

ABLIT: A Resource for Analyzing and Generating Abridged Versions of English Literature

Melissa Roemmele Kyle Shaffer Katrina Olsen Yiyi Wang Steve DeNeefe

Language Weaver (RWS Group)

{mroemmele, kshaffer, kolsen, yiyiwang, sdeneefe}@rws.com

Abstract

Creating an abridged version of a text involves shortening it while maintaining its linguistic qualities. In this paper, we examine this task from an NLP perspective for the first time. We present a new resource, ABLIT, which is derived from abridged versions of English literature books. The dataset captures passage-level alignments between the original and abridged texts. We characterize the linguistic relations of these alignments, and create automated models to predict these relations as well as to generate abridgements for new texts. Our findings establish abridgement as a challenging task, motivating future resources and research. The dataset is available at github.com/roemmele/AbLit.

1 Introduction

An abridgement is a shortened form of a text that maintains the linguistic qualities of that text¹. It is intended to make the original text faster and easier to read. In this paper, we propose abridgement as an NLP problem and describe its connection to existing inference and generation tasks. We present a novel dataset for this task, focused on abridged versions of English literature books, which we refer to as the ABLIT dataset. We demonstrate the characteristics of ABLIT in terms of the relations between original and abridged texts as well as the challenges of automatically modeling these relations. The dataset and all associated code, including a Python package for easily interfacing with the data, are available at: github.com/roemmele/AbLit.

2 The task of abridgement

2.1 Definition

We define abridgement as the task of making a text easier to understand while preserving its linguistic

¹The term “linguistic qualities” is broad, which reflects other definitions of abridgement. For instance, the [Wikipedia entry](#) for “abridgement” specifies that it “maintains the unity of the source”, but these dimensions of unity are tacitly defined.

qualities. As such, abridgement intersects with tasks that fuse natural language inference (NLI) and natural language generation (NLG), in particular summarization and simplification.

Summarization condenses the main content of a text into a shorter version in order to facilitate high-level comprehension of the content. Existing research has used the categories of *extractive* and *abstractive* to describe summaries. In the former, the summary ‘extracts’ sequences from the text, whereas in the latter the summary ‘abstracts’ out the meaning of the text and rewrites it. The degree of abstractiveness of a summary is indicated by how much novel text it contains that is not directly in the original text. Like a summary, an abridgement is shorter than its original text, but it preserves more of its language and can be seen as an alternative version rather than a meta-description. According to how summaries are characterized, abridgements are highly extractive, even if some abstraction is needed to connect the extracted components.

Some work has examined summarization of narratives, including literary text (Kazantseva and Szpakowicz, 2010; Mihalcea and Ceylan, 2007; Zhang et al., 2019). Of particular relevance to our work are datasets released by Chaudhury et al. (2019), Kryściński et al. (2021), and Ladhak et al. (2020), all of which consist of summaries of fiction books. These summaries are significantly different from abridgements in that they are highly abstractive; they convey the book’s narrative without preserving much of the text itself. Kryściński et al. provides summaries at different levels of granularity (book, chapter, and paragraph). Their analysis demonstrates that even the finer-grained summaries at the paragraph level are quite abstractive.

The task of simplification also aims to make a text easier to understand, but without significantly distilling its content. Simplification is often treated as a sentence-level task (Sun et al., 2021). Abridgement can be viewed as simplification on a docu-

ment level. It seeks to balance the goal of increasing readability with preserving as much of the original text as possible. Research on simplification has been constrained by a lack of high-quality publicly available datasets. Existing datasets have been derived from sources like Wikipedia (e.g. [Coster and Kauchak, 2011](#)) and news articles ([Xu et al., 2015](#)), but none have focused on literary text.

2.2 Practical application

There are few authors who perform abridgement, and thus relatively few abridged versions of books ([Minshull, 2001](#)). Authors have described it as challenging and time-consuming to discern what to modify without compromising the original author’s agency ([Lauber, 1998](#); [Sussman, 1988](#)). However, as touted by these authors, abridgement makes books more accessible to a larger audience, especially when delivering the content through non-text modes like audio ([Lavin, 2014](#)). Given this, automating the abridgement process could vastly expand the number of abridged versions of books and thus increase their readership. Automation does not preclude the involvement of human authors; for example, human translators use machine translation to increase their productivity (e.g. [Zhechev, 2014](#)), and the same paradigm could apply to abridgement.

3 Creating an abridgement dataset

The ABLIT dataset is derived from 10 classic English literature books, listed in [A.4](#). These books are in the public domain and available through Project Gutenberg². A single author, Emma Laybourn, wrote abridged versions of these books that are also freely available³. The author explains:

“This is a collection of famous novels which have been shortened and slightly simplified for the general reader. These are not summaries; each is half to two-thirds of the original length. I’ve selected works that people often find daunting because of their density or complexity: the aim is to make them easier to read, while keeping the style intact.”

Informed by this, we designed ABLIT to capture the alignment between passages in a text and its abridged version. We specify that an abridged and original passage are aligned if the content of the abridged passage is fully derived from the original.

After obtaining the original and abridged books from their respective sites, we detected chapter

headings to split the books into chapters (see [A.1](#) for details). We paired the original and abridged version of each chapter according to these headings. Obviously, the two versions already form a broad alignment unit, but our goal was to examine finer levels of alignment. We chose to use sentences as the minimal alignment units, since they are intuitive units of expression in text and can be detected automatically⁴. ABLIT annotates sentence boundaries by indexing their position in the text, which enables all whitespace characters (most importantly, line breaks marking paragraphs) to be preserved.

3.1 Automated alignments

We pursued an automated approach to establish initial alignments between the original and abridged sentences for each chapter. It follows the same dynamic programming scheme used to create the Wikipedia Simplification dataset ([Coster and Kauchak, 2011](#)). We refer to a group of adjacent sentences in a text as a *span*. We define the length of a span by the number of sentences it contains. Each span o of length o_n in the original version of a chapter is paired with a span a of length a_m in the abridged version. The value of a_m can be zero, allowing for the possibility that an original sentence is aligned with an empty string.

For each pair of o and a , we use a similarity metric $sim(o, a)$ to score the likelihood that they are aligned. This scoring function also considers the length of the spans in order to optimize for selecting the narrowest alignment between the original and abridged text. For instance, if a one-to-one alignment exists such that the meaning of a single sentence in the abridgement is fully encapsulated by a single original sentence, these sentences should form an exclusive alignment. To promote this, we adjust $sim(o, a)$ by a penalty factor pn applied to the size of the pair, where $size = \max(o_n, a_m)$. Ultimately, the alignment score for a given span pair (o, a) is $\max(0, sim(o, a) - ((size - 1) * pn))$. Thus, more similar pairs obtain higher scores, but the scores are increasingly penalized as their size increases. At each sentence position in the original and abridged chapters, we score spans of all lengths $[1, o_n]$ and $[0, a_m]$, then select the one that obtains the highest score when its value is combined with the accumulated score of the aligned spans prior to

²gutenberg.org

³englishliteraturebooks.com

⁴We used nltk.org for all sentence segmentation and word tokenization. For analyses pertaining to words, words are lowercased without any other normalization (e.g. lemmatization).

that position. Once all span pairs are scored, we follow the backtrace from the highest-scoring span in the final sentence position to retrieve the optimal pairs for the chapter. We refer to each resulting span pair (o, a) in this list as an alignment *row*.

3.2 Assessment of automated alignments

We applied this automated alignment approach to the first chapter in each of the ten books in ABLIT, which we designated as an *assessment* set for investigating the quality of the output rows. We instantiated $\text{sim}(o, a)$ as the ROUGE-1 (unigram) precision score⁵ between the spans, where a is treated as the hypothesis and o is treated as the reference. Here we refer to this score as $R-1_p$. It effectively counts the proportion of words in a that also appear in o . We considered values of o_n in $[1, 6]$ and a_m in $[0, 6]$ and selected $o_n = 3$ and $a_m = 5$ based on our perceived quality of a sample of output rows. We similarly optimized pn values in $[0, 0.25]$ and selected $pn = 0.175$. Smaller values of pn yielded rows that were not minimally sized (i.e. they needed to be further split into multiple rows), whereas larger values tended to wrongly exclude sentences from rows.

The output consisted of 1,126 rows, which were then reviewed and corrected by five human validators recruited from our internal team. Validators judged a row as correct if the meaning expressed by the abridged span was also expressed in the original span, consistent with how alignment is defined above. A.3 gives more detail about this task. We found that inter-rater agreement was very high (Fleiss' $\kappa = 0.984$) and the few disagreements were easily resolved through discussion. The validators reported spending 10-15 minutes on each chapter.

After establishing these gold rows for the assessment set, we evaluated the initial automated rows with reference to the gold rows. To score this, we assigned binary labels to each pair of original and abridged sentences, where pairs that were part of the same row were given a positive class label and all other pairs were given a negative class label. Given these labels for the rows automatically produced with the $R-1_p$ method compared against the labels for the gold rows, the F1 score of the automated rows was 0.967. We also evaluated other methods for computing $\text{sim}(o, a)$, but none outperformed $R-1_p$. See A.2 for the description of these alternative methods and their results.

⁵Using github.com/Diego999/py-rouge

3.3 Full dataset

Partial validation: The time spent validating this assessment set indicated it would require significant resources to review the rows for all 868 chapters across the 10 books. Meanwhile, our evaluation revealed that we can expect most automatically aligned rows to be correct. Thus, we considered how to focus effort on correcting the small percentage of erroneously aligned rows. We manually reviewed these rows in the assessment set and found that their $R-1_p$ scores were often lower than those of the correct rows. Moreover, this tended to affect two types of rows: those with two or more sentences in the abridged span, or those adjacent to another row with an empty abridged span (i.e. $a_m = 0$). We did an experiment where a human validator reviewed only the assessment rows with scores < 0.9 that qualified as one of the two above cases. Selectively applying corrections to just these rows boosted the F1 score of the assessment set from 0.967 to 0.99. We thus decided to apply this strategy of partially validating automated rows to create the training set for ABLIT.

Final sets: To construct the rest of ABLIT, we ran the automated alignment procedure on all other chapters, and then applied the above partial validation strategy. Because we previously confirmed high inter-rater agreement, a single validator reviewed each chapter. Generalizing from the assessment set, we estimate that 99% of these rows are correct. To ensure an absolute gold standard for evaluation, we set aside five chapters in each of the ten books and fully validated their rows. We designated this as the test set, and repurposed the assessment set to be a development set that we used accordingly in our experiments. All other chapters were assigned to the training set. Ultimately, ABLIT consists of 808, 10, and 50 chapters in the training, development, and test sets, respectively. Table 1 shows some examples of rows in ABLIT.

4 Characterizing abridgements

4.1 Overview

Table 2 lists the size of ABLIT in terms of rows, paragraphs, sentences, and words (see A.4 for these numbers compared by book). Here we call attention to the numbers for the fully-validated test set, but the numbers for the training set closely correspond. The development set slightly varies from the training and test set for a few statistics, likely

Original Span	Abridged Span
[<u>The letter</u> was not unproductive.] [It <u>re-established peace and kindness</u> .]	[The letter <u>re-established peace and kindness</u> .]
[<u>Mr. Guppy sitting on the window-sill</u> , nodding his head and balancing all these possibilities in his mind, continues <u>thoughtfully</u> to tap <u>it</u> , and clasp it, and measure it with his hand <u>until he hastily draws his hand away</u> .]	[<u>Mr. Guppy sitting on the window-sill</u> , <u>taps it thoughtfully</u> , until he hastily draws his hand away.]
[<u>At last the gossips thought they had found the key to her conduct, and her uncle was sure of it</u> ; and what is more, <u>the discovery</u> showed <u>his niece</u> to him in quite a new light, and he changed his whole deportment <u>to her accordingly</u> .]	[<u>At last the gossips thought they had found the key to her conduct, and her uncle was sure of it</u> .] [<u>The discovery altered his whole behaviour to his niece</u> .]
[<u>They trooped</u> down into the hall and into the carriage, Lady Pomona leading the way.] [<u>Georgiana stalked along</u> , passing <u>her father at the front door without condescending to look at him</u> .]	[<u>They trooped downstairs</u> , <u>Georgiana stalking along</u> .] [<u>She passed her father at the front door without condescending to look at him</u> .]

Table 1: Examples of alignment rows. Sentence boundaries are denoted by brackets ([]). We highlight preserved words in **blue** and underline the reordered ones. Added words are in **green**.

	Train	Dev	Test (Chpt Mean)
Chpts	808	10	50
Rows	115,161	1,073	9,765 (195)
O_{pars}	37,227	313	3,125 (62)
A_{pars}	37,265	321	3,032 (61)
O_{sents}	122,219	1,143	10,431 (209)
A_{sents}	98,395	924	8,346 (167)
$\%A_{sents}$	80.5	80.8	80.0
O_{wrds}	2,727,571	29,908	231,878 (4,638)
A_{wrds}	1,718,919	17,630	143,908 (2,878)
$\%A_{wrds}$	63.0	58.9	62.1

Table 2: Number of chapters (Chpts), alignment rows (Rows), paragraphs (pars), sentences (sents), and words (wrds) across all original (O) and abridged (A) books. The per-chapter means appear for the test set.

O_{sents}	A_{sents}	Train	Dev	Test
1	1	75.8	74.7	75.7
1	0	17.4	17.3	17.3
2+	1	4.3	4.8	4.6
1	2+	2.1	3.2	1.9
2+	2+	0.3	0.0	0.5

Table 3: Distribution of row sizes by number of sentences (sents) in original (O) and abridged (A) spans

because it is smaller. Judging by the test set, the abridged chapters have almost the same number of paragraphs as the original, but they have 80% of the number of sentences ($\%A_{sents}$) and $\approx 62\%$ of the number of words ($\%A_{wrds}$).

Table 3 pertains to the size of the original and abridged spans in each row, where size is the number of sentences in each span. The table shows the relative percentage of rows of each size. The

majority of test rows ($\approx 76\%$) contain a one-to-one alignment between an original and abridged sentence (i.e. $O_{sents} = 1, A_{sents} = 1$). Meanwhile, $\approx 17\%$ contain an original sentence with an empty abridged span ($O_{sents} = 1, A_{sents} = 0$). A minority of rows ($\approx 5\%$) have a many-to-one alignment ($O_{sents} = 2+, A_{sents} = 1$) and a smaller minority ($\approx 2\%$) have a one-to-many alignment ($O_{sents} = 1, A_{sents} = 2+$). Many-to-many alignments ($O_{sents} = 2+, A_{sents} = 2+$) are more rare (0.5%).

4.2 Lexical similarity

Score Bin	Train	Dev	Test
0.0	17.5	17.6	17.4
(0.0, 0.25]	0.1	0.2	0.1
(0.25, 0.5]	0.5	0.9	0.6
(0.5, 0.75]	2.6	4.6	2.9
(0.75, 1.0)	23.9	31.5	24.0
1.0	55.5	45.2	55.0

Table 4: Binned distribution of $R-1_p$ scores for rows

As demonstrated by the success of the $R-1_p$ metric for creating alignment rows (Section 3.2), an original and abridged span typically align if most of the words in the abridged are contained in the original. Table 4 shows the binned distribution of the $R-1_p$ scores for the rows. Rows with an exact score of 0.0 ($\approx 17\%$ of rows in the test set) consist almost exclusively of original spans aligned to empty spans, which is why this number is comparable to the second line of Table 3. Many rows have perfect scores of exactly 1.0 (55%), signifying that their abridged span is just an extraction of some or all of the original span. The abridged spans where this is not the case (i.e. they contain some words

not in the original) still copy much of the original: 24% of test rows have a $R-1_p$ score above 0.75 and below 1.0, while only a small minority ($\approx 4\%$, the sum of lines 2-4 in the table) have a score above 0.0 and below 0.75.

4.3 Lexical operations

	Train	Dev	Test
O_{rmv} (O_{prsv})	40.9 (59.1)	45.9 (54.1)	41.9 (58.1)
A_{add} (A_{prsv})	6.3 (93.7)	8.3 (91.7)	6.4 (93.6)
$Rows_{rmv}$	71.1	80.3	73.2
$Rows_{prsv}$	82.5	82.7	82.6
$Rows_{add}$	37.4	48.8	39.4
$Rows_{reord}$	11.8	16.5	11.7

Table 5: *Top*: the % of removed and added words relative to all original and abridged words, respectively. *Bottom*: the % of rows with each lexical operation.

For each row, we enumerate the common and divergent items between the words o_{wrds} in the original span and the words a_{wrds} in the abridged span. The words that appear in o_{wrds} but not a_{wrds} are removed words, i.e. $o_{rmv} = |o_{wrds} - a_{wrds}|$. All other original words are preserved in the abridgement, i.e. $o_{prsv} = |o_{wrds} - o_{rmv}|$. Accumulating these counts across all original spans $o \in O$, the top section of Table 5 indicates the percentages of removed and preserved words among all original words. In the test set, $\approx 42\%$ of words are removed, and thus $\approx 58\%$ are preserved. Next, we count the added words in the abridgement, which are those that appear in a_{wrds} and not o_{wrds} , i.e. $a_{add} = |a_{wrds} - o_{wrds}|$. All other abridged words are preserved from the original, i.e. $a_{prsv} = |a_{wrds} - a_{add}|$. Accumulating these counts across all abridged spans $a \in A$, Table 5 shows that only $\approx 6\%$ of abridged words in the test set are additions, and thus $\approx 94\%$ are preservations.

We also report the number of rows where these removal, preservation, and addition operations occur at least once. For instance, if $o_{rmv} > 0$ for the original span in a given row, we count that row as part of $Rows_{rmv}$. The bottom section of Table 5 shows the percentage of rows with each operation among all rows in the dataset. In $\approx 73\%$ of the test rows, the abridged span removes at least one word from the original. In $\approx 83\%$ of rows, the abridged span preserves at least one word from the original. In $\approx 39\%$ of rows, the abridged span adds at least one word not in the original. We considered the

possibility that preserved words could be reordered in the abridgement. To capture this, we find the longest contiguous sequences of preserved words (i.e. “slices”) in the abridged spans. A row is included in $Rows_{reord}$ if at least two abridged slices appear in a different order compared to the original span. This reordering occurs in $\approx 12\%$ of rows.

It is clear from this analysis that the abridgements are quite loyal to the original versions, but they still remove a significant degree of text and introduce some new text. The examples in Table 1 highlight these operations. We can qualitatively interpret from the examples that some added words in the abridged span are substitutions for removed original words (e.g. “tap” > “taps” in the second example, “changed” > “altered” in the third example). See A.5 for additional discussion about how some of these relations pertain to common NLI tasks.

4.4 Lexical categories

Category	% O	% O_{rmv}	% A	% A_{add}
Function	58.2	57.9	58.1	53.9
Content	41.8	42.1	41.9	46.1

Table 6: Test set distribution of lexical categories for removed words O_{rmv} compared with all original words O , and added words A_{add} compared with all abridged words A

We examined if certain types of words are more often affected by removal or addition operations. Table 6 contains a broad analysis of this for the test set. As shown, $\approx 58\%$ of original words O are function words (those with part-of-speech tags of punctuation, pronouns, adpositions, determiners, etc.), while $\approx 42\%$ are content words (nouns, verbs, adjectives, and adverbs). The category distribution of removed words O_{rmv} closely matches the O distribution, suggesting that both function and content words are removed at the same rate. The abridged words A have the same proportion of function and content words as O (again, $\approx 58\%$ and $\approx 42\%$). In comparison, $\approx 54\%$ of additions A_{add} are function words, while $\approx 46\%$ are content words. The gap between $\approx 42\%$ and $\approx 46\%$ indicates that content words are added at a slightly higher rate than the overall frequency in content words in A (and equivalently, function words are added at a lower rate). But there are few additions overall, so the abridgements retain the same word type distribution as the original texts. A.6 shows this same analysis for each specific part-of-speech tag among these types.

5 Predicting what to abridge

Garbacea et al. (2021) points out that a key (and often neglected) preliminary step in simplification is to distinguish text that could benefit from being simplified versus text that is already sufficiently simple. This is also an important consideration for abridgement, since it seeks to only modify text in places where it improves readability. Accordingly, we examine whether we can automatically predict the text in the original that should be removed when producing the abridgement. As explained in Section 4, a removed word could mean the author replaced it with a different word(s) in the abridgement, or simply excluded any representation of its meaning. However, both cases indicate some change is applied to that word.

We model this through a binary sequence labeling task. Given a passage with original tokens o_{toks} and corresponding abridged tokens a_{toks} , we assign each token t in o_{toks} the label of `preserved` ($l=0$) if it is also in a_{toks} , and otherwise the label of `removed` ($l=1$) if it is not in a_{toks} . Thus the task is to predict the label sequence $[l_1, l_2, \dots, l_n]$ from the token sequence $[t_1, t_2, \dots, t_n]$.

5.1 Model inputs

We can derive a token-label sequence from each alignment row, by which each original span corresponds to a single input instance. However, the size of these spans varies across rows. To produce models that handle texts where these span boundaries are not known in advance, we consider consistent-length passages whose boundaries can be automatically inferred. Thus the ABLIT interface can provide pairs where a fixed-length passage from the original chapter (i.e. a sentence, paragraph, or multi-paragraph chunk) is aligned to its specific corresponding abridged version.

We enable this by finding the respective positions of the longest common word sequences between the original and abridged spans. Each of these overlapping subsequences is represented as a slice of the original text with indices (o_i, o_j) mapped to a slice of the abridged text (a_i, a_j) . Then, given a passage in the original text with indices (o_l, o_m) , we find all enclosed slices (o_i, o_j) where $o_i \geq o_l$ and $o_j \leq o_m$. For each slice we retrieve its corresponding abridged slice (a_i, a_j) . Given the earliest text position $\min a_i$ and latest position $\max a_j$ among these abridged slices, the full abridgement for the passage at (o_l, o_m) is the

text covered by the indices $(\min a_i, \max a_j)$. As an example, consider the first line in Table 1. If retrieving abridgements for sentence-length passages, the first sentence in the original span “The letter was not unproductive.” will yield “The letter” as the abridgement. The second original sentence “It re-established peace and kindness” will yield the abridgement “re-established peace and kindness”. By varying the passage size, we can assess how much context beyond a single row is beneficial in modeling abridgements. See A.7 for more details.

5.2 Experiment

Model: To predict abridgement labels (`preserved/removed`), we use a ROBERTA-based sequence labeling model, which has been applied to several other NLI tasks (Liu et al., 2019). We divided chapters according to varying passage sizes and trained a separate model on the token-label sequences⁶ associated with each passage size. The passages are either sentences (detected by NLTK), paragraphs (detected by line breaks), or multi-paragraph ‘chunks’. Each chunk consists of one or more paragraphs of S sentences, such that paragraphs are combined into the same chunk when their total number of sentences does not exceed S . As an additional reference, we trained a model where each passage is an original span directly taken from a single alignment row. As explained in Section 5.1, these passages (termed Rows) vary in length. We did not train a model on the full chapters as inputs because the average length of these inputs (5,044 tokens) greatly exceeds the ROBERTA limit of 512. See A.8 for more details about the model.

Passage	Toks	P	R	F1
Rows	26	0.692	0.442	0.532
Sentences	24	0.677	0.453	0.535
Paragraphs	81	0.686	0.460	0.546
Chunks ($S=10$)	303	0.670	0.501	0.569
All=removed	-	0.415	1.000	0.583

Table 7: F1 scores of abridgement label prediction for test set with models trained on varying passage sizes. **Toks** is the mean number of tokens in each passage type.

Results: Each model is evaluated on instances of the corresponding passage size in the test set. Table

⁶A “token” in this case is a sub-token unit defined by the ROBERTA tokenizer, rather than a whitespace-separated “word” pertaining to Section 4.

7 displays the results in terms of the precision (P), recall (R), and F1 score of predicting that a token should be removed. We compare these results with the baseline of labeling all tokens in the chapter as `removed` (final line). For chunks, we tuned different values of S in [5, 11] on the development set and observed the best F1 at $S=10$. The results show that the longest passage size (Chunks) yields the best predictions, suggesting the importance of chapter context beyond that given in a single row. The consistently higher precision over recall for all models indicates they correctly predict many preservation operations, but at the expense of missing many removal operations. Consequently, they overestimate the number of tokens that should be preserved. This results in an overall F1 that is lower than what occurs when all tokens are removed.

6 Producing abridgements

The above results show that anticipating what parts of a text should be changed when writing its abridged version is not trivial. The full task of producing an abridgement implicitly involves inferring these `preserved/removed` labels while additionally predicting the specific text that dictates these labels. We examine models that have been applied to tasks related to abridgement to establish benchmarks for this new task, with the intent that these benchmarks will inspire future work.

6.1 Models

We consider the following models to produce an abridged version of an original chapter:

Naive Baselines: As a reference point for our evaluation metrics, we report the performance of very weak baselines. In particular, we copy the entire original text as the abridgement (`COPY`). Alternatively, we select T percent of original tokens (`RANDEXTOKS`) as the abridgement.

Extractive Approaches: The analysis in Section 4.2 showed that abridgements preserve much of their original text, which motivates the use of extractive summarization methods. Using the best label prediction model from Section 5, we extract all original tokens labeled as `preserved` to form the abridgement (`EXTTOKS`). To reveal the maximum performance that can be obtained with this method, we also run it using the gold labels instead of predicted labels (`PERFECTEXTTOKS`). It is not conventional to use tokens as units of extraction,

since it can compromise fluency within sentences. `EXTTOKS` and `PERFECTEXTTOKS` only serve as points of comparison for our evaluation metrics. The standard extractive approach uses sentences as extractive units. For this (`EXTSENTS`), we form an abridgement by selecting a subset of sentences in the original chapter where at least P percent of tokens are labeled as `preserved`.

Generation Models: Extractive methods cannot introduce words into the abridgement that are not in the original, so for this we need to consider generation models. In particular, we examine two transformer-based sequence-to-sequence models that have been used for various generation tasks including summarization: T5-BASE (Raffel et al., 2020) (termed `TUNEDT5` here) and BART-BASE (Lewis et al., 2020) (termed `TUNEDBART`). We fine-tuned both models on the `ABLIT` training set, specifically on inputs consisting of chunks with 10 sentences, since this passage size yielded the best results in the Section 5 experiment. To assess the impact of these models' observation of `ABLIT`, we compare them with abridgements produced by prompting the non-finetuned T5-BASE to perform zero-shot summarization (`ZEROSHOTT5`). See A.9 for more details about these models. For all models, we generated an abridgement for an original chapter by dividing the chapter into chunks, generating output for each chunk (with 5-beam decoding), then concatenating the outputs to form the complete abridgement.

6.2 Evaluation metrics

We evaluate the predicted abridgements through comparison with the human-authored reference abridgements. First, we measure the word-based similarity between the predicted abridgement a_{pred} and reference abridgement a_{ref} using ROUGE-L ($R-L$), a standard evaluation metric for summarization. We then assess how accurately a_{pred} removed and preserved words from the original. A word from the original in a_{pred} is considered correctly preserved if it also appears in a_{ref} . We report the F1 of this measure as *Prsv*. A word in the original but not in a_{pred} is considered correctly removed if it is also absent from a_{ref} . We report the F1 of this measure as *Rmv*. Finally, we evaluate the accuracy of added words, where a word not in the original is considered correctly added to a_{pred} if it is also in a_{ref} . We report the F1 of this measure as *Add*. See A.10 for formal definitions of these metrics.

Name	Description	Toks	<i>R-L</i>	<i>Prsv</i>	<i>Rmv</i>	<i>Add</i>
HUMAN	Reference (a_{ref})	2,878	-	-	-	-
COPY	Duplicate original	4,638	0.739	0.753	0.000	0.000
RANDEXTOKS ($T=0.6$)	$T\%$ randomly selected original tokens	2,787	0.753	0.800	0.694	0.000
EXTTOKS	Original tokens predicted as <code>preserved</code>	3,160	0.818	0.856	0.745	0.006
PERFECTEXTTOKS	Original tokens where gold label is <code>preserved</code> (upper bound for EXTTOKS)	2,664	0.950	0.969	0.954	0.034
EXTSENTS ($P=0.65$)	Original sentences with $\geq P\%$ tokens predicted as <code>preserved</code>	2,857	0.792	0.824	0.720	0.001
TUNEDT5	Generate from finetuned T5	3,834	0.727	0.804	0.519	0.275
TUNEDBART	Generate from finetuned BART	3,673	0.780	0.815	0.583	0.365
ZEROSHOTT5	Generate from non-finetuned T5	1,157	0.416	0.484	0.627	0.019

Table 8: Scores of predicted abridgements on evaluation metrics. For all metrics, higher scores are better.

Original	Reference	TUNEDBART
The windows were half open because of the heat, and the Venetian blinds covered the glass, – so that a gray grim light, reflected from the pavement below, threw all the shadows wrong, and combined with the green-tinged upper light to make even Margaret’s own face, as she caught it in the mirrors, look ghastly and wan.	The windows were half open because of the heat, and Venetian blinds covered the glass, giving the light a green tinge that made her face in the mirrors look ghastly and wan.	The windows were half open because of the heat, and the Venetian blinds covered the glass - so that a grey grim light, reflected from the pavement below, threw all the shadows wrong, and made even Margaret’s own face look ghastly and wan.
We must suppose little George Osborne has ridden from Knightsbridge towards Fulham, and will stop and make inquiries at that village regarding some friends whom we have left there. How is Mrs. Amelia after the storm of Waterloo? Is she living and thriving? What has come of Major Dobbin, whose cab was always hankering about her premises?	We must now make inquiries at Fulham about some friends whom we have left there. How is Mrs. Amelia? Is she living and thriving? What has come of Major Dobbin?	We must suppose little George Osborne has ridden towards Fulham, and will stop and make inquiries about some friends whom we have left there. How is Mrs. Amelia after the storm of Waterloo? Is she living and thriving? What has come of Major Dobbin, whose cab was always hankering about her premises?

Table 9: Abridgements predicted by TUNEDBART for excerpts of North and South and Vanity Fair

6.3 Results

Table 8 reports the length and metric scores of the abridgements produced by each model for the test set chapters. Where applicable, we selected the T and P parameters from tuning on the development set. The results again convey that abridgement is largely a text extraction task, though a challenging one. The low $R-L$ score of ZEROSHOTT5 confirms that ABLIT is different from the summarization datasets that T5-BASE is trained on. The high $R-L$ of PERFECTEXTTOKS validates that precisely identifying which words to remove goes far in producing the abridgement. The high $Prsv$ scores for all approaches that observe ABLIT show they can all preserve the original text reasonably well. Analogous to the results in Section 5, the lower Rmv scores indicate knowing which words to remove is harder, particularly for the generation models. The extractive methods have no opportunity to obtain an Add score that is non-trivially above 0⁷.

⁷It is possible for Add to be slightly above 0 with the extractive approaches due to tokenization; see A.10.

The generation models do show a small benefit here in correctly adding some new words to the abridgement. The examples in Table 9 qualitatively represent the outcome for the TUNEDBART model. These abridgements remove some of the same original text as the reference and also add a few words consistent with the reference, but they still retain more of the original text than the reference. Other examples are shown in A.12.

7 Conclusion

In this paper, we introduced ABLIT, a corpus of original and abridged versions of English literature. ABLIT enables systematic analysis of the abridgement task, which has not yet been studied from an NLP perspective. Abridgement is related to other tasks like summarization, but has a stricter requirement to maintain loyalty to the original text. Our experiments motivate an opportunity to better balance this goal against that of improving readability. We also envision future resources that generalize this task to other texts beyond English literature.

8 Limitations

We present ABLIT to introduce abridgement as an NLP task. However, the dataset is scoped to one small set of texts associated with a specific domain and author.

There are significant practical reasons for this limited scope. In particular, most recently published books are not included in publicly accessible datasets due to copyright restrictions, and the same restrictions typically apply to any abridgements of these books. The books in ABLIT are uniquely in the public domain due to expired copyrights, and the author chose to also provide her abridgements for free. For this reason, ABLIT consists of British English literature from the 18th and 19th centuries. Some of the linguistic properties of these original books do not generalize to other types of English texts that would be useful to abridge. We do not yet know what aspects of abridgement are specific to this particular domain.

Moreover, as described in Section 2.2, creating abridgements is a rare and highly skilled writing endeavor. The ABLIT abridgements are written exclusively by one author. Without observing alternative abridgements for the same books by a different author, it is unclear what features are specific to the author’s preferences. This conflation between task and author is a concern for many NLP datasets (Geva et al., 2019). More generally, obtaining human writing expertise is a challenge shared by all language generation research as it becomes more ambitious (e.g. Wu et al., 2021).

9 Ethical Considerations

As stated in the introduction, all data and code used in this work is freely available. The text included in the dataset is in the public domain. Additionally, we explicitly confirmed approval from the author of the abridged books to use them in our research.

For the data validation task, the validators were employed within our institution and thus were compensated as part of their normal job role. Given that the dataset is derived directly from published books, it is possible that readers may be offended by some content in these books. The validators did not report any subjective experience of this.

With regard to our modeling approaches, large pretrained models like the ones we use here for generating abridgements have a well-known risk of producing harmful content (e.g. Gehman et al., 2020). For the generation models fine-tuned on ABLIT,

we did not subjectively observe any such text in the sample output we assessed. We judge that our controlled selection of training data reduces this risk, but does not eliminate it. Accordingly, future applications of abridgement can similarly consider careful data curation for mitigating this risk.

References

- Atef Chaudhury, Makarand Tapaswi, Seung Wook Kim, and Sanja Fidler. 2019. The shmoop corpus: A dataset of stories with loosely aligned summaries. *arXiv preprint arXiv:1912.13082*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32.
- William Coster and David Kauchak. 2011. [Simple English Wikipedia: A new text simplification task](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 665–669, Portland, Oregon, USA. Association for Computational Linguistics.
- 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.
- Cristina Garbacea, Mengtian Guo, Samuel Carton, and Qiaozhu Mei. 2021. [Explainable prediction of text complexity: The missing preliminaries for text simplification](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1086–1097, Online. Association for Computational Linguistics.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [RealToxicityPrompts: Evaluating neural toxic degeneration in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. [Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets](#). 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)*, pages 1161–1166, Hong Kong, China. Association for Computational Linguistics.

- Anna Kazantseva and Stan Szpakowicz. 2010. [Summarizing short stories](#). *Computational Linguistics*, 36(1):71–109.
- Wojciech Kryściński, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong, and Dragomir Radev. 2021. Booksum: A collection of datasets for long-form narrative summarization. *arXiv preprint arXiv:2105.08209*.
- Faisal Ladhak, Bryan Li, Yaser Al-Onaizan, and Kathleen McKeown. 2020. [Exploring content selection in summarization of novel chapters](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5043–5054, Online. Association for Computational Linguistics.
- Lynn Lauber. 1998. [Bookend; confessions of \(an\) abridger](#). *The New York Times*.
- Brittany Lavin. 2014. [Abridgement: What it can mean for your book](#). *Opyrus*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Rada Mihalcea and Hakan Ceylan. 2007. [Explorations in automatic book summarization](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 380–389, Prague, Czech Republic. Association for Computational Linguistics.
- Duncan Minshull. 2001. [The incredible shrinking book](#). *The Guardian*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- Renliang Sun, Hanqi Jin, and Xiaojun Wan. 2021. [Document-level text simplification: Dataset, criteria and baseline](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7997–8013, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vic Sussman. 1988. [The fine art of abridgement](#). *The Washington Post*.
- Jeff Wu, Long Ouyang, Daniel M Ziegler, Nisan Stiennon, Ryan Lowe, Jan Leike, and Paul Christiano. 2021. Recursively summarizing books with human feedback. *arXiv preprint arXiv:2109.10862*.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. [Problems in current text simplification research: New data can help](#). *Transactions of the Association for Computational Linguistics*, 3:283–297.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Weiwei Zhang, Jackie Chi Kit Cheung, and Joel Oren. 2019. Generating character descriptions for automatic summarization of fiction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7476–7483.
- Ventsislav Zhechev. 2014. Analysing the post-editing of machine translation at autodesk. *Post-editing of machine translation: Processes and Applications*, pages 2–24.

A Appendix

A.1 Detecting chapter boundaries

There is a one-to-one relation between each chapter in an original book and each chapter in its corresponding abridged version. Both the original and abridged version of the books include headings separating chapters. We automatically detected these headings through a set of regular expressions (e.g. matching lines specifying a chapter number and name with the regular expression “`^Chapter [0-9]+: *[a-zA-Z\s]*$`”).

However, there is variability in the format of the headings: some can span multiple lines, or specify a book and volume number in addition to the chapter identifier, or have numbers written in non-numerical form, for instance. The format also varies between the original and abridged version of the same book. Thus, we manually reviewed all detected chapter boundaries and fixed any erroneous or missed boundaries. Ultimately we ensure that each chapter in an original book is paired exactly with its abridged counterpart.

A.2 Additional automated alignment results

Table 10 shows the results of all methods we assessed for computing similarity between original and abridged spans to create alignment rows, compared alongside the best method of unigram ROUGE precision ($R-1_p$) reported in Section 3.2. A clear drawback to using unigram overlap to measure similarity is that it does not account for differences in word order. However, taking this into account by using bigrams instead of unigrams to calculate ROUGE precision (i.e. $R-2_p$) reduced the F1 to 0.935, likely because it added more sparsity to the overlap units. In addition to the word-based ROUGE metric, we assessed vector-based similarity encoded by different configurations of pretrained language models: BERT (Devlin et al., 2019), XLNET (Yang et al., 2019), XLM (Conneau and Lample, 2019), and ROBERTA (Liu et al., 2019). We used the HuggingFace Transformers implementation of these models: <https://huggingface.co/docs/transformers/index>. For each model we report the best result among size penalty (p_n) values in $[0, 0.25]$. As displayed, the vectors that obtained the best F1 came from BERT (Devlin et al., 2019), particularly BERT-BASE-UNCASED, which consists of 110M parameters. See additional details about this

model here: <https://huggingface.co/bert-base-uncased>. Ultimately, however, the result from BERT-BASE-UNCASED was still outperformed by $R-1_p$. As reported in Section 3.3, the resulting rows were further improved by applying the described partial validation strategy (final line of table).

A.3 Details about validation task

For each row, validators assessed whether the abridged span in the row was correctly aligned with the corresponding original span. As described in Section 3.2, a row is correct if the meaning of the abridged span can be derived from the original span. For a given row, if the abridged span expressed some meaning not contained in the original span, it either meant that some sentence(s) in the abridged chapter were incorrectly placed in that row, or some sentence(s) in the original chapter were incorrectly placed in a different row. In both cases, validators moved the wrongly placed sentence(s) to a row resulting in correctly aligned spans. We utilized Google Sheets as an interface for this task, which enabled validators to easily review and correct the rows. We produced a single spreadsheet per chapter, where each spreadsheet row corresponded to an alignment row. For the partial validation strategy, we designed a Google Apps Script (<https://developers.google.com/apps-script>) that visually highlighted spreadsheet rows qualifying for partial validation so that validators could specifically attend to those rows.

For the development (assessment) and test sets, there were a few cases where the validators edited the spans themselves in order to correct sentence segmentation errors (e.g. wrongly segmenting after honorifics like “Mr.”).

A.4 Size of ABLIT compared by book

Table 11 shows characteristics of the data for each book in terms of number of alignment rows, original words, and abridged words.

A.5 NLI challenges in ABLIT

Table 12 shows some examples of rows in ABLIT where modeling the relation between the original and abridged span involves NLI challenges like abstractive paraphrasing, figurative language interpretation, commonsense reasoning, and narrative understanding.

Similarity Metric	pn	P	R	F1
<i>Vector cosine similarity</i>				
BERT-BASE-UNCASED	0.21	0.963	0.952	0.957
BERT-BASE-CASED	0.22	0.948	0.934	0.940
BERT-LARGE-UNCASED	0.21	0.934	0.919	0.926
BERT-LARGE-CASED	0.21	0.944	0.935	0.939
XLNET-BASE-CASED	0.22	0.753	0.731	0.742
XLNET-LARGE-CASED	0.21	0.583	0.564	0.573
XLM-MLM-EN	0.21	0.821	0.816	0.818
ROBERTA-BASE	0.21	0.738	0.717	0.727
ROBERTA-LARGE	0.21	0.592	0.573	0.582
<i>Word overlap similarity</i>				
$R-1_p$	0.175	0.964	0.969	0.967
$R-2_p$	0.175	0.912	0.958	0.935
$R-1_p$	0.175	0.990	0.991	0.990
+ partial human validation				

Table 10: Extended results for accuracy of automated alignment methods

Book (Orig Author)	Train			Dev			Test		
	Rows (Chpts)	O_{wrds}	$\%A_{wrds}$	Rows	O_{wrds}	$\%A_{wrds}$	Rows	O_{wrds}	$\%A_{wrds}$
Bleak House (Charles Dickens)	17,948 (62)	390,857	63.2	24	935	20.0	1,746	38,132	62.9
Can You Forgive Her? (Anthony Trollope)	16,494 (74)	350,092	62.2	94	3,216	49.5	1,339	27,660	61.2
Daniel Deronda (George Eliot)	12,735 (64)	333,283	61.6	158	3,524	61.9	786	25,334	49.1
Mansfield Park (Jane Austen)	5,744 (42)	159,863	67.0	91	3,564	62.1	795	22,607	66.1
North and South (Elizabeth Gaskell)	8,922 (46)	193,355	67.9	184	4,907	68.5	1,169	23,159	70.0
Shirley (Charlotte Bronte)	12,027 (31)	235,888	63.2	253	5,987	57.4	1,031	23,369	60.4
The Way We Live Now (Anthony Trollope)	19,355 (94)	392,554	60.3	166	4,345	53.7	1,122	23,238	60.7
Tristram Shandy (Laurence Sterne)	4,805 (305)	216,984	66.7	5	439	77.0	69	3,972	72.3
Vanity Fair (W. M. Thackeray)	11,682 (62)	334,783	59.8	18	717	60.9	738	23,609	57.4
Wuthering Heights (Emily Bronte)	5,449 (28)	119,912	66.3	80	2,274	68.3	970	20,798	71.0
All	115,161 (808)	2,727,571	63.0	1,073	29,908	58.9	9,765	231,878	62.1

Table 11: Statistics for each book in the AbLit dataset, in terms of number of alignment rows, total original word (O_{wrds}), and proportional length of abridgement relative to original ($\%A_{wrds}$). The number of chapters in the training set for each book is shown; there is 1 chapter per book in the development set and 5 chapters per book in the test set.

Original Span	Abridged Span	Type of Challenge
Still there was not a word.	No one spoke.	Paraphrasing: abridgement has same meaning as original but no word overlap
But it is time to go home; my appetite tells me the hour.	But it is time to go home; I am hungry.	Interpretation of figurative language: abridgement replaces phrase “appetite tells me the hour” with more literal term “hungry”
“Daniel, do you see that you are sitting on the bent pages of your book?”	“Daniel, you are sitting on the bent pages of your book.”	Change in dialogue act: question in original is transformed into statement in abridgment
While she was at Matching, and before Mr. Palliser had returned from Monkshade, a letter reached her, by what means she had never learned. “A letter has been placed within my writing-case,” she said to her maid, quite openly. “Who put it there?”	While she was at Matching, a letter reached her, by what means she never learned, although she suspected her maid of placing it inside her writing-case.	Dialogue interpretation: abridgement summarizes the narrative event (suspecting maid of placing letter) conveyed by the spoken utterances in the original text (“A letter has been placed... she said to her maid.”)
“If you will allow me, I have the key,” said Grey. Then they both entered the house, and Vavasor followed his host up-stairs.	Mr. Grey unlocked the door of his house, and Vavasor followed him upstairs.	Commonsense inference: abridgement involves knowledge that doors are unlocked by keys, which is not explicit in the original text
George Osborne was somehow there already (sadly "putting out" Amelia, who was writing to her twelve dearest friends at Chiswick Mall), and Rebecca was employed upon her yesterday’s work.	George Osborne was there already, and Rebecca was knitting her purse.	Narrative inference: “knitting her purse” in the abridgement is the event referenced by “yesterday’s work” in the original, and resolving this requires knowledge of the previous text in the chapter
But Kate preferred the other subject, and so, I think, did Mrs. Greenow herself.	But Kate preferred the subject of the Captain, and so, I think, did Mrs. Greenow herself.	Elaboration: abridgement specifies “Captain” is the “other subject” implied in the original

Table 12: Examples of rows where alignment involves a language inference challenge

A.6 Extended lexical category analysis

Section 4.4 summarized the frequency of lexical categories for removed and added words in the ABLIT test set, relative to these frequencies among all words in the original and abridged texts. Table 13 additionally displays these percentages for all part-of-speech tags within the function and content word classes, along with examples of common words associated with each tag. We used the spacy library to perform part-of-speech tagging: <https://spacy.io/usage/linguistic-features>.

A.7 Comment about passage size variation

Because the method for converting rows into passages of a consistent length (i.e. sentences, para-

graphs, chunks) relies on string matching, the boundaries of the abridged passage may be off by one or a few words, which occurs less frequently as the size of the passages increase. This tends to occur when a word at the end of the original passage is replaced by a synonym in the abridged passage. However, a manual review of our assessment set revealed that only 0.4% of sentences in the original text yielded abridgements with imprecise boundaries, and no paragraphs (and consequently no chunks) had this issue.

A.8 Details about binary prediction model

For all passage sizes, we initialized models with the ROBERTA-BASE weights using the HuggingFace Transformers implemen-

Category	% <i>O</i>	% <i>O_{rmv}</i>	% <i>A</i>	% <i>A_{add}</i>	Examples of Common Words
Function words	$\Sigma=58.2$	$\Sigma=57.9$	$\Sigma=58.1$	$\Sigma=53.9$	
Punctuation	14.0	12.9	15.7	23.3	, " . -- ; - ? !
Pronoun	11.0	10.2	11.5	8.7	i he it his her
Adposition	10.2	11.5	9.0	7.0	of in to with for
Determiner	7.8	8.3	7.1	3.9	the a an no all
Aux. Verb	6.4	6.0	6.5	3.8	was had be is been
Coord. Conj.	3.7	4.0	3.4	2.1	and but or nor both
Particle	2.7	2.6	2.7	2.2	to not 's n't '
Subord. Conj.	2.3	2.5	2.2	2.8	that as if when upon
Content words	$\Sigma=41.8$	$\Sigma=42.1$	$\Sigma=41.9$	$\Sigma=46.1$	
Noun	14.5	15.4	13.7	14.0	time man day way hand
Verb	10.4	10.3	11.1	17.1	said had know do have
Adjective	6.6	6.9	6.2	5.3	little own other such good
Adverb	5.0	5.3	4.7	5.3	so very as now then
Proper Noun	4.4	3.3	5.2	3.4	mr. mrs. sir miss lady
Other	1.0	0.9	1.0	1.1	one two oh no yes

Table 13: Distribution of part-of-speech categories for the set of all removed words O_{rmv} and all added words A_{add} in the ABLIT test chapters. These numbers are respectively compared alongside those for the total set of all original words O and all abridged words A . (Aux.=Auxiliary, Coord.=Coordinating, Conj.=Conjunction, Subord.=Subordinate)

tation: https://huggingface.co/docs/transformers/v4.16.2/en/model_doc/roberta#transformers.

RobertaModel. ROBERTA-BASE consists of 125M parameters (<https://huggingface.co/roberta-base>).

The maximum sequence length allowed by this model is 512, so we truncated all input tokens beyond this limit. We fine-tuned each model for 5 epochs, saving model weights after each epoch of training, and selected the model with the highest F1 score on the development set to apply to our test set. We used the AdamW optimizer (Loshchilov and Hutter, 2017) and a batch size of 16. It took ≈ 2 hours to train each model on a g4dn.2xlarge AWS instance. During evaluation, any input tokens beyond the model length limit were assigned the default label of `preserved`. The result for each model reported in Table 7 is based on a single run of the training procedure.

A.9 Details about generation models

Both TUNEDT5 and TUNEDBART were fine-tuned using the HuggingFace transformers library, in particular this script: http://github.com/huggingface/transformers/blob/master/examples/pytorch/summarization/run_summarization.py. TUNEDT5

was initialized from T5-BASE (Raffel et al., 2020), which consists of ≈ 220 M parameters (<https://huggingface.co/t5-base>).

For this model, we prepended the prefix “summarize: ” to the target (i.e. the abridged passage), consistent with how T5-BASE was trained to perform summarization. TUNEDBART was initialized from BART-BASE (Lewis et al., 2020), which consists of 140M parameters (<https://huggingface.co/facebook/bart-base>). For both TUNEDT5 and TUNEDBART, we used a maximum length of 1024 for both the source (original passage) and target (abridged passage), and truncated all tokens beyond this limit. We evaluated each model on the development set after each epoch and concluded training when cross-entropy loss stopped decreasing, thus saving the model weights with the optimal loss. We used a batch size of 4. For all other hyperparameters we used the default values set by this script, which specifies AdamW for optimization. It took ≈ 3 hours to train each model on a g4dn.4xlarge AWS instance. The result for each model reported in Table 8 is based on a single run of the training procedure.

A.10 Details about evaluation metrics

Preservation: The formal definition of the preservation metric $Prsv$ is as follows. If

$o_{prsv}(a_{pred})$ are the words in the original that are preserved in the predicted abridgement, and $o_{prsv}(a_{ref})$ are the words in the original that are preserved in the reference abridgement, then we consider the number of correctly preserved words: $Correct_Prsv = |o_{prsv}(a_{pred}) \cap o_{prsv}(a_{ref})|$.

The precision of this measure $Prsv_p = \frac{Correct_Prsv}{o_{prsv}(a_{pred})}$ is the proportion of correctly preserved words among all preserved words in the predicted abridgement. The recall $Prsv_r = \frac{Correct_Prsv}{o_{prsv}(a_{ref})}$ is the proportion of correctly preserved words among all preserved words in the reference abridgement. $Prsv$ is the F1 of these precision and recall measures: $Prsv = 2 \frac{Prsv_p \cdot Prsv_r}{Prsv_p + Prsv_r}$.

Removal: The formal definition of the removal metric is as follows. If $o_{rmv}(a_{pred})$ are the words in the original that are removed in the predicted abridgement, and $o_{rmv}(a_{ref})$ are the words in the original that are removed in the reference abridgement, then we consider the number of correctly removed words: $Correct_Rmv = |o_{rmv}(a_{pred}) \cap o_{rmv}(a_{ref})|$. The precision of this measure $Rmv_p = \frac{Correct_Rmv}{o_{rmv}(a_{pred})}$ is the proportion of correctly removed words among all removed words for the predicted abridgment. The recall $Rmv_r = \frac{Correct_Rmv}{o_{rmv}(a_{ref})}$ is the proportion of correctly removed words among all removed words for the reference abridgement. Rmv is the F1 of these precision and recall measures: $Rmv = 2 \frac{Rmv_p \cdot Rmv_r}{Rmv_p + Rmv_r}$.

Addition: The formal definition of the addition metric is as follows. If $a_{add}(a_{pred})$ are the words in the predicted abridgement that do not appear in the original, and $a_{add}(a_{ref})$ are the words in the reference abridgement that do not appear in the original, then we consider the number of correctly added words: $Correct_Add = |a_{add}(a_{pred}) \cap a_{add}(a_{ref})|$. The precision of this measure $Add_p = \frac{Correct_Add}{a_{add}(a_{pred})}$ is the proportion of correctly added words among all added words in the predicted abridgement. The recall $Add_r = \frac{Correct_Add}{a_{add}(a_{ref})}$ is the proportion of correctly added words among all added words in the reference abridgement. Add is the F1 of these measures: $Add = 2 \frac{Add_p \cdot Add_r}{Add_p + Add_r}$.

A.11 Comment about addition scores

Regarding the above-zero scores of the extractive methods on the *Add* metric, there are two reasons for this. One reason is that the prediction model uses sub-tokens while the *Add* metric analyzes whitespace-separated words. Consequently, one sub-token may be predicted as `preserved` while others within the same word are predicted as `removed`. Isolated from these other sub-tokens, the `preserved` sub-token will be recognized as a new added word in the abridgement. The other reason is that a single word in the original may be split by the tokenizer into two words in the abridgement, or vice-versa. For example, we observed that “Mr.” gets split into two tokens (“Mr”, ‘.’) in some contexts and is treated as one token (“Mr.”) in others. If the original text represents this item as two tokens and both the extracted and reference abridgement represent it as a single token, then this single token will be counted as an added word in the extracted abridgement.

A.12 Examples of produced abridgements

Tables 14 and 15 below show excerpts of the abridgements produced by the EXTSSENT and TUNEDBART models, alongside the original chapter text and human-authored reference abridgement. The sentences in each excerpt are lined up to better visualize their differences.

Original	Reference	EXTSENTS	TUNEDBART
<p>Seven days glided away, every one marking its course by the henceforth rapid alteration of Edgar Linton's state.</p> <p>The havoc that months had previously wrought was now emulated by the inroads of hours.</p> <p>Catherine we would fain have deluded yet; but her own quick spirit refused to delude her: it divined in secret, and brooded on the dreadful probability, gradually ripening into certainty.</p> <p>She had not the heart to mention her ride, when Thursday came round; I mentioned it for her, and obtained permission to order her out of doors:</p> <p>for the library, where her father stopped a short time daily-the brief period he could bear to sit up-and his chamber, had become her whole world.</p> <p>She grudged each moment that did not find her bending over his pillow, or seated by his side.</p> <p>Her countenance grew wan with watching and sorrow, and my master gladly dismissed her to what he flattered himself would be a happy change of scene and society;</p> <p>drawing comfort from the hope that she would not now be left entirely alone after his death.</p>	<p>In the next seven days Edgar Linton's state grew rapidly worse.</p> <p>Catherine could no longer be deluded:</p> <p>she brooded on the dreadful probability of her father's death, gradually ripening into certainty.</p> <p>She had not the heart to mention her ride when Thursday came round. I obtained permission to send her out of doors:</p> <p>for her father's chamber had become her whole world.</p> <p>She grudged each moment that she did not spend bending over his pillow, or seated by his side.</p> <p>She grew pale with watching, and my master gladly dismissed her to what he thought would be a happy change of scene;</p> <p>drawing comfort from the hope that she would not now be left entirely alone after his death.</p>	<p>but her own quick spirit refused to delude her:</p> <p>She had not the heart to mention her ride, when Thursday came round; I mentioned it for her, and obtained permission to order her out of doors:</p> <p>Her countenance grew wan with watching and sorrow, and my master gladly dismissed her to what he flattered himself would be a happy change of scene and society;</p> <p>drawing comfort from the hope that she would not now be left entirely alone after his death.</p>	<p>Seven days glided away, every one marking its course by the rapid alteration of Edgar Linton's state.</p> <p>The havoc that months had previously wrought was now emulated by the inroads of hours.</p> <p>Catherine we would fain have deluded yet; but her own quick spirit refused to delude her: it brooded on the dreadful probability, gradually ripening into certainty.</p> <p>She had not the heart to mention her ride, when Thursday came round; I ordered her out of doors:</p> <p>for the library, where her father stopped daily - the brief period he could bear to sit up - and his chamber, had become her whole world.</p> <p>She grudged each moment that did not find her bending over his pillow, or seated by his side.</p> <p>Her countenance grew wan with watching and sorrow, and my master gladly dismissed her to what he flattered himself would be a happy change of scene and society;</p> <p>drawing comfort from the hope that she would not now be left entirely alone after his death.</p>

Table 14: Abridgements for an excerpt of Wuthering Heights, Chapter 27
3732

Original	Reference	EXTSENTS	TUNEDBART
<p>It happened that when I came home from Deal I found a note from Caddy Jellyby (as we always continued to call her), informing me that her health, which had been for some time very delicate, was worse and that she would be more glad than she could tell me if I would go to see her.</p> <p>It was a note of a few lines, written from the couch on which she lay and enclosed to me in another from her husband, in which he seconded her entreaty with much solicitude.</p> <p>Caddy was now the mother, and I the godmother, of such a poor little baby—such a tiny old-faced mite, with a countenance that seemed to be scarcely anything but cap-border, and a little lean, long-fingered hand, always clenched under its chin.</p> <p>It would lie in this attitude all day, with its bright specks of eyes open, wondering (as I used to imagine) how it came to be so small and weak.</p> <p>Whenever it was moved it cried, but at all other times it was so patient that the sole desire of its life appeared to be to lie quiet and think.</p>	<p>When I came home from Deal I found a note from Caddy, informing me that her health, which had been for some time very delicate, was worse and that she would be very glad if I would go to see her.</p> <p>It was a short note, written from her bed.</p> <p>Caddy was now the mother, and I the godmother, of such a poor little baby - such a tiny old-faced mite, with a little lean, long-fingered hand always clenched under its chin.</p> <p>It would lie in this attitude all day, with its bright specks of eyes open, wondering (I used to imagine) how it came to be so small and weak.</p> <p>Whenever it was moved it cried, but at all other times it lay quiet.</p>	<p>It happened that when I came home from Deal I found a note from Caddy Jellyby (as we always continued to call her), informing me that her health, which had been for some time very delicate, was worse and that she would be more glad than she could tell me if I would go to see her.</p> <p>Caddy was now the mother, and I the godmother, of such a poor little baby—such a tiny old-faced mite, with a countenance that seemed to be scarcely anything but cap-border, and a little lean, long-fingered hand, always clenched under its chin.</p> <p>It would lie in this attitude all day, with its bright specks of eyes open, wondering (as I used to imagine) how it came to be so small and weak.</p> <p>Whenever it was moved it cried, but at all other times it was so patient that the sole desire of its life appeared to be to lie quiet and think.</p>	<p>It happened that when I came home from Deal I found a note from Caddy Jellyby informing me that her health, which had been for some time very delicate, was worse and that she would be more glad than she could tell me if I would go to see her.</p> <p>It was a note of a few lines, written from the couch on which she lay and enclosed to me in another from her husband, in which he seconded her entreaty with much solicitude.</p> <p>Caddy was now the mother, and I the godmother, of such a poor little baby - such a tiny old-faced mite, with a countenance that seemed to be scarcely anything but cap-border, and a little lean, long-fingered hand, always clenched under its chin.</p> <p>It would lie in this attitude all day, with its bright specks of eyes open, wondering how it came to be so small and weak.</p> <p>Whenever it was moved it cried, but at all other times it was so patient that the sole desire of its life appeared to be to lie quiet and think.</p>

Table 15: Abridgements for an excerpt of Bleak House, Chapter 50