

Chapter Captor: Text Segmentation in Novels

Charuta Pethe, Allen Kim, Steven Skiena

Department of Computer Science,
Stony Brook University, NY, USA

{cpethe, allekim, skiena}@cs.stonybrook.edu

Abstract

Books are typically segmented into chapters and sections, representing coherent sub-narratives and topics. We investigate the task of predicting chapter boundaries, as a proxy for the general task of segmenting long texts. We build a Project Gutenberg chapter segmentation data set of 9,126 English novels, using a hybrid approach combining neural inference and rule matching to recognize chapter title headers in books, achieving an F1-score of 0.77 on this task. Using this annotated data as ground truth after removing structural cues, we present cut-based and neural methods for chapter segmentation, achieving an F1-score of 0.453 on the challenging task of *exact* break prediction over book-length documents. Finally, we reveal interesting historical trends in the chapter structure of novels.

1 Introduction

Text segmentation (Hearst, 1994; Beeferman et al., 1999) is a fundamental task in natural language processing, which seeks to partition texts into sequences of coherent segments or episodes. Segmentation tasks differ widely in scale, from partitioning sentences into clauses to dividing large texts into coherent parts, where each segment is ideally an independent event occurring in the narrative.

Text segmentation plays an important role in many NLP applications including summarization, information retrieval, and question answering. In the context of literary works, event detection is a central concern in discourse analysis (Joty et al., 2019). In order to obtain representations of events, it is essential to identify narrative boundaries in the text, where one event ends and another begins.

In novels and related literary works, authors often define such coherent segments by means of sections and chapters. Chapter boundaries are typically denoted by formatting conventions such as page breaks, white-space, chapter numbers, and

titles. This physical segmentation improves the readability of long texts for human readers, providing transition cues for breaks in the story.

In this paper, we investigate the task of identifying chapter boundaries in literary works, as a proxy for that of large-scale text segmentation. The text of thousands of scanned books are available in repositories such as Project Gutenberg (Gutenberg, n.d.), making the chapter boundaries of these texts an attractive source of annotations to study text segmentation. Unfortunately, the physical manifestations of the printed book have been lost in the Gutenberg texts, limiting their usefulness for such studies. Chapter titles and numbers are retained in the texts but not systematically annotated: indeed they sit as hidden obstacles for most NLP analysis of these texts.

We develop methods for extracting ground truth chapter segmentation from Gutenberg texts, and use this as training/evaluation data to build text segmentation systems to predict the natural boundaries of long narratives. Our primary contributions¹ include:

- **Project Gutenberg Chapter Segmentation Resource:** To create a ground-truth data set for chapter segmentation, we developed a hybrid approach to recognizing chapter formatting which is of independent interest. It combines a neural model with a regular expression based rule matching system. Evaluation on a (noisy) silver-standard chapter partitioning yields a mean value F1 score of 0.77 of a test set of 640 books, but manual investigation shows this evaluation receives an artificially low recall score due to incorrect header tags in the silver-standard.

Our data set consists of 9,126 English fiction books in the Project Gutenberg corpus. To

¹All code and links to data are available at <https://github.com/cpethe/chapter-captor>.

encourage further work on text segmentation for narratives, we make the annotated chapter boundaries data publicly available for future research.

- **Local Methods for Chapter Segmentation:**

By concatenating chapter text following the removal of all explicit signals of chapter boundaries (white space and header notations), we create a natural test bed to develop and evaluate algorithms for large-document text segmentation. We develop two distinct approaches for predicting the location of chapter breaks: an unsupervised weighted-cut approach minimizing cross-boundary cross-references, and a supervised neural network building on the BERT language model (Devlin et al., 2019). Both prove effective at identifying likely boundary sites, with F1 scores of 0.164 and 0.447 respectively on the test set.

- **Global Break Prediction using Optimization:**

Social conventions encourage authors to maintain chapters of modest yet roughly equal length. By incorporating length criteria into the desired optimization criteria and using dynamic programming to find the best global solution enables us to control how important it is to keep the segments equal. We find that a balance between equal segments and model-influenced segments gives us the best segmentation, with minimal error. Indeed, augmenting the BERT-based local classifier with dynamic programming yielded an F1 score of 0.453 on the challenging task of *exact* break prediction over book-length documents, while simultaneously beating challenging baselines on two other error metrics.

Incorporating chapter length criteria require an independent estimate of the number of chapters in a given text. We demonstrate that there are approximately five times as many likely break candidates as there are chapter breaks in the weighted cut approach, reflecting the number of sub-events within an average book chapter.

- **Historical Analysis of Segmentation Conventions**

– We exploit our data analysis of segmented books in two directions. We demonstrate that novels grew in length to an average of roughly 30 chapters/book by 1800, and retained this length until 1875 before beginning

a steady decline. Second, an analysis of regular expression patterns reveal the wide variety of chapter header conventions and which forms dominate.

2 Previous Work

Many approaches have been developed in recent years to address variants of the task of identifying structural elements in books.

McConnaughey et al. (2017) attempt this task at the page-level, by assigning a label (e.g. Preface, Index, Table of Contents, etc.) to each page of the book. Wu et al. (2013) address the task of recognizing and extracting tables of contents from book documents, with a focus on identifying its style. Participants of the Book Structure Extraction competition at ICDAR 2013 (Doucet et al., 2013) attempted to use various approaches for the task. These include making use of the table of contents, OCR information, whitespace, and indentation. Déjean and Meunier (2005) present approaches to identify a table of contents in a book, and Déjean and Meunier (2009) attempt to structure a document according to its table of contents.

However, our approach relies only on text, and does not require positional information or OCR coordinates to extract front matter and headings from book texts.

For text segmentation, many approaches have been developed over the past years, suitable for different types of data, such as news articles, scientific article, Wikipedia pages, and conversation transcripts.

The TextTiling algorithm (Hearst, 1994) makes use of lexical frequency distributions across blocks of a fixed number of words. Dotplotting (Reynar, 1994) is a graphical technique to locate discourse boundaries using lexical cohesion across the entire document.

Yamron et al. (1998) and Beeferman et al. (1999) propose methods to identify story boundaries in news transcripts.

The C99 algorithm (Choi, 2000) uses a global lexical similarity matrix and a ranking scheme for divisive clustering. Choi et al. (2001) further proposed the use of Latent Semantic Analysis (LSA) to compute inter-sentence similarity.

Utiyama and Isahara (2001) proposed a statistical model to find the maximum probability segmentation. The Minimum Cut model (Barzilay and Malioutov, 2006) addresses segmentation as a graph partitioning task.

This problem has also been addressed in a Bayesian setting (Eisenstein and Barzilay, 2008; Eisenstein, 2009). TopicTiling (Riedl and Biemann, 2012) is a modification of the TextTiling algorithm, and makes use of LDA for topic modeling.

Segmentation using sentence similarity has been extensively explored using affinity propagation (Kazantseva and Szpakowicz, 2011; Sakahara et al., 2014). More recent approaches (Alemi and Ginsparg, 2015; Glavaš et al., 2016) involve the use of semantic representations of words to compute sentence similarities. Koshorek et al. (2018) and Badjatiya et al. (2018) propose neural models to identify break points within the text.

Sims et al. (2019) address the slightly different, but relevant task of event prediction using a neural model, on a human-annotated dataset of short events.

3 Header Annotation

In order to create a ground-truth dataset for chapter segmentation, we first build a system to recognize chapter headings, using a hybrid approach combining a neural model with a regular expression (regex)-based rule matching system.

3.1 Data

In the absence of human-annotated gold standard data with annotated front matter and chapter headings, we derive silver-standard ground truth from Project Gutenberg. We identify 8,400 English fiction books available in HTML format, and extract (noisy) HTML header elements from these books. We use a train-test split of 90-10%.

3.2 Methodology

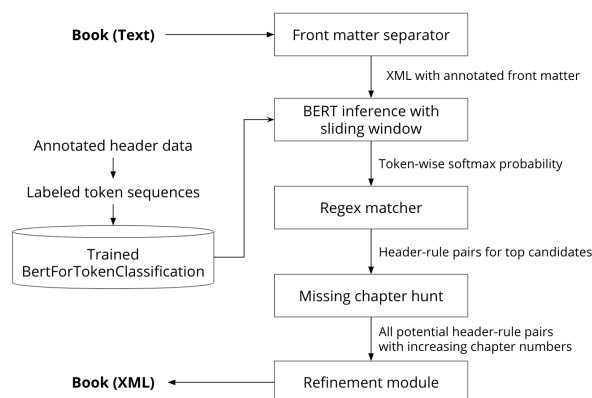


Figure 1: Header Annotation Pipeline

The annotation pipeline has five components, as shown in Figure 1. First, we make use of white-space cues and string matching for keywords such

as ‘Preface’, ‘Table of contents’ etc. to identify front matter. We tag all such content up to the first chapter heading as the front matter, and identify the remaining content as body.

3.2.1 BERT Inference

We fine-tune a pretrained BERT model (Devlin et al., 2019) with a token classification head, to identify the lines which are likely to be headers.

Training: For each header extracted from the Project Gutenberg HTML files, we append content from before and after the header, to generate training sequences of fixed length. We empirically select a sequence length of 120. We use a custom BERT Cased Tokenizer with a special token for the newline character, to tokenize the input sequences. The training samples are of the format:

Sequence: $[p_1, \dots, p_x, h_1, \dots, h_k, q_1, \dots, q_y]$

Labels: $[0, \dots, 0, 1, \dots, 1, 0, \dots, 0]$

where p_1, \dots, p_x are x tokens before the header, h_1, \dots, h_k are k tokens from the header, and q_1, \dots, q_y are y tokens after the header. x and y are randomly generated numbers, such that $x + k + y = 120$. This is done in order to prevent header tokens from appearing only in the center of the input sequence.

We fine-tune a pre-trained model for token classification using headers from 6,515 books in our training set for 4 epochs using the BertAdam optimizer. A compute server with a 2.30 GHz CPU and TeslaV100 GPU was used for all experiments.

Inference: For inference on a test set example, we tokenize the text using the custom BERT Cased Tokenizer, and use the model to generate a confidence score for each token. We do this using a sliding window approach, wherein we run inference on a text window of 120 tokens, and slide the window forward by 60 tokens in each iteration. We then perform token-wise max pooling to obtain a single confidence score per token. Further, we detokenize the output by concatenating sub-word tokens and mean-pooling their confidence scores.

We choose the top 10% tokens with the highest confidence scores, and use the lines containing these tokens as potential header candidates for regex matching.

3.2.2 Regex Rule Matching

We compile a list of regular expressions for constituent elements in chapter headings:

- Keywords like ‘Chapter’, ‘Section’, ‘Volume’

- Punctuation marks and whitespace
- Title (uppercase and mixed case)
- Roman numerals (uppercase and lowercase)
- Cardinal, ordinal, and digital numbers.

Using the rules for these constituent elements, we further generate a list of 1,015 regex rules for valid permutations of these elements.

For every potential header candidate generated using the BERT model, we pick the best matching regex rule as the longest rule that captures constituent elements in order of priority, and discard the candidate if there is no matching rule.

3.2.3 Missing Chapter Hunt

Once we have the list of candidates and their corresponding matching rules, we search for chapter headings the BERT model may have missed. For each matched rule that contains a number in some format, we search for chapter headings in the same format with the missing number. In order to account for chapter numbering restarts in different sections of the book, we search for missing headers within all increasing subsequences in the list of chapter numbers found.

3.2.4 Refinement

We get rid of false positive matches, by removing headers between consecutive chapter numbers, which do not match the same rule.

3.3 Evaluation

Table 1 shows the stage-wise performance of the annotation pipeline. Stage 1 contains all candidates generated using the BERT model, Stage 2 contains headers predicted after applying regex rules and searching for missing chapters, Stage 3 contains headers after removing false positives.

Stage	Precision	Recall	F1
1	0.02	0.67	0.05
2	0.75	0.79	0.76
3	0.78	0.78	0.77

Table 1: Stage-wise performance for header annotation

Figure 2 shows the distribution of evaluation metrics on the test set of 640 books, evaluated on the ground truth extracted from HTML files. The mean value of the F1 score is 0.77. Manual investigation of a sub-sample of the test set shows that several books get a low recall score due to false negatives, caused due to incorrect header tags in the




Figure 2: F1 score distribution for 640 test set books

silver-standard ground truth. Thus we have even greater confidence in our testbed than the F1-score suggests.

3.4 Popularly used rule formats




Figure 3: Number of books in which the most frequent header formats occur the most frequently

For each book, we count the number of occurrences of each header format. Figure 3 shows the number of books in which the respective header format occurs most frequently, namely “Chapter # TITLE”.

3.5 Historical Trends




Figure 4: Trend in the number of chapters in a book

Figure 4 presents the number of chapters in

each book as obtained by our annotation pipeline, against the author’s year of birth. For authors born before 1875, novels were roughly 30 chapters long, after which there has been a steady decline in the number of chapters per book.

4 Local Methods for Segmentation

After removing all explicit signals of chapter boundaries from the texts, we now evaluate algorithms for segmenting text into chapters.

We formulate our task as follows:

Given: Sentences S_0, S_1, \dots, S_{N-1} in the book, and P , the number of breaks to insert

Compute: P break points

$B_0, B_1, \dots, B_{P-1} \in \{0, 1, \dots, N - 1\}$ corresponding to chapter breaks.

4.1 Weighted Overlap Cut (WOC)

The motivation behind this technique is based on the intuition that chapters are relatively self-contained in the words that they use. For example, consider a chapter that refers to a “cabin in the woods”. We would expect references to this cabin to be higher within the same chapter as compared to other chapters. Hence, our hypothesis is that there will be fewer words in common across a break point separating two chapters, as compared to words within the same chapter.

Considering sentences as nodes, and common words as edges, we can compute the density of a potential break point as the sum of the number of edges going across it, weighted by their distance from the break point. As per our hypothesis, we expect the break point between two chapters to appear as local minima in density as a function of sentence number.

We restrict potential break points to the points between paragraphs, and compute the local minima in density. For each local minimum, we compute its prominence as the vertical distance between the minimum and its highest contour line. We then pick the top P most prominent local minima as the break points.

Note that the same hypothesis can also be made at the paragraph level. However, a major limitation of this approach is that paragraph sizes vary widely, ranging from a single word to a considerably huge block of text. Hence, we have taken the approach of computing sentence-level density and then restricting the potential break points to points between paragraphs.

Preprocessing: We use the Stanford CoreNLP pipeline (Manning et al., 2014) for sentence tokenization and lemmatization. We consider paragraphs as text separated by two or more newline characters.

Computation: For every potential break point i between sentences S_i and S_{i+1} , we compute the density of the break point, which is essentially a weighted sum of the number of overlapping word lemmas within a certain window before and after the break point (weighted by the distance of the word occurrence from the break point). We compute the density d_i of break point candidate i as:

$$d_i = \sum_{x=i-w}^i \left(\sum_{y=\max(x+1,i)}^{x+w} \frac{\text{overlap}_{xy}}{|i-x||i-y|} \right)$$

where w is the window size and overlap_{xy} is the number of common lemmas in sentences S_x and S_y , excluding stopwords and punctuation. (Note that we use only valid sentence indices during summation, considering the first and last sentences of the book as cutoffs.)

Experiments: We perform experiments on 2,546 books in the test set, using window sizes of 50, 100, 150, and 200 sentences.

Figure 5a shows the computed densities and local minima for window size 200, for a sample book (“*The Rover Boys Out West*”, by Edward Stratemeyer). The figure shows that chapter breaks roughly correspond to prominent local minima in density.


4.2 BERT for Break Prediction (BBP)

We fine-tune a pre-trained BERT model for the Next Sentence Prediction task, to classify pairs of sequences in which the second sequence is a coherent continuation of the first. Intuitively, for text sequences which are separated by a chapter break, we expect the second sequence to not be a continuation of the first, i.e. the output label should be 0. Whereas for consecutive text sequences within the same chapter, the output label should be 1, denoting that it is a logical continuation.


Training: We generate training sequences from 7,582 books in the training set. We generate training examples in the following format:

[CLS] <Seq A> [SEP] <Seq B> [SEP]

To generate negative training samples (i.e. class 0, meaning chapter break), we consider all the chapter boundaries, and construct the input using the



(a) WOC density (Local minima point sizes are proportional to their prominences)



(b) Processed BERT confidence scores

Figure 5: Breakpoint probability scores as a function of sentence number, for a sample book (“*The Rover Boys Out West*”, by Edward Stratemeyer). Vertical red lines denote chapter breaks in ground truth. Predictions are computed using the dynamic programming approach (described in Section 5) with $\alpha = 0.8$.

text just before the chapter break as Seq A , and text just after the break as Seq B . To generate positive training samples (i.e. class 1, meaning no break) we consider the break points between paragraphs within the same chapter, and construct the input sequence similarly. We use these sequence pairs to fine-tune a pre-trained model for next sentence prediction. Note that class 0 is of interest to us in this task, as lack of continuity between the sequences denotes the possibility of a chapter break.

Inference: During inference on a book, we consider all break points between paragraphs, and generate input sequences as described above. We run each pair of input sequences through the classifier, and generate confidence scores per class. We then use the confidence score for class 0 as the probability of a break. We select the top P break points with the highest confidence scores.

Experiments: We perform experiments using the following variants of training sequences to fine-tune the BERT model:

- *Single paragraph:* We use only one paragraph from before, and one paragraph from after the break point.
- *Full window:* We use 254 tokens each, from

before and after the break point. (If the paragraph length exceeds 254 tokens, we cut off the text before/after that point, depending on which side the paragraph lies.)

Figure 5b shows the modified BERT scores for the full window configuration, for a sample book in the test set. The figure shows that BERT is able to capture points close to chapter breaks in most cases, indicating a good recall as well as precision.

4.3 Evaluation

We evaluate our algorithms using three metrics:

P_k (Beeferman et al., 1999): To compute this metric, k is set to half of the average true segment size. Using a moving window of length k , a penalty is computed based on whether the two ends of the window are in the same or different segments, and whether the ground truth segmentation is in agreement.

WindowDiff (WD) (Pevzner and Hearst, 2002): This metric also uses a moving window, and compares the number of ground truth segmentation boundaries that fall in the window, with the number of boundaries assigned by the algorithm. A penalty is added if the counts are not equal.

F1 score : We use the F1 score to evaluate *exact* break prediction, and consider a match only if the break matches with the ground truth exactly, i.e. predictions near the true break points are not counted.

Lower values of P_k and WindowDiff, and a high value for F1 score are indicative of better performance.

Table 2 shows the evaluation metrics P_k , WD (WindowDiff), and the F1 score for the WOC and BBP configurations described above.

Algorithm	P_k	WD	F1
Equidistant breaks	0.482	0.492	0.052
TextTile (Hearst, 1994)	0.587	0.714	0.085
C99 (Choi, 2000)	0.493	0.517	0.049
P(Badjatiya et al., 2018)	0.485	0.555	0.111
L(Badjatiya et al., 2018)	0.493	0.569	0.087
WOC (window=50)	0.442	0.465	0.144
WOC (window=100)	0.425	0.450	0.158
WOC (window=150)	0.418	0.447	0.162
WOC (window=200)	0.416	0.446	0.164
BBP (single para.)	0.454	0.509	0.126
BBP (full window)	0.303	0.384	0.447

Table 2: Evaluation metrics for chapter break insertion approaches (For P_k , WD: lower is better. For F1: higher is better.)

We compare our approaches against the following baselines:

- **Equidistant:** We divide the book into $P + 1$ segments, such that each segment has the same number of sentences.
- **TextTiling:** We run the TextTiling algorithm (Hearst, 1994), using mean words per sentence as pseudosentence size, and number of paragraphs as block size for each book. The average number of breaks per book inserted by this algorithm is 574, which clearly does not reflect the actual number of chapters, resulting in poor performance.
- **C99:** We run the C99 algorithm (Choi, 2000) on our dataset, and choose the first P breaks obtained while performing divisive clustering.
- **Perceptron (P):** We train a 3-layer baseline perceptron model with 300 neurons in each layer, for 10 epochs, as described by Badjatiya et al. (2018). We use mean-pooled 300-dimensional word2vec embeddings (Mikolov et al., 2013) trained on the Google News dataset, as input to the perceptron.

- **LSTM (L):** We train a neural model as described by Badjatiya et al. (2018), using the same pre-trained word2vec embedding matrix. The network consists of an Embedding layer, followed by an LSTM layer, a dropout layer, a dense layer and finally, a sigmoid activation layer.

Our models outperform the baselines on all metrics, with the BERT (full window) model for break prediction model giving the best results. The approaches by Reynar (1994) and Utiyama and Isahara (2001), and the neural models proposed by Badjatiya et al. (2018) and Koshorek et al. (2018) are global models, and are prohibitively expensive on long documents.

5 Global Break Prediction

In the approaches described above, we simply select the highest scoring P points. However, this selection does not conform to spatial constraints. For example, the model may place two breaks close to one another, when realistically, chapter breaks are spaced fairly apart in practice.

To validate this, we compute the coefficient of variance (CV) for each book, in terms of the number of sentences per chapter. Figure 6 shows the distribution of the CV over all books in our dataset. Most books in our dataset have a low CV (distributions with CV less than 1 are considered to be low-variance), reflecting the fact that chapters breaks are spaced fairly equally apart.




Figure 6: Frequency of distribution for the coefficient of variance of number of sentences per segment

Hence, we propose a dynamic programming approach, in order to incorporate a weight for keeping the chapter breaks equidistant.

We formulate the task in the same way as described previously, with an additional parameter α , which determines the importance of the confidence scores as compared to equidistant breaks. α ranges from 0 to 1, where 1 indicates

that full weight is given to confidence scores, and 0 indicates that full weight is given to keeping the breaks equidistant. We define the cost of inserting a break at point n and inserting k breaks in points 0 to $n - 1$ recursively as:

$$\text{cost}(n, k) = \min_{i \in [0, n-1]} \left(\text{cost}(i, k-1) + (1-\alpha) \frac{|n-i|}{L} \right) - \alpha \cdot s_n$$

where s_n is the confidence score for n being a break point, and L is the ideal chapter length, i.e. number of sentences in each chapter if the book is split into $P + 1$ equal parts. At each step, we use the break point which results in cost minimization as the next break point, and repeat the recursive call.

5.1 Experiments

We apply dynamic programming for global break prediction, to both the approaches described above. We conduct experiments for α from 0 to 1, with a step increase of 0.2.

5.1.1 WOC

We use the prominences of local minima obtained using WOC, with window sizes 50, 100, 150, and 200 respectively. We apply min-max normalization on the prominences.




Figure 7: WindowDiff error metric for WOC

Figure 7 shows the WindowDiff error metric for the WOC approach, with differing window sizes, for different values of α . (Note that the window sizes here are in terms of the number of neighboring sentences used to compute density, and not used while calculating the WindowDiff metric.) The figure shows that an increase in window size results in lower error, and for all window sizes, $\alpha = 0.8$ shows the best performance.

5.2 BBP

We use the confidence scores for class 0 obtained using the BERT model. We observe that confidence

scores are clustered close to 0 and 1. Higher confidence scores are of more interest to us, as they are indicative of potential chapter boundaries. Hence, in order to distribute the values closer to 1 further apart, we apply the log function and compute the modified confidence score as $-\ln(1 - \text{score})/c$, where c is a normalizing constant. In practice, we use $c = 10$ to limit a majority of the values between 0 and 1. We optimize for the best value of alpha independently of this constant.




Figure 8: WindowDiff error metric for BBP

Figure 8 shows the WindowDiff error metric for the BERT-based approach for the single paragraph and full window models respectively. We use thresholds of 0.9 and 0.99 for each of the models, meaning that we consider only those break points with confidence scores above the threshold as potential break point candidates.

The full window model shows the least error at $\alpha = 0.8$. Note that a higher threshold of 0.99 shows a performance almost equal to that of 0.9, since a higher threshold means fewer potential break point candidates, and hence a lower runtime. Figures 5a and 5b depict predictions from the WOC and BBP approaches respectively, with $\alpha = 0.8$.

Algorithm	P_k	WD	F1
Best BBP (local)	0.303	0.384	0.447
WOC (window=50)	0.443	0.456	0.144
WOC (window=100)	0.426	0.440	0.158
WOC (window=150)	0.421	0.434	0.162
WOC (window=200)	0.420	0.433	0.164
BBP (single para.)	0.441	0.455	0.128
BBP (full window)	0.284	0.305	0.453

Table 3: Metrics for global chapter break insertion

Table 3 shows the evaluation metrics for global chapter break insertion. The dynamic programming approach consistently improves the WindowDiff and F1 metrics. The BERT model (full window) gives the best performance in terms of all

Algorithm	MSE	MAE	R2	Pk	WD	F1
Baseline (# sent)	205.97	8.928	0.44	-	-	-
WOC (win=50)	203.26	8.797	0.45	0.46	0.50	0.13
WOC (win=100)	203.23	8.804	0.45	0.45	0.49	0.14
WOC (win=150)	203.19	8.805	0.45	0.44	0.49	0.14
WOC (win=200)	203.17	8.804	0.45	0.44	0.49	0.14
BBP (thr=0.9)	192.22	8.366	0.48	0.33	0.38	0.41
BBP (thr=0.99)	188.08	8.155	0.49	0.32	0.38	0.41

Table 4: Evaluation metrics for regression to predict number of chapter breaks in a book (Window size [WOC] and threshold [BBP] denoted in parentheses)




Figure 9: Error distribution over test set using predictions for number of chapters from BBP (threshold=0.99)

three metrics.

5.3 Estimating the Number of Breaks

The models described above require the number of chapter boundaries to be specified. We now address the independent question of estimating how many chapter breaks to insert.

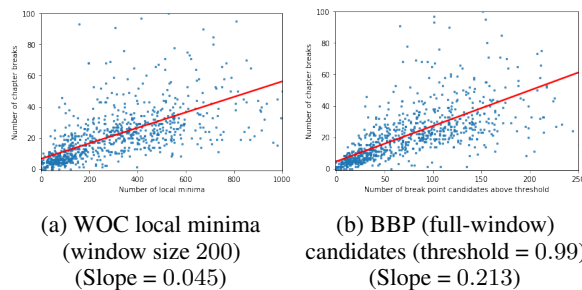


Figure 10: Number of chapter breaks as a function of the number of candidate break points

Figure 10 shows the number of chapters against the number of break point candidates for both the approaches. The number of local minima in WOC are approximately 20 times the number of chapter breaks, reflecting potential event boundaries within chapters. The number of break point candidates obtained using BERT are approximately 5 times the number of chapter breaks. This can also be seen in Figures 5a and 5b. Although the BBP model performs better at exact break prediction, the WOC model provides more information in terms of events within chapters.

We now use a regression model to predict the number of breaks, with the number of candidate break points and the total number of sentences in the book, as features. For the number of candidate breaks, we use:

- **WOC:** The total number of local minima
- **BBP:** The number of candidate break points above a certain threshold.

We perform experiments on 2,626 books in the test set, so as to keep the results comparable for both the approaches. We perform a train-test split of 67-33%. We predict the number of chapter breaks using this regression, and further evaluate global break prediction with $\alpha = 0.8$.

Table 4 shows the evaluation metrics on the test set, for regression using the models described above. The full-window BERT model shows the best performance in predicting the number of chapter breaks as well as break locations. Figure 9 shows the error distribution over the test set for the best performing model.

6 Conclusion and Future Work

We build a chapter segmentation dataset resource consisting of 9,126 English fiction novels, using a hybrid approach combining neural inference and regular expression-based rule matching. We achieve an F1 score of 0.77 on this task. Further, we use this dataset, remove structural cues, and address the task of predicting chapter boundaries. We present two methods for chapter segmentation. Our supervised approach achieves the best performance in exact break prediction, while our unsupervised approach provides information about potential sub-chapter break points.

Our work opens up avenues for further research in text segmentation, with potential applications in summarization and discourse analysis. Potential future work includes combining the neural and cut-based approaches into a stronger method. Finally, it would be interesting to do a deeper dive into variations of author strategies in chapterization, focusing more intently on books with large numbers of short chapters as being more reflective of episode boundaries.

Acknowledgments

We thank the anonymous reviewers for their helpful feedback. This work was partially supported by NSF grants IIS-1926751, IIS-1927227, and IIS-1546113.

References

- Alexander A Alemi and Paul Ginsparg. 2015. Text Segmentation based on Semantic Word Embeddings. *arXiv preprint arXiv:1503.05543*.
- Pinkesh Badjatiya, Litton J Kurisinkel, Manish Gupta, and Vasudeva Varma. 2018. Attention-based neural text segmentation. In *European Conference on Information Retrieval*, pages 180–193. Springer.
- Regina Barzilay and Igor Malioutov. 2006. Minimum Cut Model for Spoken Lecture Segmentation. In *In Proceedings of the Annual Meeting of the Association for Computational Linguistics (COLING-ACL 2006)*. Citeseer.
- Doug Beeferman, Adam Berger, and John Lafferty. 1999. Statistical Models for Text Segmentation. *Machine learning*, 34(1-3):177–210.
- Freddy YY Choi. 2000. Advances in Domain Independent Linear Text Segmentation. *arXiv preprint cs/0003083*.
- Freddy YY Choi, Peter Wiemer-Hastings, and Johanna D Moore. 2001. Latent Semantic Analysis for Text Segmentation. In *Proceedings of the 2001 conference on empirical methods in natural language processing*.
- Hervé Déjean and Jean-Luc Meunier. 2005. Structuring Documents According to Their Table of Contents. In *Proceedings of the 2005 ACM symposium on Document engineering*, pages 2–9.
- Hervé Déjean and Jean-Luc Meunier. 2009. On tables of contents and how to recognize them. *International Journal of Document Analysis and Recognition (IJ DAR)*, 12(1):1–20.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Antoine Doucet, Gabriella Kazai, Sebastian Colutto, and Günter Mühlberger. 2013. Overview of the IC-DAR 2013 Competition on Book Structure Extraction. In *2013 12th International Conference on Document Analysis and Recognition*, pages 1438–1443. IEEE.
- Jacob Eisenstein. 2009. Hierarchical Text Segmentation from Multi-scale Lexical Cohesion. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 353–361.
- Jacob Eisenstein and Regina Barzilay. 2008. Bayesian Unsupervised Topic Segmentation. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 334–343.
- Goran Glavaš, Federico Nanni, and Simone Paolo Ponzetto. 2016. Unsupervised Text Segmentation using Semantic Relatedness Graphs. Association for Computational Linguistics.
- Project Gutenberg. n.d. www.gutenberg.org. Accessed: May 2020.
- Marti A Hearst. 1994. Multi-Paragraph Segmentation of Expository Text. In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics*, pages 9–16. Association for Computational Linguistics.
- Shafiq Joty, Giuseppe Carenini, Raymond Ng, and Gabriel Murray. 2019. **Discourse analysis and its applications**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 12–17, Florence, Italy. Association for Computational Linguistics.
- Anna Kazantseva and Stan Szpakowicz. 2011. Linear Text Segmentation using Affinity Propagation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 284–293. Association for Computational Linguistics.
- Omri Koshorek, Adir Cohen, Noam Mor, Michael Rotman, and Jonathan Berant. 2018. Text Segmentation as a Supervised Learning Task. *arXiv preprint arXiv:1803.09337*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. **The Stanford CoreNLP natural language processing toolkit**. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Lara McConnaughey, Jennifer Dai, and David Bamman. 2017. The Labeled Segmentation of Printed Books. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 737–747.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Lev Pevzner and Marti A Hearst. 2002. A Critique and Improvement of an Evaluation Metric for Text Segmentation. *Computational Linguistics*, 28(1):19–36.

- Jeffrey C Reynar. 1994. An Automatic Method of Finding Topic Boundaries. In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics*, pages 331–333. Association for Computational Linguistics.
- Martin Riedl and Chris Biemann. 2012. Text Segmentation with Topic Models. *Journal for Language Technology and Computational Linguistics*, 27(1):47–69.
- Makoto Sakahara, Shogo Okada, and Katsumi Nitta. 2014. Domain-independent Unsupervised Text Segmentation for Data Management. In *2014 IEEE International Conference on Data Mining Workshop*, pages 481–487. IEEE.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. Literary Event Detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634.
- Masao Utiyama and Hitoshi Isahara. 2001. A Statistical Model for Domain-independent Text Segmentation. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, pages 499–506.
- Zhaohui Wu, Prasenjit Mitra, and C Lee Giles. 2013. Table of Contents Recognition and Extraction for Heterogeneous Book Documents. In *2013 12th International Conference on Document Analysis and Recognition*, pages 1205–1209. IEEE.
- Jonathan P Yamron, Ira Carp, Larry Gillick, Steve Lowe, and Paul van Mulbregt. 1998. A Hidden Markov Model Approach to Text Segmentation and Event Tracking. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, volume 1, pages 333–336. IEEE.