be solved with any type of classifier. In this work, we train two types of transformer-based classifiers for this task, one from the pre-trained language models [23] and another Siamese network [36] variation, which is more suitable for encoding pairwise similarity. Subsequently, the two variations are discussed.

**3.1.1 RoBERTa** is a replication study of BERT pre-training with optimized hyper-parameters that applies minor adjustments to the BERT language model to achieve better performance [23]. BERT and RoBERTa both belong to the family of pre-trained transformer-based language models. The transformer is an architecture for shaping one sequence into another one with the help of the self-attention mechanism, which helps the model to extract features from each word relative to all the other words in the sequence. The encoder stacks in BERT and RoBERTa consist of one or multiple self-attention blocks followed by a feed-forward network. During pre-training, two sentences are taken as input, and models are trained on two tasks of language modeling, by predicting masked words in the input and next sentence prediction, and by classifying whether the two sentences are sequential. By these means, the models learn task-independent features from a vast amount of unlabelled text that can then be used in a fine-tuning stage for various natural language understanding tasks. Since the performance difference between most transformer-based language models is negligible, we choose RoBERTa as the representative of this family. In the fine-tuning process, the model receives two chunks as input and learns to predict whether they belong to the same topic or not. To distinguish between two chunks in training a [CLS] token is inserted at the beginning of the first chunk and a [SEP] token is inserted at the end of both the first and second chunk. The embedding of the [CLS] is what is used for pre-training the next sentence prediction task and contains RoBERTa’s understanding at the sentence-level.



**Figure 2: Demonstration of how the transformer classifier is used during inference, by comparing consecutive paragraphs to detect section boundaries.**

This token is used by a simple classification layer, learned during fine-tuning, for the same topic prediction task. Since the input size for both chunks combined is limited to a maximum of 512 tokens, shorter than many sections and paragraphs in our dataset, any longer chunk of text has to be truncated to fit.

**3.1.2 Sentence-Transformers (SRoBERTa)** aims to enhance the sentence embeddings by modification of RoBERTa using a Siamese architecture to derive semantically meaningful sentence embeddings [36]. Their method is available for several transformer models. We choose a RoBERTa-based variant to make the results comparable to the first approach. SRoBERTa enables RoBERTa to be used for certain new tasks, such as large-scale semantic similarity comparison. Their modifications result in faster inference and better representation for sentence-pair tasks. Moreover, because of the Siamese structure and coupling of two RoBERTa networks, the input size doubles, which allows for longer sequences and thus more context. In this setup, each sentence is passed through a separate RoBERTa network with an input limit of 512 tokens. The sentence embeddings are derived from a pooling operation over the output of two models with tied weights. Sentence-Transformers introduce several learning objectives, out of which we use the classification objective function with binary cross-entropy loss to classify the chunks into the same topics.

### 3.2 Sequential Inference

For inference, we use the classifiers of the previous step as topic change detectors for text segmentation. We read each paragraph of the document sequentially and classify the adjacent paragraphs for topical mutuality. More concretely, given a document  $d \in D$  divided into consecutive paragraphs  $P = (p_1, \dots, p_k)$ , section breaks are marked as where the paragraph’s topic changes. Considering a transformer  $TF$  as our classifier and two consecutive paragraphs as our input, the classifier outputs the probability of the two paragraphs belonging to the same topic, independent of their surrounding context, e.g.,  $TF(p_1, p_2) = P(Topic(p_1) = Topic(p_2))$ . Therefore, given sequences of paragraphs  $p_1, \dots, p_k$ , and the corresponding predicted labels  $y = (y_1, \dots, y_{k-1})$ , a segmentation of the document is given by  $k - 1$  predictions of  $TF$ , where  $y_i = 0$  denotes the end of a segment by  $p_i$ . It is worth noting that regardless of the chunk type used during the training of the classifiers (section or paragraph inputs) the segmentation module operates on paragraphs only. Figure 2 shows

the inference on a sample document with four paragraphs and two sections, where the paragraph colors show the topics. The *TF* classifier is applied on a paragraph pair and can ideally recognize the topic change from *P3* to *P4* and mark the beginning of the new section.

### 3.3 Legal Applications

To put the presented segmentation into a legal context, we focus on three main application areas: (i) As mentioned, a section-based semantic segmentation can be used as a pre-processing step for a passage retrieval context. This, however, would require additional data with relevance annotations for both sentence- and paragraph-level relevance to compare the specific benefits of our approach, which we leave to future work in this area. (ii) However, semantically coherent sections can also be used as a basis for similarity search. This is especially helpful when looking for, e.g., related sections in existing contracts [42]. Here, we focus on Terms-of-Service documents that are widely available, and contain sections that follow a general pattern of similar topics. (iii) Lastly, the section separation can be used for generating outlines of documents, which has previously been shown to work well on other domains such as Wikipedia [45]. During our document crawl, we also encountered several documents not including any sectional headings, which makes it especially hard to understand the legal contexts for laymen users.

## 4 TERMS-OF-SERVICE (TOS) DATASET

Due to data governance policies in many countries, it is generally mandated that commercial websites contain the necessary legal information for site users. Specifically, these must be easily reachable via the landing page, which makes it comparatively easy to be crawled. For each Terms-of-Service document, we automatically extract the content divided into paragraphs and respective hierarchical section headings. Further, ToS documents allow us to experiment with a large-scale dataset that comes with a shared set of topics, while still maintaining a heterogeneous set of topics due to the different types of websites. In the following, we will discuss the detailed mining process, and implicate limitations of this approach.

### 4.1 Crawling

As seeds to our crawler, we use the Alexa 1M URL dataset.<sup>3</sup> For each URL in the dataset, we try to access the website both with and without the `www` prefix. First, the landing page is downloaded and parsed using the *Beautiful Soup* Python package. We then search for hyperlinks with texts *Terms of Service*, *Terms of Use*, *Terms and Conditions*, and *Conditions of Use*, and follow them to get to the respective terms-of-service pages. Levenshtein distance with a threshold of 0.75 is used to allow for spelling mistakes and different wording. The raw Hypertext Markup Language (HTML) content of the Terms-of-Service page is downloaded and stored for further processing. In case of an error, e.g., if the website is temporarily unreachable, we retry the same website 2 additional times before skipping it. The unprocessed dataset contains HTML code for roughly 74,000 websites. Note that due to limitations of the current crawler implementation, websites that rely on JavaScript to display content are not supported.

<sup>3</sup>Available at: <http://s3.amazonaws.com/alexa-static/top-1m.csv.zip>

## 4.2 Section Extraction

Despite the fact that HTML is a structured format, it is a non-trivial task to extract text and hierarchies. The main reasons are that Web pages often contain a lot of boilerplate (e.g., navigational elements, advertisements, etc.), generally have heterogeneous appearances and implementations, and that they simply do not always conform to the HTML standard. Here, only a rough outline of the pipeline is given. For further reference, please refer to the implementation in our repository.

*Boilerplate Removal.* For boilerplate removal, we use the *boilerpipe* package by Kohlschütter et al. [20], which is based on shallow text features for classifying the text elements on a Web page. The result is an HTML page with all navigational elements, advertisements, and template code removed. Importantly, relevant hierarchical information is retained past this step.

*HTML Cleanup.* To deal with websites that do not conform to HTML standards, we perform several cleanup steps. This includes, for example, fixing mistakes such as text appearing without a corresponding paragraph (`<p>` tag), or incorrectly nested tags (e.g., section headings within a `<p>` tag). We fix such mistakes by adding missing tags and adjusting nested tags similar to how a web-browser would interpret the code.

*Language Detection.* Since the Alexa dataset also contains many non-English websites, we reject extracted terms-of-service, where the majority of text most likely has a language different from English. We use the *langid* Python package for detecting the language of each individual paragraph (`<p>` tag).

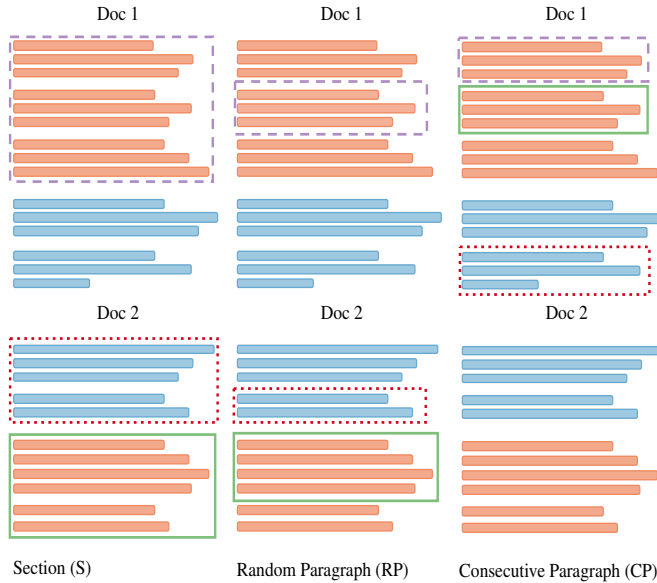
*Extracting Hierarchy.* To obtain the hierarchy, we split the document into smaller chunks. Splits are done in the following order: first we split on each section heading (`<h1>-<h6>` tags), then on bold text (`<b>` tag) starting with an enumeration pattern, then on enumerations (`<li>` tags), then on underline text (`<u>` tag) starting with an enumeration pattern, and lastly on regular text (`<p>` tag) starting with an enumeration pattern. To prevent spurious splits, each criterion is only used if there are at least 5 occurrences within the document. Each time the document is split, we save the corresponding headings, which then form the hierarchy. As enumeration patterns we recognize Latin numbers, roman numerals, and letters, optionally prefixed with *Part*, *Section*, or *Article*. The majority of documents contain at most two levels of section hierarchy.

## 4.3 Data Set Statistics

In addition to the full dataset, we provide a cleaned subset for which we manually grouped sections into distinct topics based on similar spelling or meaning. We manually merged 554 section titles, which corresponds to all titles with at least 250 distinct occurrences in the corpus. After merging, 82 topics were obtained, and only sections that have at least one of these aliases as a heading were kept. The dataset contains different levels of section hierarchy. For our work, we group document content into top-level sections only, any further hierarchies are discarded, but are present in the raw data and available for future work. After removing documents without any valid sections, belonging to the predefined 82 topics, we are left with approximately 43,000 documents for the same section prediction

**Table 1: Top 10 section topics by document frequency. Additionally, the number of associated paragraphs is given.**

Topic Label	Document Frequency	Paragraph Frequency
limitation of liability	21,317	68,517
indemnification	16,698	25,683
law and jurisdiction	15,113	29,790
links to other websites	13,752	24,727
termination	12,855	33,978
warranty	9,926	41,403
privacy	8,958	25,022
disclaimer	8,575	29,265
general terms	7,936	54,693

**Figure 3: Visualization of three distinct setups for same topic classification, where the example chunk of text is depicted in dashed purple, a positive sample in line of green, and a negative sample in dotted red. Each vertical line is a sentence and the grouping of them shows a paragraph, different sections in the document are shown with different colors, denoting the same topic across all paragraphs of the same section. From left to right: Same Section prediction, Random Paragraph, and Consecutive Paragraph sampling.**

task, and around 40,000 documents for our paragraph-level setup. We randomly split the data 80/10/10 into train, validation, and test set. The average number of sections per document is 6.56, and each document consists of 22.32 paragraphs on average, which results in a mean of 2.92 paragraphs per section. Table 1 shows the top 10 section topic labels. The average number of paragraphs per section varies between different topics.

## 5 EVALUATION

We demonstrate the capabilities of transformer-based architectures for topical change detection using dataset consisting of online Terms-of-Service (ToS) documents, which was discussed above. Results are compared for the introduced IID STP task (see Section 3.1) as well as a downstream comparison of text segmentation results to a range of baselines and existing methods. Results show a great improvement in the performance for all transformer-based models.

### 5.1 Evaluation of Models

We compare our methods against a range of baselines, including averaging over Global Vectors (*GLVavg*) [33], tf-idf vectors (*tf-idf*), and Bag of Words (*BoW*) [16]. For transformer language models, we evaluate the standard [CLS] sequence classification with *roberta-base* (*Ro-CLS*). For Sentence-Transformers [36] we use the Siamese transformer setup with a variant of *roberta-base* (*ST-Ro*) and an additional model that has been pre-trained on NLI sentence similarity tasks (*ST-Ro-NLI*) to investigate performance of further pretraining.

Transformer models are trained using the popular HuggingFace transformer library [44] for the [CLS] models, as well as the sentence-transformers package from [36] for Siamese variants. We use two Nvidia Titan RTX GPUs for training, and each model variant has been trained with five different random seeds. Details for the training parameters can be found in our public repository. Due to the length limitation of 512 tokens, we employ an iterative truncation strategy for two-sentence inputs. Due to the coupling of two transformers for the Sentence-Transformers the input size doubles, accepting a total input of 1024 tokens.

### 5.2 Prediction Tasks

As previously introduced, we train models with an independent classification setup, which is generally much faster than more complicated hierarchical sequential models. Specifically, we highlight the differences in the setup for the same section prediction task, compared to the two paragraph-based methods. We point out that results depicted in Table 2, for the prediction accuracy are not directly comparable between sampling methods, as they generate different development and test sets based on the employed sampling strategy. We show in the subsequent section, however, that downstream performance for the text segmentation is in line with results on topic prediction. Figure 3 visualizes the different sampling strategies. In the following, we describe each strategy in detail and highlight their difference. Across all strategies, we added three positive and three negative samples for each individual section/paragraph.

**5.2.1 Section (S) Topic Prediction.** In this setup, we use sections as input chunks to the transformer classifiers. The section task showcases how different levels of granularity can affect outcomes in the prediction results. Specifically, the extremely long input sequences test the limits of what transformers can predict from partial observations since the majority of inputs will be heavily truncated. To ensure an equal distribution of samples from within the same and different sections, we match each section with three samples from the same topic, and three from different topics. The positive and negative sections can be sampled from a different document.

**Table 2: Prediction accuracy for the independent topic prediction tasks, Same Topic Prediction (STP), Random Paragraph (RP), Consecutive Paragraph (CP) with different sampling strategies. Standard deviation is reported over 5 runs and the best model on each respective set is depicted in bold.**

		GLVavg	tf-idf	BoW	Ro-CLS	ST-Ro	ST-Ro-N
STP	Dev	89.70 $\pm$ .07	82.10 $\pm$ .05	50.94 $\pm$ .33	96.42 $\pm$ .52	96.38 $\pm$ .03	<b>96.39 <math>\pm</math>.03</b>
	Test	90.01 $\pm$ .06	82.54 $\pm$ .07	51.05 $\pm$ .51	96.58 $\pm$ .52	96.45 $\pm$ .06	<b>96.46 <math>\pm</math>.02</b>
RP	Dev	76.63 $\pm$ .04	70.94 $\pm$ .07	50.34 $\pm$ .04	57.63 $\pm$ 10.4	<b>87.50 <math>\pm</math>.13</b>	87.39 $\pm$ .08
	Test	76.16 $\pm$ .06	70.41 $\pm$ .09	50.31 $\pm$ .37	57.48 $\pm$ 10.2	<b>87.19 <math>\pm</math>.64</b>	86.88 $\pm$ .11
CP	Dev	77.64 $\pm$ 6.6	74.94 $\pm$ .11	56.34 $\pm$ .83	89.63 $\pm$ .12	<b>91.17 <math>\pm</math>.05</b>	91.12 $\pm$ .04
	Test	78.63 $\pm$ 6.8	76.17 $\pm$ .07	56.58 $\pm$ 1.1	90.34 $\pm$ .08	91.17 $\pm$ .04	<b>91.69 <math>\pm</math>.02</b>

The important point is that the positive samples should come from the *same topic* and negative samples from different ones. The first column of Figure 3 visualizes the section sampling, where the first section of *Doc*<sub>1</sub> is paired with the second section of *Doc*<sub>2</sub> to form a positive sample and the first section of *Doc*<sub>2</sub> to form a negative sample, respectively. The same strategy is employed for the generation of the development and test set.

Despite the constraints with respect to the input length, we find that all transformers perform on a near-perfect level, compare Table 2. Comparing these results to already very well-performing baselines, we suspect that certain keywords give away similar sections, but highlight the fact that the explicit representation of different topics is not given during training in the binary classification task, which makes this a suitable method for dealing with imbalanced topics.

**5.2.2 Random Paragraph (RP) Topic Prediction.** In contrast to the section-level task, we revert to a more fine-grained distinction of paragraphs in a text. In the Random Paragraph setting, we still generate samples similarly, meaning we include three paragraphs from a random document with the same topic and three negative samples from random paragraphs with different topics. The main difference between the Section prediction and Random Paragraph is in the level of granularity and not how the samples are chosen. The second column of Figure 3 highlights this difference, where the samples are paragraphs inside the sections rather than the entire section. Paragraph-based sampling is closer to our inference setup, where each input document is considered one paragraph at a time. However, results show a sharp drop in the performance, which can come from a much narrower context of the paragraphs, as well as a differing selection of test samples compared to the section task. Solely the BoW model seems to be largely unaffected, which is simply due to its low performance in either setting.

**5.2.3 Consecutive Paragraph (CP) Topic Prediction.** To boost performance and account for the coherent structures in the text, we employ a sampling strategy inspired by Ein Dor et al. [12]. For their triplet loss, samples are generated inside the same document only, which can be translated into sampling from intra-document paragraphs. Note that this strategy also no longer requires any merging and annotation of topics across documents, as all relevant information is now contained within a single document. This fact opens up much larger generation of training data, which we omit in our current work for the sake of comparability with the RP

model. To generate samples, we look at all paragraphs of a section and pair them as positive samples. Negative samples are picked from paragraphs of different sections in the same document. The third column of Figure 3 depicts the consecutive paragraph setup, where the samples are limited to paragraphs of *Doc*<sub>1</sub>. Note that despite their similar setup, results of RP and CP runs in Table 2 are not evaluated on the same test set and thus are not comparable, since the test sets are each generated with the respective sampling strategies (RP or CP) as well. However, we are able to compare their downstream performance on the subsequently introduced text segmentation task (see Section 5.3 and Table 3).


The result of different sampling strategies along with the performance of the baselines is shown in Table 2, where the transformer-based models all outperform the baselines by a significant margin. Among the baselines BoW has the worst performance overall, with the accuracy close to random, showing that distinct word occurrences are not a sufficient indicator. Average GloVe has the best performance of all baselines, but is still behind the transformers by a large margin. Despite the NLI-pretrained SRoBERTa model (ST-Ro-N) achieving better scores than the base model (ST-Ro) for most setups, the difference is insignificant, indicating that the pre-training on sentence similarity tasks does not directly influence our topic prediction setup.

### 5.3 Text Segmentation

By generating a text segmentation over the paragraphs of a full document, the independent prediction results from the previous section can now be compared across several approaches. Specifically, we compare the paragraph-based training methods CP and RP. As an evaluation metric, we follow related literature and adopt the  $P_k$  metric introduced by Beeferman et al. [1], which is the error rate of two segments at  $k$  sentences apart being classified incorrectly. We use the default window size of half the document length for our evaluation, again following related work. Furthermore, we count the number of explicit misclassifications, and use the accuracy  $acc_k$  of “up to  $k$  mistakes per document” as an evaluation metric. Due to the coarser nature of paragraphs and the lower number of predictions per document compared to the sentence-level segmentation, this is a more illustrative metric. This also relates to the “exact match” metric  $EM_{outline}$  employed by Zhang et al. [45],

**Table 3: Boundary error rate  $P_k$  for compared models (lower is better), based on sampling strategies Random Paragraph (RP), Consecutive Paragraph (CP) and their Ensemble variates,  $RP_{Ens}$  and  $CP_{Ens}$ , respectively. Ensemble (“Ens”) predictions are obtained by majority voting over model runs.**

	RP	CP	$RP_{Ens}$	$CP_{Ens}$
GLVavg	29.97 ±.09	26.23 ±6.2	29.55	23.06
tf-idf	39.87 ±.24	29.70 ±.28	39.36	28.60
BoW	45.76 ±.67	43.46 ±1.5	46.20	41.80
Random Oracle	35.08 ±.15	-	31.88	-
GraphSeg	-	32.48 ±.46	-	32.28
WikiSeg	-	48.29 ±.30	-	48.29
Ro-CLS	37.26 ±4.8	15.15 ±.00	41.15	15.15
ST-Ro	15.72 ±.11	14.06 ±.14	<b>14.62</b>	13.14
ST-Ro-N	15.97 ±.14	13.97 ±.19	14.81	<b>12.95</b>
Ens consec	-	-	-	12.50



**Figure 4: Mistake rate of per-model ensembles, where the suffix CP indicates the consecutive paragraph sampling and RP the random paragraph sampling for each model. The baseline is Rand Oracle (Random Oracle), GLVavg (average GloVe vectors), tf-idf, and BoW (Bag-of-Words). Ro-CLS is the fine-tuned CLS token for Roberta and Sentence transformer models are ST-Ro and ST-Ro-N, where the latter is pre-trained on NLI task. The best performing model is the Ens All (Ensemble of all models).**

where  $acc_0 = EM_{outline}$ .

Here, we also include the performance of related works where public and up-to-date code repositories are available. Specifically, we compare to the unsupervised segmentation algorithm GraphSeg [13],

and the supervised model by Koshorek et al. [21], which we dub “WikiSeg”. Both approaches are trained on a sentence-level approach, though, and predictions have to be translated back to a paragraph level for comparison of results. We train each model with the suggested parameters in their publicly available repositories. For an additional pseudo-sequential baseline, we use an informed random oracle that has a-priori information on the number of topics in the document, and samples from a distribution with adjusted probability  $P(\text{“next section”}) = \#sections/\#paragraphs$ . Note that no additional parameters are learned for any model, and predictions are binarized with a simple 0.5 threshold over the same topic predictions. We provide ensembling results for the majority voting decisions by the five seed runs of each model variant (Ens), which provides further improvements. Best results are obtained by ensembling all consecutive transformer-based methods (Ens consec).

Table 3 shows the results of the evaluation, where one can see that results in the sequential segmentation are directly linked to the performance on the independent classification task seen in Table 2. To verify our initial assumption of cross-document comparability of content from similar sections, we make the following observations: (i) Evaluation performance for the STP setup is consistent for both training strategies (RP and CP) when using Sentence-Transformer models (see Table 2). (ii) Similarly, both CP and RP-trained Sentence-Transformer segmentations achieve results within 2 percentage points of the respective  $P_k$  scores. (iii) In general, CP training setup yields slightly better  $P_k$  scores, likely because the intra-document dependencies are captured better with this sampling strategy, which is a more appropriate sampling for the segmentation task. (iv) We find convergence problems for RP training with the [CLS] models, as well as the tf-idf model. Due to the size of our general training corpus, we therefore conclude that it is realistic to expect topical similarity within a section, even across documents. However, due to the seemingly inconsistent convergence of RP models, we caution against blindly using this strategy, especially when dealing with more heterogeneous corpora. The oracle baseline performs unexpectedly better than both tf-idf and BoW, indicating that additional information about the sections of a document can greatly boost task performance, which might be relevant for future work. Additional pre-training of ST models (ST-Ro-N) does not show any significant improvement over the standard ST-Ro models.

To our surprise, sentence-based implementations (GraphSeg and WikiSeg) show significantly lower performance, and fall even behind the simpler baselines. For GraphSeg, an unsupervised segmentation approach, the lack of explicit training on the different granularity seems to significantly prohibit correct predictions on longer segments. WikiSeg heavily preprocesses the data and discards many samples, thus significantly shrinking the training set. Since performance on the reduced training set is still decent, this indicates that training a network from scratch is not suitable with the smaller training set of a reduced corpus and tends to overfit. We expect a significant increase in performance if the training would instead be performed without such strict preprocessing criteria, or continuing fine-tuning on pre-trained weights from a paragraph-level WikiSeg model. For either baseline model, it is also important to note that these models predict on the entirety of the



sequence, which theoretically allows information sharing between different sections in the current sample. However, they show no improvement over our binary prediction setup which does not share this information. It would be of interest to compare results to sequential transformer-based architectures, such as they are used by Glavas et al. [14]. However, their model again requires training from scratch, which has proven to be inconsistent in our experiments with WikiSeg.

Lastly, the plots for  $acc_k$  for various models in Figure 4 indicate a correlation between the  $acc_k$  and  $P_k$  measures, which does not apply to sentence-level segmentations. Overall, the best-performing ensembles classify around 25% of documents without any mistake ( $acc_0$ ), and around 70% with less than three mistakes ( $acc_2$ ) over the entire document. We therefore suggest  $acc_k$  as an interpretable addition to the classic evaluation of segmentation approaches when dealing with paragraph-level segmentations.

## 6 CONCLUSION AND FUTURE WORK

Despite a multitude of previous works, structural text segmentation methods have always focused on very finely segmented text chunks in the form of sentences. In this work, we have shown that a relaxation of this problem to coarser text structures reduces the complexity of the problem, while still allowing for semantic segmentation. Further, we reformulate the oftentimes expensive-to-train sequential setup of text segmentation as a supervised Same Topic Prediction task, which reduces training time, while allowing for a near-trivial generation of samples from automatically crawled text documents. To show the applicability of our method, we present a new domain-specific and large corpus of online Terms-of-Service documents, and train transformer-based models that vastly outperform a number of text segmentation baselines.

We are currently investigating the setup for deeper hierarchical sections, which our dataset already contains annotations for, to see whether such notions can also be picked up by an independent classifier and benefit a legal retrieval system. Also, the findings from our Consecutive Paragraph model already indicate that training requires no further information than the ground truth segmentation, which can generally be inferred from structured input formats, such as HTML or XML, making this an attractive option for a larger-scale study of cross-domain document collections. Finally, an interface build on top of our framework, enabling the users to judge the usefulness of segmentation for legal use cases, such as a collection of documents from mergers and acquisitions, could be used to determine the efficacy of our improved segmentation.

## ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their insightful comments.

## REFERENCES

- [1] Doug Beeferman, Adam L. Berger, and John D. Lafferty. 1999. Statistical Models for Text Segmentation. *Mach. Learn.* 34, 1-3 (1999), 177–210.
- [2] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (2003), 993–1022.
- [3] Luther Karl Branting. 2017. Automating Judicial Document Analysis. In *Proceedings of the Second Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 16th International Conference on Artificial Intelligence and Law (ICAIL 2017), London, UK, June 16, 2017 (CEUR Workshop Proceedings, Vol. 2143)*, Kevin D. Ashley, Katie Atkinson, Luther Karl Branting, Enrico Francesconi, Matthias Grabmair, Marc Lauritsen, Vern R. Walker, and Adam Zachary Wyner (Eds.). CEUR-WS.org.
- [4] Thorsten Brants, Francine Chen, and Ioannis Tsochantaridis. 2002. Topic-based Document Segmentation with Probabilistic Latent Semantic Analysis. In *Proceedings of the 2002 ACM CIKM International Conference on Information and Knowledge Management, McLean, VA, USA*. ACM, 211–218.
- [5] Harr Chen, S. R. K. Branavan, Regina Barzilay, and David R. Karger. 2009. Global Models of Document Structure using Latent Permutations. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, 2009, Boulder, Colorado, USA*. 371–379.
- [6] Freddy Y. Y. Choi. 2000. Advances in Domain Independent Linear Text Segmentation. In *Proceedings of the 1st North American Chapter of the Association for Computational Linguistics Conference (Seattle, Washington) (NAACL 2000)*. ACL, USA, 26–33.
- [7] Freddy Y. Y. Choi. 2000. Advances in Domain Independent Linear Text Segmentation. In *6th Applied Natural Language Processing Conference, ANLP, Seattle, Washington, USA, 2000*. ACL, 26–33.
- [8] Jack G. Conrad, Khalid Al-Kofahi, Ying Zhao, and George Karypis. 2005. Effective Document Clustering for Large Heterogeneous Law Firm Collections. In *The Tenth International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 6-11, 2005, Bologna, Italy*, Giovanni Sartor (Ed.). ACM, 177–187.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, Minneapolis, USA, 2019, Volume 1*. 4171–4186.
- [10] Satya Dharanipragada, Martin Franz, J. Scott McCarley, Salim Roukos, and Todd Ward. 1999. Story Segmentation and Topic Detection for Recognized Speech. In *Sixth European Conference on Speech Communication and Technology, EUROSPEECH 1999, Budapest, Hungary*. ISCA.
- [11] Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2019. Topic Modeling in Embedding Spaces. *CoRR* abs/1907.04907 (2019). arXiv:1907.04907
- [12] Liat Ein Dor, Yosi Mass, Alon Halfon, Elad Venezian, Ilya Shnayderman, Ranit Aharonov, and Noam Slonim. 2018. Learning Thematic Similarity Metric from Article Sections Using Triplet Networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. ACL, Melbourne, Australia, 49–54.
- [13] Goran Glavas, Federico Nanni, and Simone Paolo Ponzetto. 2016. Unsupervised Text Segmentation Using Semantic Relatedness Graphs. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics, \*SEM@ACL, Berlin, Germany, 2016*. The \*SEM 2016 Organizing Committee.
- [14] Goran Glavas and Swapna Somasundaran. 2020. Two-Level Transformer and Auxiliary Coherence Modeling for Improved Text Segmentation. *CoRR* abs/2001.00891 (2020). arXiv:2001.00891
- [15] Thomas L. Griffiths, Mark Steyvers, David M. Blei, and Joshua B. Tenenbaum. 2004. Integrating Topics and Syntax. In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, 2004, British Columbia, Canada]*. 537–544.
- [16] Zellig S. Harris. 1954. Distributional Structure. *WORD* 10, 2-3 (1954), 146–162.
- [17] Marti A. Hearst. 1994. Multi-Paragraph Segmentation of Expository Text. In *32nd Annual Meeting of the Association for Computational Linguistics, 1994, Las Cruces, New Mexico, USA, Proceedings*. ACL, 9–16.
- [18] Marti A. Hearst. 1997. TextTiling: Segmenting Text into Multi-Paragraph Subtopic Passages. *Comput. Linguist.* 23, 1 (March 1997), 33–64.
- [19] Thomas Hofmann. 2017. Probabilistic Latent Semantic Indexing. *SIGIR Forum* 51, 2 (2017), 211–218.
- [20] Christian Kohlschütter, Peter Fankhauser, and Wolfgang Nejdl. 2010. Boilerplate Detection Using Shallow Text Features. In *Proceedings of the Third International Conference on Web Search and Web Data Mining, WSDM, New York, USA 2010*. ACM, 441–450.
- [21] Omri Koshorek, Adir Cohen, Noam Mor, Michael Rotman, and Jonathan Berant. 2018. Text Segmentation as a Supervised Learning Task. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, 2018, Volume 2 (Short Papers)*. ACL, 469–473.
- [22] Hideki Kozima. 1993. Text Segmentation Based on Similarity between Words. In *31st Annual Meeting of the Association for Computational Linguistics, 1993, Ohio State University, Columbus, Ohio, USA, Proceedings*. ACL, 286–288.
- [23] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR* abs/1907.11692 (2019). arXiv:1907.11692
- [24] Qiang Lu, Jack G. Conrad, Khalid Al-Kofahi, and William Keenan. 2011. Legal document clustering with built-in topic segmentation. In *Proceedings of the 20th*

- ACM Conference on Information and Knowledge Management, CIKM 2011, Glasgow, United Kingdom, October 24–28, 2011, Craig Macdonald, Iadh Ounis, and Ian Ruthven (Eds.). ACM, 383–392.
- [25] Alex Lyte and Karl Branting. 2019. Document Segmentation Labeling Techniques for Court Filings. In *Proceedings of the Third Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 17th International Conference on Artificial Intelligence and Law (ICAIL 2019)*, Montreal, QC, Canada, June 21, 2019 (CEUR Workshop Proceedings, Vol. 2385), Kevin D. Ashley, Katie Atkinson, Luther Karl Branting, Enrico Francesconi, Matthias Grabmair, Bernhard Walzl, Vern R. Walker, and Adam Zachary Wyner (Eds.). CEUR-WS.org.
- [26] Igor Malioutov and Regina Barzilay. 2006. Minimum Cut Model for Spoken Lecture Segmentation. In *ACL, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, Sydney, Australia, 2006. ACL.
- [27] Eneldo Loza Mencía. 2009. Segmentation of legal documents. In *The 12th International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 8–12, 2009, Barcelona, Spain*. ACM, 88–97.
- [28] Nada Mimouni. 2013. Modeling Legal Documents as Typed Linked Data for Relational Querying. In *Proceedings of the First JURIX Doctoral Consortium and Poster Sessions in conjunction with the 26th International Conference on Legal Knowledge and Information Systems, JURIX 2013, Bologna, Italy, December 11–13, 2013 (CEUR Workshop Proceedings, Vol. 1105)*, Monica Palmirani and Giovanni Sartor (Eds.). CEUR-WS.org.
- [29] Hemant Misra, François Yvon, Olivier Cappé, and Joemon M. Jose. 2011. Text Segmentation: A Topic Modeling Perspective. *Inf. Process. Manag.* 47, 4 (2011), 528–544.
- [30] Hemant Misra, François Yvon, Joemon M. Jose, and Olivier Cappé. 2009. Text segmentation via Topic Modeling: an Analytical Study. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM, Hong Kong, China, 2009*. 1553–1556.
- [31] Marie-Francine Moens. 2001. Innovative techniques for legal text retrieval. *Artif. Intell. Law* 9, 1 (2001), 29–57.
- [32] Christopher E. Moody. 2016. Mixing Dirichlet Topic Models and Word Embeddings to Make lda2vec. *CoRR* abs/1605.02019 (2016). arXiv:1605.02019
- [33] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*. 1532–1543.
- [34] Prakash Poudyal, Teresa Gonçalves, and Paulo Quaresma. 2019. Using Clustering Techniques to Identify Arguments in Legal Documents. In *Proceedings of the Third Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 17th International Conference on Artificial Intelligence and Law (ICAIL 2019)*, Montreal, QC, Canada, June 21, 2019 (CEUR Workshop Proceedings, Vol. 2385), Kevin D. Ashley, Katie Atkinson, Luther Karl Branting, Enrico Francesconi, Matthias Grabmair, Bernhard Walzl, Vern R. Walker, and Adam Zachary Wyner (Eds.). CEUR-WS.org.
- [35] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language Models are Unsupervised Multitask Learners. (2018).
- [36] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP, Hong Kong, China*. ACL, 3980–3990.
- [37] Martin Riedl and Chris Biemann. 2012. How Text Segmentation Algorithms Gain from Topic Models. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, 2012, Montréal, Canada*. ACL, 553–557.
- [38] Martin Riedl and Chris Biemann. 2012. TopicTiling: A Text Segmentation Algorithm based on LDA. In *Proceedings of the Student Research Workshop of the 50th Meeting of the Association for Computational Linguistics*. Republic of Korea, 37–42.
- [39] Masao Utiyama and Hitoshi Isahara. 2001. A Statistical Model for Domain-Independent Text Segmentation. In *Association for Computational Linguistic, 39th Annual Meeting and 10th Conference of the European Chapter, Proceedings of the Conference, 2001, Toulouse, France*. ACL, 491–498.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, Long Beach, CA, USA*. 5998–6008.
- [41] Hanna M. Wallach. 2006. Topic modeling: Beyond Bag-of-Words. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML), Pittsburgh, Pennsylvania, USA, 2006 (ACM International Conference Proceeding Series, Vol. 148)*. 977–984.
- [42] Hannes Westermann, Jaromír Savelka, Vern R. Walker, Kevin D. Ashley, and Karim Benyekhlef. 2020. Sentence Embeddings and High-Speed Similarity Search for Fast Computer Assisted Annotation of Legal Documents. In *Legal Knowledge and Information Systems - JURIX 2020: The Thirty-third Annual Conference, Brno, Czech Republic, December 9–11, 2020 (Frontiers in Artificial Intelligence and Applications, Vol. 334)*, Villata Serena, Jakub Harasta, and Petr Kremen (Eds.). IOS Press, 164–173.
- [43] Ross Wilkinson. 1994. Effective Retrieval of Structured Documents. In *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, Dublin, Ireland, 1994 (Special Issue of the SIGIR Forum)*. ACM/Springer, 311–317.
- [44] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv* abs/1910.03771 (2019).
- [45] Ruqing Zhang, Jiafeng Guo, Yixing Fan, Yanyan Lan, and Xueqi Cheng. 2019. Outline Generation: Understanding the Inherent Content Structure of Documents. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR, Paris, France, 2019*. 745–754.