

# Novel Chapter Abstractive Summarization using Spinal Tree Aware Sub-Sentential Content Selection

Anonymous ACL submission

## Abstract

Summarizing novel chapters is a difficult task due to the length of the chapter to be summarized and the fact that summary sentences draw content from multiple sentences in the chapter. We present a pipelined extractive-abstractive approach where the extractive step filters the content that is passed to the abstractive component. Extremely lengthy input also results in a dataset highly skewed towards negative instances and we thus adopt a margin ranking loss for extraction to encourage separation between positive and negative input. To generate summary sentences that fuse information from different sentences, our extraction component operates at the constituent level; our novel approach to this problem enriches the text with spinal tree information which provides context to the extraction model. We show an improvement of 3.71 Rouge-1 points over the state-of-the-art on an existing novel chapter dataset.

## 1 Introduction

Research on summarizing novels (Mihalcea and Ceylan, 2007; Wu et al., 2017; Ladhak et al., 2020; Kryściński et al., 2021; Wu et al., 2021) has recently gained popularity following advancements in sequence to sequence pre-trained models (Zhang et al., 2019a; Lewis et al., 2019; Raffel et al., 2019) and in summarization of newswire datasets (Narayan et al., 2018; Hermann et al., 2015; Grusky et al., 2018). Novel chapters present challenges not commonly encountered when summarizing news articles. Phrases from multiple, non-contiguous sentences within the chapter are often fused to form new sentences for the summary. One would be inclined to use an abstractive approach, but the length of chapters (on average, seven times longer than news articles (Ladhak et al., 2020)) makes it unfeasible to use state of the art generative models, such as BART (Lewis et al., 2019) and even Longformer (Beltagy et al., 2020). Chapter length causes the additional problem of an imbalanced

dataset, as a much higher percentage of the input will not be selected for the summary than is typical in domains such as news.

To address these challenges, we adopt an extractive-abstractive architecture, where content is first selected by extracting units from the input and then an abstractive model is used on the filtered input to produce fluent text. Kryściński et al. (2021) benchmarked the extractive-abstractive architecture, first proposed by Chen and Bansal (2018), for novel summarization, but did not extend it. In this work, we propose several novel extensions to improve its performance on the novel chapter summarization task. The large amount of compression in novel chapter summarization (372 summary words per 5,165 chapter words on average) creates an extreme imbalance in the training data; a successful extractive summarization algorithm would have to discard most of the text. The standard practice of using Cross-Entropy loss (Good, 1992) when training a neural network model backfires in our case: a network that opts to discard everything will achieve near-perfect performance. We alleviate the issue by improving the margin structure of the minority class boundary using the Margin Ranking loss (Rosasco et al., 2004), which encourages separation between the two classes. Other study, such as Cruz et al. (2016), also shows that a pairwise ranking improves model performances on imbalanced data.

Second, in order to model the fusion of chapter phrases into summary sentences, we carry out extraction at the constituent level. Ladhak et al. (2020) also tried this approach, but with mixed results. They noted that sometimes the sub-sentential unit can be too small and, therefore, lack meaningful content (e.g., phrases such as “what has?” in the extractive summary, Table 1). These small unintelligible pieces can negatively affect the performance of the extractive model and, more importantly, the subsequent abstractive model. We hypothesize that




Figure 1: The encoding part of the model for spinal tree encoding. We don't show the CLS and SEP tokens here for space-saving purposes, but they are treated as in BertSumm (Liu and Lapata, 2019).

we can improve the performance of the extractive model—and, consequently, that of the downstream abstractive model—by augmenting the meaning of the extracted sub-sentential units using additional information from the sentence. To that end, we propose an enrichment process, during model training, where we augment the sub-sentential units with linguistic information. For this purpose, we use a spinal tree (Carreras et al., 2008; Ballesteros and Carreras, 2015) which carries information about both the dependency and the constituent structure of the segment. We encode the spine's information using a recurrent network and concatenate its output to the embedding of the token, as illustrated in Figure 1.

Our contributions are threefold: (1) We adopt an extractive-abstractive architecture, improving the decision boundary of the content selection by using a Margin Ranking loss, (2) perform extraction at the constituent level, introducing an enrichment process that uses spinal tree information and (3) show that our approach improves over the state-of-the-art with a 3.71 gain in Rouge-1 points.

## 2 Related Work

Several previous works on novel chapter summarization, such as Mihalcea and Ceylan (2007), Wu et al. (2017), Ladhak et al. (2020), Kryściński et al. (2021) and Wu et al. (2021), are closely related to ours. Mihalcea and Ceylan (2007) uses MEAD, an

unsupervised extractive summarization described in Radev et al. (2004); this approach includes features focusing on terms weighting that take into account the different topics in the text. In this work, topic boundaries are determined using a graph-based segmentation algorithm that uses normalized cuts (Malioutov, 2006). A similar line of work, including Mihalcea and Ceylan (2007) and Wu et al. (2017), also performs topic modelling with Latent Dirichlet Allocation (Blei et al., 2003) followed by greedy unsupervised extraction.

Conversely, Ladhak et al. (2020) experiment with extracting information at the sentence and at the syntactic constituent level, via a supervised learning approach. To train their model, they use an aligning process based on the weighted ROUGE scores between the reference and novel text to assign proxy extract labels, in the absence of manually annotated ground truth. Their results at the constituent level are mixed; human evaluation shows a lower performance of constituent extraction models presumably because the summaries are not very readable. Kryściński et al. (2021) construct a novel chapter dataset that is slightly larger than that of Ladhak et al. (2020) and benchmark existing summarization algorithms on the dataset. Wu et al. (2021), on the other hand, use a human-in-the-loop approach to obtain summaries via behaviour cloning and reward modelling.

## 3 Novel Chapter Summarization

We use a two-step process where we first run an extractive model (Mihalcea and Ceylan, 2007; Wu et al., 2017; Ladhak et al., 2020) to select informative content and then run a separate abstractive model (Lewis et al., 2019; Zhang et al., 2019a; Raffel et al., 2019) to produce a coherent and readable version of this content.

### 3.1 Dataset and Pre-processing

For our novel dataset, we use summary-chapter pairs collected by Ladhak et al. (2020) from Project Gutenberg and various study guide sources. The size of the dataset is 8,088 chapter/summary pairs<sup>1</sup>. The average length of the chapters is 5,165 words with the longest being 33,167 words<sup>2</sup>.

<sup>1</sup>Train/dev/test splits are 6,288/938/862

<sup>2</sup>We are aware that there is a larger dataset called BookSum (Kryściński et al., 2021), which uses similar sources; however, due to licensing issues, we are unable to use it in our work.

| Extracts from our best performance Extractive Model  |
|--|
| tess went down the hill to trantridge cross , and inattentively waited to take her seat in the van returning from chaseborough to shaston .<q>her mother had advised her to stay here for the night , at the house of a cottage-woman<q>what has ? ”<q>“ they say – mrs d’urberville says –<q>that she wants you to look after a little fowl-farm which is her hobby .<q>cried joan to her husband .   |
| Abstracts from our best performance Abstractive Model  |
| tess goes down the hill to trantridge cross and waits to take her seat in the van returning from chaseborough to shaston . her mother has advised her to stay for the night at the house of a cottage-woman who has a fowl-farm . joan tells her husband that mrs. d’urberville has written a letter asking her daughter to look after her poultry-farm  |
| Reference  |
| when tess returns home the following day. a letter from mrs. d’urberville offering her a job tending fowl awaits her . despite her mother ’s ecstatic eagerness , tess is displeased and looks instead for local jobs to earn money to replace the family ’s horse .alec d’urberville stops by and prompts her mother for an answer about the job . her efforts to find alternative work prove fruitless and so tess accepts d’urberville ’s offer . she remarks that mrs. d’urberville ’s handwriting looks masculine . |

Table 1: The outputs from two different models. The extract is obtained through a content selection model while the abstract is obtained by passing the extract into BART (Lewis et al., 2019) language generation model. The <q> tokens in the extract are the delimiters for constituents.

In order to prepare the data for the experiments, we follow the same pre-processing steps as Ladhak et al. (2020) to obtain the sub-sentential units and their alignment to reference summaries. In addition, we truncate chapters to 30k tokens to fit into the GPU memory<sup>3</sup>; as a result, a single chapter of the dataset is actually truncated.

### 3.2 Extractive Model

The extractive summarization task can be posed as a classic regression and ranking problem where the model produces a score for each of a given set of units and then ranks them based on that score. The top  $k$  units are then used as an extract. The input of our model is the sub-sentential units of the novel chapter text. We train the model with the oracle labels which we obtain from the alignment between sub-sentential units and reference summaries.

**Baseline** Our baseline is BERTSUMMEXT model (Liu and Lapata, 2019) modified as follows. First, we replaced the underlying Transformer models (Vaswani et al., 2017) with Longformers, which can better capture long context and requires less computing memory than BERT (Devlin et al., 2019). Second, we removed the inter-sentence Transformer layers stacked on top of the BERT output, to further reduce memory usage.

**Spinal Tree** A spinal tree is a dependency structure of a sentence that is augmented with constituent information (Carreras et al., 2008; Ballesteros and Carreras, 2015). For each sub-sentential unit, we retrieve the spinal tree parse by first using the constituency parser (Manning et al., 2014) and then apply Collins Head-Word Finder (Collins,

1997) to calculate the spines. We then encode<sup>4</sup> the spinal tree using bidirectional-GRU networks (Cho et al., 2014)<sup>5</sup>. We construct the input of the Longformer by concatenating the embeddings of the tokens<sup>6</sup>, the corresponding positional embeddings per token, and the encoding of the spines for each token via the bidirectional-GRU encoders, as illustrated in Figure 1.

**Ranking Loss** The baseline model uses the Cross-Entropy (CE) loss function and minimizes the loss via gradient descent. However, the CE loss function focuses on optimizing both the negative and positive labels at the same time. To compensate for the imbalance in our dataset, we add a Margin Ranking (MR) loss that gives the positive labels higher ranks than the negative labels<sup>7</sup>.

**Re-ordering Scheme** The default baseline of Liu and Lapata (2019) produces extracts with sub-sentential units that are ordered based on their score. This scheme, however, destroys the plot of the story. Hence, we re-order the units according to the original positional order in the source text, thus preserving the correct plot order in the story.

### 3.3 Abstractive Model

Since the extractive model outputs are sometimes incoherent and hard to read, we forward them to an abstractive model, with the goal to produce a more fluent and coherent result.

<sup>4</sup>We use the hidden size of 512

<sup>5</sup>We experimented with other architectures including bi-LSTM and found that bidirectional-GRU were the best.

<sup>6</sup>We use the embedding size of 768

<sup>7</sup>We also have tried the weighted CE loss function but we get worse results. We also found that training our model first with the CE loss function until convergence and then continuing using the MR loss gives the best result.

<sup>3</sup>We use Amazon AWS EC2 P4dn 40GB GPU memory

We use BART (Lewisa et al., 2019) as our engine for abstractive summarization. To train BART, we use the oracle extracts as the input source and the reference summaries as the target. During prediction, we use the output of our content selection model as the input source.

| Model                         | R1           | R2           | RL           | WMD          | BERTScore    |
|-------------------------------|--------------|--------------|--------------|--------------|--------------|
| Extractive                    |              |              |              |              |              |
| Oracle Ext                    | <b>46.75</b> | <b>14.27</b> | <b>45.64</b> | <b>0.633</b> | <b>0.823</b> |
| CB const R-wtd (Ladhak, 2020) | 36.62        | 6.9          | 35.4         | N/A          | N/A          |
| Longformer Ext                | 39.24        | 7.61         | 38.29        | 0.712        | 0.803        |
| + Spinal Information          | 39.35        | 7.62         | 38.45        | 0.711        | 0.802        |
| + Ranking Loss                | <b>39.48</b> | 7.63         | <b>38.58</b> | <b>0.708</b> | 0.802        |
| + Re-ordering                 | <b>39.48</b> | <b>7.70</b>  | <b>38.58</b> | <b>0.708</b> | <b>0.806</b> |
| Abstractive                   |              |              |              |              |              |
| Oracle Abs                    | <b>45.82</b> | <b>14.14</b> | <b>42.74</b> | 0.641        | 0.828        |
| BART Abs                      | 39.77        | 9.28         | 37.56        | 0.693        | 0.807        |
| + Spinal Information          | 39.83        | 9.33         | 37.61        | 0.691        | 0.807        |
| + Ranking Loss                | 39.88        | <b>9.35</b>  | 37.68        | 0.691        | 0.807        |
| + Re-ordering                 | <b>40.33</b> | 9.10         | <b>37.95</b> | <b>0.690</b> | <b>0.810</b> |

Table 2: ROUGE, Word Mover Distance and BERTScore for extractive and abstractive models.

## 4 Results

Examples of outputs from our best abstractive and extractive models are shown in Table 1. Here we report results from an automatic and a manual evaluation. From previous work, Ladhak et al. (2020) we included three models: CB Sent K, CB Sent R-wtd, and CB Const R-wtd. Their models are extractive and use the same pre-processing approach as ours. We also included the oracle for both the extractive and abstractive models.

### 4.1 Automatic Evaluation

We use three different metrics for automatic evaluation: ROUGE (Lin, 2004), BERTScore (Zhang et al., 2019b) and Word Mover Distance (WMD) (Kusner et al., 2015). ROUGE measures syntactic similarities between system and reference summaries and BERTScore and WMD measure semantic similarities. BERTScore measures similarities at the sentence level while WMD at the token level. We run each experiment three times using different random seeds and we report the mean score.

Table 2 shows our models performance against the baseline and previous works. Our best extractive model (Longformer Ext+spinal+Ranking+Re-ordering) outperforms previous work (CB const R-wtd) by 2.86 ROUGE-1, 0.8 ROUGE-2, and 3.18 ROUGE-L points. Meanwhile, the abstractive

model (BART Abs+spinal+Ranking+Re-ordering) outperforms previous work (CB const R-wtd) by 3.71 ROUGE-1, 2.2 ROUGE-2, and 2.55 ROUGE-L points. We have also shown that both the best abstractive and extractive models exceed their corresponding baselines (Longformer Ext and BART Abs) in all metrics. Our models still have room to grow as shown by the oracle results. The limitation of our work is that the dataset is small<sup>8</sup> and investigation on larger datasets would be necessary to further validate our conclusions.

### 4.2 Human Evaluation

For human evaluation, we use the lightweight Pyramid (Shapira et al., 2019). We randomly selected 99 samples<sup>9</sup> from the test dataset for human evaluation. We also re-run Ladhak et al. (2020)’s output’s using the same samples in order to compare ours with their work.

| Model                            | Pyramid |
|----------------------------------|---------|
| CB Const R-wtd (Ladhak, 2020)    | 17.91   |
| BART Abs                         | 22.03   |
| BART Abs+Spinal+Rank+Re-ordering | 22.86   |

Table 3: Pyramid score for our best abstractive performance model compared to previous works

Table 3 shows that our models outperform previous work by at least 2 points. We also show that the application of spinal tree enrichment, ranking loss and re-ordering show an improvement of 0.83 points in the human evaluation.

## 5 Conclusion and Future Work

We have built a novel chapter summarization that produces abstract summaries using a spinal tree aware sub-sentential content selection method. Our results show that we have improved over the state-of-the-art of an existing novel chapter dataset in both automatic and human evaluations.

For future work, we propose an approach where the segmentation of sub-sentential units is jointly trained with the content selection instead of pre-processed before the training process. We hypothesize that this could improve the alignment with reference summaries, therefore, increasing the performance of the overall models.

<sup>8</sup>It is difficult to show significance using a small dataset.

<sup>9</sup>We prepared 100 samples, but one sample got corrupted during the evaluation.

## 6 Ethical Impact

We don't foresee any ethical issues with our approach. One could argue that our system might ultimately take jobs away from the people who currently write such summaries. However, given the number of books being written, it is more likely that some summaries would never be written and a good system for novel chapter summarization might help to increase the amount of summaries that are available online.

## References

- Miguel Ballesteros and Xavier Carreras. 2015. Transition-based spinal parsing. In *Proceedings of the 19th Conference on Computational Language Learning (CoNLL 2015)*. 2015 July 30-31; Beijing, China.[Stroudsburg]: ACL, 2015. p. 289-99. [repository.upf.edu](https://www.repository.upf.edu).
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. *Longformer: The Long-Document transformer*.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.
- Xavier Carreras, Michael Collins, and Terry Koo. 2008. Tag, dynamic programming, and the perceptron for efficient, feature-rich parsing. In *CoNLL 2008: Proceedings of the Twelfth Conference on Computational Natural Language Learning*, pages 9–16.
- Yen-Chun Chen and Mohit Bansal. 2018. *Fast abstractive summarization with Reinforce-Selected sentence rewriting*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734.
- Michael Collins. 1997. Three generative, lexicalised models for statistical parsing. *arXiv preprint [cmp-lg/9706022](https://arxiv.org/abs/cmp-lg/9706022)*.
- Ricardo Cruz, Kelwin Fernandes, Jaime S Cardoso, and Joaquim F Pinto Costa. 2016. Tackling class imbalance with ranking. In *2016 International joint conference on neural networks (IJCNN)*, pages 2182–2187. IEEE.
- 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. 338  
339  
340
- Irving John Good. 1992. Rational decisions. In *Breakthroughs in statistics*, pages 365–377. Springer. 341  
342
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. *Newsroom: A Dataset of 1.3 Million Summaries with Diverse Extractive Strategies*. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 708–719, Stroudsburg, PA, USA. Association for Computational Linguistics. 343  
344  
345  
346  
347  
348  
349  
350
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. *Teaching Machines to Read and Comprehend*. In *Neural Information Processing Systems*, pages 1–14. 351  
352  
353  
354  
355
- Wojciech Kryściński, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong, and Dragomir Radev. 2021. *Booksum: A collection of datasets for long-form narrative summarization*. 356  
357  
358  
359
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR. 360  
361  
362  
363
- Faisal Ladhak, Bryan Li, Yaser Al-Onaizan, and Kathleen McKeown. 2020. *Exploring content selection in summarization of novel chapters*. 364  
365  
366
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. *BART: Denoising Sequence-to-Sequence pre-training for natural language generation, translation, and comprehension*. 367  
368  
369  
370  
371  
372
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics. 373  
374  
375  
376
- Yang Liu and Mirella Lapata. 2019. *Text summarization with pretrained encoders*. 377  
378
- Igor Igor Mikhailovich Malioutov. 2006. *Minimum cut model for spoken lecture segmentation*. Ph.D. thesis, Massachusetts Institute of Technology. 379  
380  
381
- 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. 382  
383  
384  
385  
386  
387
- 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* 388  
389  
390  
391

392 and Computational Natural Language Learning  
393 (EMNLP-CoNLL), pages 380–389, Prague, Czech  
394 Republic. Association for Computational Linguistics.

395 Shashi Narayan, Shay B Cohen, and Mirella Lapata.  
396 2018. [Don't give me the details, just the summary!](#)  
397 [Topic-Aware convolutional neural networks for ex-](#)  
398 [treme summarization.](#)

399 Dragomir R Radev, Hongyan Jing, Małgorzata Styś, and  
400 Daniel Tam. 2004. Centroid-based summarization  
401 of multiple documents. [Information Processing &  
402 Management](#), 40(6):919–938.

403 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine  
404 Lee, Sharan Narang, Michael Matena, Yanqi Zhou,  
405 Wei Li, and Peter J Liu. 2019. [Exploring the lim-](#)  
406 [its of transfer learning with a unified Text-to-Text](#)  
407 [transformer.](#)

408 Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto,  
409 Michele Piana, and Alessandro Verri. 2004. Are  
410 loss functions all the same? [Neural computation](#),  
411 16(5):1063–1076.

412 Ori Shapira, David Gabay, Yang Gao, Hadar Ronen,  
413 Ramakanth Pasunuru, Mohit Bansal, Yael Am-  
414 sterdamer, and Ido Dagan. 2019. [Crowdsourcing](#)  
415 [lightweight pyramids for manual summary eval-](#)  
416 [uation.](#) In [Proceedings of the 2019 Conference](#)  
417 [of the North American Chapter of the Association](#)  
418 [for Computational Linguistics: Human Language](#)  
419 [Technologies, Volume 1 \(Long and Short Papers\)](#),  
420 pages 682–687, Minneapolis, Minnesota. Associa-  
421 tion for Computational Linguistics.

422 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob  
423 Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz  
424 Kaiser, and Illia Polosukhin. 2017. Attention is  
425 all you need. In [Advances in neural information](#)  
426 [processing systems](#), pages 5998–6008.

427 Jeff Wu, Long Ouyang, Daniel M Ziegler, Nisan Sti-  
428 ennon, Ryan Lowe, Jan Leike, and Paul Christiano.  
429 2021. Recursively summarizing books with human  
430 feedback. [arXiv preprint arXiv:2109.10862](#).

431 Zongda Wu, Li Lei, Guiling Li, Hui Huang, Chen-  
432 gren Zheng, Enhong Chen, and Guandong Xu. 2017.  
433 [A topic modeling based approach to novel docu-](#)  
434 [ment automatic summarization.](#) [Expert Systems with](#)  
435 [Applications](#), 84:12–23.

436 Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Pe-  
437 ter J Liu. 2019a. [PEGASUS: Pre-training with ex-](#)  
438 [tracted gap-sentences for abstractive summarization.](#)

439 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Wein-  
440 berger, and Yoav Artzi. 2019b. [Bertscore: Eval-](#)  
441 [uating text generation with bert.](#) [arXiv preprint](#)  
442 [arXiv:1904.09675](#).

## A Human Evaluation 443

444 For performing Human Evaluation, we use crowd-  
445 sourcing service provided by Appen.<sup>10</sup> We follow  
446 [Ladhak et al. \(2020\)](#)'s approach including the in-  
447 structions and number of crowdworkers. For each  
448 crowdworker, we calculate the payment based on  
449 the minimum wage in the US (\$ 15 per hour).

<sup>10</sup><https://appen.com/solutions/crowd-management/>