

Improving Human Text Comprehension through Semi-Markov CRF-based Neural Section Title Generation

Sebastian Gehrmann^{1,2}, Steven Layne^{2,3}, Franck Deroncourt²

¹Harvard SEAS, ²Adobe Research, ³University of Illinois Urbana Champaign
gehrmann@seas.harvard.edu, solayne2@illinois.edu,
franck.deroncourt@adobe.com

Abstract

Titles of short sections within long documents support readers by guiding their focus towards relevant passages and by providing anchor-points that help to understand the progression of the document. The positive effects of section titles are even more pronounced when measured on readers with less developed reading abilities, for example in communities with limited labeled text resources. We, therefore, aim to develop techniques to generate section titles in low-resource environments. In particular, we present an extractive pipeline for section title generation by first selecting the most salient sentence and then applying deletion-based compression. Our compression approach is based on a Semi-Markov Conditional Random Field that leverages unsupervised word-representations such as ELMO or BERT, eliminating the need for a complex encoder-decoder architecture. The results show that this approach leads to competitive performance with sequence-to-sequence models with high resources, while strongly outperforming it with low resources. In a human-subjects study across subjects with varying reading abilities, we find that our section titles improve the speed of completing comprehension tasks while retaining similar accuracy.

1 Introduction

Section titles in long documents that explain the content of the section improve the recall of content (Dooling and Lachman, 1971; Smith and Swinney, 1992) while simultaneously increasing the reading speed (Bransford and Johnson, 1972). Additionally, they can provide a context to allow ambiguous words to be understood more easily (Wiley and Rayner, 2000) and to better understand the overall text (Kintsch and Van Dijk, 1978). However, most documents do not include titles for short segments or only provide a very abstract description of their topics, e.g. “Geography”

Lascaux cave complex discovered	When another old cave is discovered in the south of France, it is not usually news. Rather, it is an ordinary event. Such discoveries are so frequent these days that hardly anybody pays heed to them. However, when the <i>Lascaux cave complex</i> was <i>discovered</i> in 1940, the world was amazed. Painted directly on its walls were hundreds of scenes showing how people lived thousands of years ago. The scenes show people hunting animals, such as bison or wild cats. Other images depict birds and, most noticeably, horses, which appear in more than 300 wall images, by far outnumbering all other animals.
Paintings exposed to destructive action	Early artists drawing these animals accomplished a monumental and difficult task. They did not limit themselves to the easily accessible walls but carried their painting materials to spaces that required climbing steep walls or crawling into narrow passages in the Lascaux complex. Unfortunately, the <i>paintings</i> have been <i>exposed to the destructive action</i> of water and temperature changes, which easily wear the images away. Because the Lascaux caves have many entrances, air movement has also damaged the images inside. Although they are not out in the open air, where natural light would have destroyed them long ago, many of the images have deteriorated and are barely recognizable.
Site closed to tourists	To prevent further damage, the <i>site</i> was <i>closed to tourists</i> in 1963, 23 years after it was discovered.

Figure 1: Output example from our model (left) for an out-of-domain text (right).

or “Introduction”. This makes them more inaccessible especially to readers with less developed reading skills, who have trouble identifying relevant information in text (Englert et al., 2009) and therefore more strongly rely on text-markups (Bell and Limber, 2009).

This paper introduces an approach to generate section titles by extractively compressing the most salient sentence of each paragraph, as shown in Figure 1. While there has been much recent work on abstractive headline generation from a single sentence (Nallapati et al., 2016), abstractive models require larger datasets, which are not available in many domains and languages. Moreover, abstractive text-generation models tend to generate incorrect information for complex inputs (See et al., 2017; Wiseman et al., 2017). Misleading headlines can have unintended effects, affecting readers’ memory and reasoning skills and even bias them (Ecker et al., 2014; Chesney et al., 2017). Especially in times of sensationalism and click-baiting, the unguided generation of titles can be considered unethical and we thus focus on the investigation of deletion-only approaches to title-generation. While this restricts this approach to languages that, similar to English, do not

lose grammatical soundness when clauses are removed, this approach is highly data-efficient and preserves the original meaning in most cases.

We approach the problem with a two-part pipeline where we aim to generate a title for each paragraph of a text, as illustrated in Figure 2. First, a SELECTOR selects the most salient sentence within a paragraph and then the COMPRESSOR compresses the sentence. The selector is an extractive summarization algorithm that assigns a score to each sentence corresponding to its likelihood to be used in a summary (Gehrmann et al., 2018). Algorithms for word deletion typically rely on linguistic features within a tree-pruning algorithm that identifies which phrases can be excluded (Filippova and Altun, 2013). Following recent work that shows the efficiency of contextual span-representations (Lee et al., 2016; Peters et al., 2018), we develop an alternative approach based on a Semi-Markov Conditional Random Field (SCRF) (Sarawagi and Cohen, 2005). The SCRF is further extended by a language model that ranks multiple compression candidates to generate grammatically correct compressions.

We evaluate this approach by comparing it to strong sequence-to-sequence baselines on an English sentence-compression dataset and show that our approach performs almost as well on large datasets while outperforming the complex models with limited training data. We further show the results of a human study to compare the effects of showing no section titles, human-generated titles, and titles generated with our method. The results corroborate previous findings in that we find a significant decrease in time required to answer questions about a text and an increase in the length of summaries written by test subjects. We also observe that the extractive algorithmic titles have a stronger effect on question answering tasks, whereas abstractive human titles have a stronger effect on the summarization task. This indicates that the inherent differences in how humans and our approach summarize the content of a section play a major role in how reading comprehension is affected.

2 Methods

2.1 SELECTOR

To select the most important sentence, we adapt an approach to the problem of content-selection in summarization, which has been shown to be ef-




Figure 2: Overview of the three steps. The SELECTOR detects the most salient sentence. Then, the COMPRESSOR generates compressions and the RANKER scores the them.

fective and data-efficient (Gehrmann et al., 2018). An advantage of this approach over other extractive summarizers (e.g. Zhou et al. (2018)) is that those often model dependencies between selected sentences, which is not applicable to this problem since we only aim to extract a single sentence.

Let $x = x_1, \dots, x_n$ denote a sequence of words within a paragraph, and $y = y_1, \dots, y_m$ a multi-sentence summary with $n \gg m$. Further, let $t = t_1, \dots, t_n$ be a binary alignment variable where $t_i=1$ iff $x_i \in y$. Using this alignment, the word-saliency problem is defined as learning a SELECTOR model that maximizes $\log p(t|x) = \sum_{i=1}^n \log p(t_i|x)$. Using this model, we calculate the relevance of a sentence $sent := x_{start}, \dots, x_{end}$, with $1 \leq start < end \leq n$, with a saliency function defined as

$$\text{saliency}(sent) = \frac{1}{|sent|} \sum_{i=1}^{|sent|} p(t_i|sent).$$

The sentence selection problem thus reduces to sentence with the most relevant words within a paragraph $para$,

$$\operatorname{argmax}_{sent \in para} \text{saliency}(sent).$$

We first represent each word using two different embedding channels. The first is a contextual word representation using ELMo (Peters et al., 2018), and the second one uses GloVe (Pennington et al., 2014). Preliminary experiments corroborated the findings by Peters et al. that the combination of the embeddings help the model converge faster, and perform better with limited training data. Both embeddings for a word x_i are concatenated into one vector e_i , and used as input to a bidirectional LSTM (Hochreiter and Schmidhuber, 1997). Finally, the output of the LSTM h_i is used to compute the probability that a word is selected $\sigma(W_s^T h_i + b_s)$ with the trainable parameters W_s and b_s .

2.2 COMPRESSOR

We next define the problem of deletion-only compression of a single sentence. For simplicity of notation, let x_1, \dots, x_n refer to the words within a single sentence, and y_1, \dots, y_n be a binary indicator whether a word is kept in the compressed form. The compression $c(x, y)$ becomes

$$c(x, y) = x_i \forall (x_1, y_1), \dots, (x_n, y_n) \text{ iff } y_i = 1.$$

The challenge here is that choices are not made independently from another. Consider the sentence *The round ball flew into the net last weekend*, which should be compressed to *Ball flew into net*. Here, the choice to include *into* depends on first selecting the corresponding verb *flew*. One approach to this problem is to use an autoregressive sequence-to-sequence model, in which choices are conditioned on preceding ones. However, these models typically require too many training examples for many languages or domains. Therefore, we relax the problem by assuming that it obeys the Markov property of order L , and train a COMPRESSOR model to maximize $p(y|x) = \sum_{i=1}^n \log p(y_i | \mathbf{x}, y_{i-L:i-1})$. To still retain grammaticality, we define the additional problem of estimating the likelihood of a compression $p(y)$ with a RANKER model, described below.

We compare multiple approaches to deletion-based sentence compression. Throughout, we apply the same embedding as for the SELECTOR. Since a grammatical compression problem relies on the underlying linguistic features (Filippova and Altun, 2013), we process the source documents with spaCy (Honnibal and Montani, 2017) to include the following features in addition to contextualized word embeddings: (1) Part-of-speech tags, (2) Syntactic dependency tags, (3) original word shape, (4) Named entities. Each of these information is encoded in an additional embedding that is concatenated to the word embeddings. Over the final embedding vector, we use a bidirectional LSTM to compute the hidden representations h_1, \dots, h_n for this task.

Naive Tagger The simplest approach we consider is a naive tagger similar to the SELECTOR. We assume full independence between values in y . The probability $p(y_i=1|x)$ is computed as $\sigma(W_c^T w_1 + b_c)$ with trainable parameters W_c and b_c .

Conditional Random Field Compressed sentences have to remain grammatical, which implies a dependence between values in y . Without any restrictions on the dependence between choices, it is intractable to marginalize over all possible y with a scoring function for a given (x, y) pair,

$$p(y|x) = \frac{\exp \text{Score}(x, y)}{\sum_{y'} \exp \text{Score}(x, y')}. \quad (1)$$

Therefore, we assume that only neighboring values in y are dependent, and apply a linear-chain CRF (Lafferty et al., 2001) that uses h as its features to this problem.

We define a scoring function $\text{Score}(x, y) = \sum_i \log \phi_i(x, y)$ that computes the (log) potentials at a position i using a function ϕ . The emission potential for a word x_i is computed as $\phi_i^E(x, y) = W_{e2}(\tanh(W_{e1}h_i + b_{e1})) + b_{e2}$, using only the local information. The transition potential ϕ_i^T depends on the previous and current choice (y_{i-1}, y_i) , and can be looked up in a 2×2 matrix that is learned during training. The complete scoring function can be expressed as

$$\text{Score}(x, y) = \sum_i^{|y|} (\phi_i^E(x, y) + \phi_i^T(y_{i-1}, y_i)).$$

During training, we can minimize the negative log-likelihood in which the partition function is computed with the forward-backward algorithm. During inference, this formulation allows for exact inference with the Viterbi algorithm.

Semi-Markov Conditional Random Field Although compressed sentences should often include entire phrases, the CRF does not take into account local dependencies beyond neighboring words. Therefore, we relax the Markov assumption and score longer spans of words. This can be achieved with a SCRf (Sarawagi and Cohen, 2005). Following a similar approach as Ye and Ling (2018), let $s = \{s_1, s_2, \dots, s_p\}$ denote the segmentation of x . Each segment s_i is represented as a tuple $\langle \text{start}, \text{end}, \tilde{y}_i \rangle$, where start and the end denote the indices of the boundaries of the phrase, and \tilde{y} the corresponding label for the entire phrase. To ensure the validity of the representation, we impose the restrictions that $\text{start}_1 = 1$, $\text{end}_p = |x|$, $\text{start}_{i+1} = \text{end}_i + 1$, and $\text{start}_i \leq \text{end}_i$. Additionally, we set a fixed maximum length L . We

extend Eq. 1 to account for a segmentation s instead of individual tags such that

$$p(s|x) = \frac{\exp \text{Score}(s, x)}{\sum_{s'} \exp \text{Score}(s', x)},$$

marginalizing over all possible segmentations. The CRF represents a special case of the SCRF with $L = 1$. To account for the segment-level information, we extend the emission potential function to

$$\phi_i^E(x, \langle \tilde{y}, start, end \rangle) = \sum_{i=start}^{end} W_e^T h'_i, \quad (2)$$

where h'_i is the concatenation of h_i , $h_{start} - h_{end}$, and a span length embedding e_{len} to account for both individual words and global segment information. We also extend ϕ^T to include transitions to and from longer tagged sequences by representing targets as BIEUO tags (Ratinov and Roth, 2009). This formulation allows for similar training by minimizing the negative log-likelihood.

2.3 RANKER

The inference of the CRF and the SCRF require an estimation of the best possible segmentation $s^* = p(s|x)$, which can be computed using the Viterbi algorithm. However, CRFs and SCRFs are typically employed in sequence-level tagging with no inter-segment dependencies. The sentence compression task differs since a resulting compressed sentence should be grammatical. Therefore, we employ a language model (LM) to rank compression candidates based on the likelihood of the compressed sentence $\text{compress}(x, s)$. Namely, we extend the inference target to

$$s_{LM}^* = \underset{s}{\operatorname{argmax}} (p(s|x) + \lambda p(\text{compress}(x, s))),$$

using a weighting parameter λ . Since exact inference for this target is intractable, we approximate this inference by constraining the re-ranking to the K best segmentations according to a K -best Viterbi algorithm.

The RANKER uses the same word embeddings as the COMPRESSOR, and a bidirectional LSTM in order to maximize the probability of a sequence

$\sum_i p(x_i|x_{1:i-1}) + p(x_i|x_{i+1:n})$. Using the hidden representation h_i , we compute a distribution over the vocabulary as $\text{softmax}(W_l h_i + b_l)$. We additionally use the same weights for word embeddings and W_l , which has been shown to improve language model performance (Inan et al., 2016). We prevent affording an advantage to shorter compressions c during the inference by applying length normalization (Wu et al., 2016) with a length penalty α .

2.4 Sequence-to-Sequence Baselines

The most common approach to summarization and sentence compression uses sequence-to-sequence (S2S) models that learn an alignment between source and target sequences (Sutskever et al., 2014; Bahdanau et al., 2014). S2S models are autoregressive and generate one word at a time by maximizing the probability $p(y|x) = \sum_i p(y_i|x, y_{1:i-1})$. Since this condition is stronger than that of CRF-based approaches, we hypothesize that S2S models perform better with unlimited training data. However, since S2S models need to jointly learn the alignment and generate words, they typically perform worse with limited data.

To test this hypothesis, we define two S2S baselines we compare to our models. First, we use a standard S2S model with attention as described by Luong et al. (2015). In contrast to the other approaches, this model is abstractive and has the ability to paraphrase and re-order words. We constrain these abilities in a second S2S approach as described by Filippova et al. (2015). This model is a sequential pointer-network (Vinyals et al., 2015) and can only generate words from the source sentence. Instead of using an attention mechanism to compute which word to copy, the model enforces a monotonically increasing index of copied words to prevent the re-ordering. We compare both against their reported numbers and our own implementation.

3 Data and Experiments

The SELECTOR is trained on the CNN-DM corpus (Hermann et al., 2015; Nallapati et al., 2016), which is the most commonly used corpus for news summarization (Dernoncourt et al., 2018). Each summary comprises a number of bullet points for an article, with an average length of 66 tokens and 4.9 bullet points. The COMPRESSOR is trained on the Google sentence compression dataset (Filip-

pova and Altun, 2013), which comprises 200,000 sentence-headline pairs from news articles. The deletion-only version of the headlines was created by pruning the syntactic tree of the sentence and aligning the words with the headline. The largest comparable corpus Gigaword (Rush et al., 2015) does not include deletion-only headlines.

We limit the vocabulary size to 50,000 words for both corpora. Both SELECTOR and COMPRESSOR use a two-layer bidirectional LSTM with 64 hidden dimensions for each direction, and a word-embedding size of 200. Each linguistic feature is embedded into 30-dimensional space. During training, the dropout probability is set to 0.5 (Srivastava et al., 2014). The model is trained for up to 50 epochs or until the validation loss does not decrease for three consecutive epochs. We additionally halve the learning rate every time the validation loss does not decrease for two epochs. We use Adam (Kingma and Ba, 2014) with AMS-Grad (Reddi et al., 2018), an initial learning rate of 0.003, and a l_2 -penalty weight of 0.001. The RANKER uses the same LSTM configuration, but we optimize it with SGD with 0.9 momentum, and an initial learning rate of 0.25.

The S2S models have 64 hidden dimensions for each direction of the encoder, and 128 dimensions for the decoder LSTM. They use one layer, and the decoder is initialized with the final state of the encoder. Our optimizer for this task is adagrad with an initial learning rate of 0.15, and an accumulator value of 0.1 (Duchi et al., 2011).

3.1 Automated Evaluation

In the automated evaluation, we focus on the compression models and first conduct experiments with the full dataset to compute an upper bound on the performance of our approach. This experiment functions as a benchmark to investigate how much better the S2S based approaches perform with sufficient data. The next experiment investigates a scenario, in which data availability is limited and ranges from 100 to 1000 training examples. We compare results with and without linguistic features to further evaluate whether these features improve the performance or whether contextual embeddings are a sufficient representation. In each experiment, we measure precision, recall, and F1-score of the predictions compared to the human reference, as well as the ROUGE-score. We additionally measure the length of the compressions




Figure 3: Two paragraphs within the interface of our human evaluation with titles in the left-top margin of a paragraph. A side-effect of the data-efficient deletion-only approach, some titles look ungrammatical, as shown in the second example “*Warming is putting more moisture.*”

to investigate whether methods delete a sufficient number of words.

3.2 Human Evaluation

We evaluated the effect of our generated titles in a between-subjects study on Amazon Mechanical Turk. We compared three different conditions: no titles, human-generated titles, and algorithmically generated titles by our SCRF+Ranking model. Every participant kept their randomly assigned condition throughout all tasks. We defined the following three tasks to approximately measure the effect of short section titles on (1) retention of text, (2) comprehension of text, and (3) retrieval of information. **(Retention)** We first presented a text and then asked participants three questions about facts in the text. **(Comprehension)** We showed a text and then asked the participants to generate a three-sentence summary of the text. **(Retrieval)** We first presented two questions and then the text, prompting participants to find the answers.

Previous findings indicate that titles help with retention only when presented towards the beginning of a text (Dooling and Mullet, 1973). Thus, we place texts in the left margin at the top of a paragraph as shown in the example in Figure 3. This further avoids interrupting the reading flow of the long text while being integrated into the natural left-to-right reading process. Although reading comprehension is well studied in natural language processing, most datasets focus on machine comprehension (Richardson et al., 2013; Rajpurkar et al., 2016). Therefore, we adapted texts from the interactive reading practice by National Geographic, written by Helen Stephenson¹. The 33 texts are based on articles and comprise three versions for each story; elementary, intermediate, and advanced, from which we selected intermediate and advanced versions. Topics of the

¹<http://www.ngllife.com/student-zone/interactive-reading-practice>

Model	Features	P	R	F1	Length
Filippova et al. (2015)	Yes	82.0			
S2S w/o copy	No	83.2 ±4.5	73.0 ±5.9	75.8 ±4.1	9.4 ±2.9
Sequential Pointer		87.1 ±4.1	76.0 ±5.1	79.1 ±3.5	9.4 ±2.6
Naive Tagger		81.4 ±3.7	74.6 ±6.0	75.5 ±3.9	9.7 ±2.7
SCRF		85.2 ±3.9	71.6 ±8.6	73.6 ±5.5	9.0 ±3.5
SCRF+Ranking		86.1 ±3.9	72.9 ±8.5	74.8 ±5.5	9.1 ±3.4
S2S w/o copy	Yes	84.6 ±4.1	75.1 ±5.6	77.6 ±3.8	9.5 ±3.1
Sequential Pointer		89.6 ±3.5	74.2 ±5.2	79.8 ±3.6	8.7 ±2.3
Naive Tagger		84.1 ±3.5	75.4 ±5.3	79.5 ±3.7	9.5 ±2.5
SCRF		86.3 ±3.7	73.0 ±8.4	79.1 ±3.5	9.3 ±2.7
SCRF+Ranking		87.2 ±3.7	73.9 ±7.8	79.6 ±3.3	9.1 ±2.8

Table 1: Results of our models on the large dataset comprising 200,000 compression examples.

texts include Geography, Science, Anthropology, and History; their length ranges from four to seven paragraphs. Each text is accompanied by reading comprehension questions, which we utilized in the retention and retrieval tasks. We first excluded those questions where the answer was part of either human- or algorithmically generated summary. Of the remaining questions, we randomly selected three questions for each of the retention and retrieval tasks. The same questions were shown in either of the conditions.

Every participant completed six tasks, two for every possible task, one with intermediate and one with advanced difficulty. To account for the different backgrounds of participants, we also asked participants about their perceived difficulty for each task on a 5-point Likert scale. The total time to complete all tasks was limited to 30 minutes, and Turkers were paid \$5. In total, we recruited 144 participants who self-reported that they fluently spoke English, uniformly distributed over the three conditions. They answered on average 68.25% of questions correctly and took 16.5 minutes to complete all six tasks. This is approximately 30% faster than the fastest graduate student we recruited for pilot-testing, indicating that Turkers aimed to complete the tasks as fast as possible, possibly by only skimming the text. We omitted results from participants with an answer accuracy of less than 25% ($n=21$), and excluded individual replies given in under 15 seconds ($n=10$) or over 10 minutes ($n=5$), leaving a total of 701 completed tasks. After excluding outliers, the correct answer average was 75.64%, while the time to completion increased by 15 seconds to 16.75 minutes.




Figure 4: Index of extraction within a paragraph.

4 Results

Selector We compare the performance of the selector against the LEAD-1 baseline that naively selects the first sentence of a news article. This provides a strong comparison since a news article typically aims to summarize its content in the first sentence. LEAD-1 achieves ROUGE (1/2/L)-scores of 27.5/9.6/23.7 respectively. In contrast, our selector achieves scores of 30.2/12.2/26.45 which presents an improvement of over 10% in each category. We illustrate the source of this improvement in Figure 4, which shows the locations of selected sentences and observe a negative correlation between a later location within a text and the probability of being selected. However, in most cases, the first sentence is not the most relevant according to the model.

Unrestricted Data Table 1 shows the results of the different approaches on the large dataset. As expected, the copy-model performs best due to its larger modeling potential. It is closely followed by SCRF+Ranking, which comes within 0.2 F1-score when using additional features. This difference is not statistically significant, given the high variance of the results. Compared to the best reported re-




Figure 5: F1-scores of the different models with an increasing number of training examples.

sult in the literature by [Filippova et al. \(2015\)](#), our models perform almost as well, despite the fact that their model is trained on 2,000,000 unreleased datapoints compared to our 200,000. We further observe that all models generate compressed sentences of almost the same length between 8.7 and 9.5 tokens per compression.

The Naive Tagger also achieves comparable performance to the SCRF in F1-Score. To test whether our model leads to a higher fluency compared to it, we additionally measure the ROUGE score. In ROUGE-2, we find that the SCRF+Ranking leads to an increase from 58.1 to 60.1, with an increase in bigram precision by 5 points from 64.5 to 69.3. The Naive Tagger is more efficient at identifying individual words, with a ROUGE-1 of 71.3 when the ranking approach only achieves 68.7. While these differences lead to similar ROUGE-L scores between 69.9 and 70.2, the fact that the ranking-based approach matches longer sequences indicates higher fluency. In an analysis of samples from the Naive Tagger, we found that it commonly omits crucial verb-phrases from compressed sentences.

We show compressions from two paragraphs in the NatGeo data in Figure 3. This example illustrates the robustness of the compression approach to out-of-domain data when including linguistic features. Despite the fact that the example text is not a news article like the training data, it performs well and generates mostly grammatical compressions.

Limited Data We present results on limited data in Figure 5. The results show the major advantage of the simpler training objective. All of the tagging-based models outperform the S2S baselines by a large margin due to their data-efficiency. We did not observe a significant difference between the different tagging approaches in the limited data condition. In our experiments, we found that the S2S models start outperforming the sim-

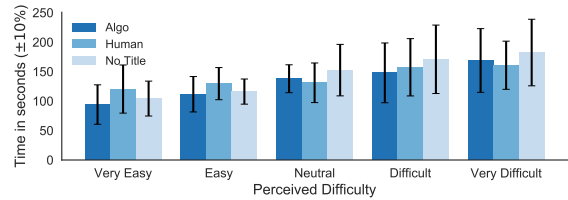


Figure 6: Mean and the 90% confidence interval of the time taken by Turkers to complete the tasks, grouped by the perceived difficulty.

pler models at around 20,000 training examples. Despite its high F1-Score, the Naive Tagger suffers from ungrammatical output in this condition as well, with the readability scoring significantly lower than the SCRF outputs. We argue that the SCRF+Ranking approach represents the best trade-off of our presented models since it performs well with limited data while performing almost as well as complex models in unlimited data situations. This makes it most flexible to apply to a wide range of tasks.

Human Evaluation In the human study, we notice an immediate effect of the difficulty of texts. Between the intermediate and advanced versions of the texts, the mean time to complete the tasks increases by 8 seconds (from 125 to 133 seconds). Additionally, the mean perceived difficulty increases from 2.40 to 2.55 (in between Easy and Neutral). The largest observed effect of text difficulty is on the accuracy of answers, which decreases from 84.4% to 67.5%, indicating that within a similar time frame, the difficult texts were harder to understand.

We present a breakdown of time spent on a task by perceived difficulty for each of the conditions in Figure 6. There is a positive correlation between perceived difficulty and time spent on a task across all conditions. Interestingly, tasks rated Very Easy and Easy were completed slower in the human condition than in the no-title condition, but faster with the algorithmically generated ones. This effect alleviates in the higher difficulties, in which the no-title condition takes longest. Indeed, a comparison between the algorithmic and no title conditions reveals a decrease in time by 19.8 seconds in the retention task, significant with $p < 0.005$ according to a χ^2 test. Interestingly, we can observe opposite effects in the human condition. Here, the comprehension task is completed 15.4 seconds faster ($p < 0.005$), but the other tasks only show minor effects.

Task	Measure	Intervention	Effect Size	p-value
Retention	Time Taken (sec)	Human	-2.2	0.63
		Algo	-27.1	0.01*
	Accuracy	Human	-0.01	0.07
		Algo	-0.01	0.13
Retrieval	Time Taken (sec)	Human	-0.9	0.87
		Algo	-4.5	0.03*
	Accuracy	Human	+0.01	0.20
		Algo	-0.01	0.15
Comprehension	Time Taken (sec)	Human	-20.9	0.03*
		Algo	-2.6	0.04*
	Summary Length (words)	Human	+8.6	0.02*
		Algo	+5.3	0.03*
	Readability	Human	-0.1	0.24
		Algo	-0.1	0.50
	Relevance	Human	-0.02	0.92
		Algo	+0.01	0.63

Table 2: The causal effects of the human and algorithmic section titles on different measures differ across tasks. All the shown effect sizes are measured in comparison to the baseline without any shown titles. Significant p-values at a 0.05 level are marked with a *.

To further investigate these effects, we analyze the causal effect of our three conditions by measuring the average treatment effect while controlling for both actual and perceived difficulty of the tasks with an Ordinary Least Squares analysis. Whenever possible, we additionally condition on the total time taken. An overview of our tests is presented in Table 2. In the causal tests, we observe similar effects in the retention task – the algorithmically generated titles lead to a decrease in time required for the task with an effect size of 27.1 seconds. In contrast, human-generated titles only lead to a non-significant 2.2-second decrease. We observe a non-significant decrease of approximately 4% in accuracy with added titles. We observed no effect of adding titles on the perceived difficulty of this task. In the retrieval task, added titles result in a weaker effect. The algorithmic titles decrease time by only 4.5 seconds, and the human titles by non-significant 0.9 seconds. Similar to the retention task, there is no significant change in accuracy, with all accuracy levels within 1% from another and we observe no effect on the perceived difficulty.

In the comprehension task, it is not possible to measure accuracy. Instead, we evaluate readability and relevance as judged by human raters on a five-point Likert scale, two commonly used metrics for abstractive summarization (Paulus et al.,

Condition	Readability		Relevance	
	μ	σ	μ	σ
No title	4.66	0.65	4.11	0.86
Human	4.55	0.76	4.09	0.95
Algo	4.52	0.72	4.12	1.02

Table 3: Human ratings for human-generated summaries while showing different section titles.

2017). We present the average ratings in Table 3 and observe that there is almost no difference in relevance and only a minor (not significant) decrease in readability with either condition. Curiously, the previous effect on speed reverses in this task – algorithmic titles only lead to a 2.6-second decrease in time, while human titles lead to a 20.9-second decrease. Both conditions additionally lead to longer summaries; algorithmic titles by 5.3 words and human titles by 8.6 words. One potential explanation for this behavior could be that subjects copied the presented section titles into the summary text field. This was not the case, since, on average, only 2.8-3.7% of the bigrams in the titles were used in the summaries, across both conditions and difficulties (0.6-1.5% of trigrams, 0.1-0.8% of 4-grams).

Given the similar relevance scores, we thus argue that presenting the titles leads to more

detailed descriptions of the texts. Similarly to the other tasks, the perceived difficulty does not change significantly. However, we note that some subjects noted that there were some ungrammatical generated titles, which is an artifact of the deletion-only approach. Future work may investigate how abstractive approaches that are not restricted to deletion-based approaches can be applied to the same problem.

Overall, the results of the human-subject study reveal an effect that is well studied in the literature. Namely, that the type of title influences what is being remembered about a text (Schallert, 1975), and that different headline styles affect readers in different ways (Lorch Jr et al., 2011). Kozminsky (1977) found that the immediate free recall of information is biased towards topics emphasized in titles. The better performance in memorization tasks in the algorithmic condition can be explained by the fully extractive approach that immediately shows information judged most relevant by the model. In contrast, human-generated titles show a higher level of abstraction and generalization, which is more helpful for the overall comprehension but does not emphasize any piece of information.

5 Related Work

The aim of SCRFs is to learn a segmentation of a sequential input and assigning the same label to an entire segment. While they were originally developed for information extraction (Sarawagi and Cohen, 2005), it is most commonly applied to speech recognition within the acoustic model to improve segmentation between different words (He and Fosler-Lussier, 2015; Lu et al., 2016; Kong et al., 2015). Similar to this work, it has also been shown that coupling an LM with an SCRF can improve segmentation through multi-task training (Lu et al., 2017; Liu et al., 2017). SCRFs have also been applied to sequence tagging tasks, for example, the extraction of phrases that indicate opinions (Yang and Cardie, 2012). In this work, we built upon an approach by Ye and Ling (2018) who recently introduced a hybrid SCRF that uses both word- and phrase-level information. Alternative approaches for similar tasks are CRFs that estimate pairwise potentials rather than using a fixed transition matrix (Jagannatha and Yu, 2016) or high-order CRFs which outperform SCRFs in some sequence labeling tasks (Cuong et al., 2014).

While this work is the first to apply SCRFs to sentence compression, Grootjen et al. (2018) also use extractive summarization techniques to improve reading comprehension by highlighting relevant sentences. Most similar to our compression approach is Hedge Trimmer (Dorr et al., 2003), which compresses sentences through deletion, but uses an iterative shortening algorithm based on linguistic features. Extending this work, Filipova and Altun (2013) apply a similar approach on linguistic features, but learn weights for the shortening algorithm. Both approaches also do not consider the selection of the sentence to be compressed, unlike our proposed model.

6 Conclusion

In this work, we have presented a novel approach to section title generation that uses an efficient sentence compression model. We demonstrated that our approach performs almost as well as sequence-to-sequence approaches with unlimited training data while outperforming sequence-to-sequence approaches in low-resource domains. A human evaluation showed that our section titles lead to strong improvements across multiple reading comprehension tasks. Future work might investigate end-to-end approaches, or develop alternative approaches that generate titles more similar to how humans write titles.

Acknowledgments

We are grateful for the helpful feedback from the three anonymous reviewers. We additionally thank Anthony Colas and Sean MacAvaney for the multiple rounds of feedback on the ideas presented in this paper.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Kenneth E Bell and John E Limber. 2009. Reading skill, textbook marking, and course performance. *Literacy research and instruction*, 49(1):56–67.
- John D Bransford and Marcia K Johnson. 1972. Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of verbal learning and verbal behavior*, 11(6):717–726.

- Sophie Chesney, Maria Liakata, Massimo Poesio, and Matthew Purver. 2017. Incongruent headlines: Yet another way to mislead your readers. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 56–61.
- Nguyen Viet Cuong, Nan Ye, Wee Sun Lee, and Hai Leong Chieu. 2014. Conditional random field with high-order dependencies for sequence labeling and segmentation. *The Journal of Machine Learning Research*, 15(1):981–1009.
- Franck Dernoncourt, Mohammad Ghassemi, and Walter Chang. 2018. A repository of corpora for summarization. In *Eleventh International Conference on Language Resources and Evaluation (LREC)*.
- D James Dooling and Roy Lachman. 1971. Effects of comprehension on retention of prose. *Journal of experimental psychology*, 88(2):216.
- D James Dooling and Rebecca L Mullet. 1973. Locus of thematic effects in retention of prose. *Journal of Experimental Psychology*, 97(3):404.
- Bonnie Dorr, David Zajic, and Richard Schwartz. 2003. Hedge trimmer: A parse-and-trim approach to headline generation. In *Proceedings of the HLT-NAACL 03 on Text summarization workshop-Volume 5*, pages 1–8. Association for Computational Linguistics.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.
- Ulrich KH Ecker, Stephan Lewandowsky, Ee Pin Chang, and Rekha Pillai. 2014. The effects of subtle misinformation in news headlines. *Journal of experimental psychology: applied*, 20(4):323.
- Carol Sue Englert, Troy V Mariage, Cynthia M Okolo, Rebecca K Shankland, Kathleen D Moxley, Carrie Anna Courtad, Barbara S Jocks-Meier, J Christian O’Brien, Nicole M Martin, and Hsin-Yuan Chen. 2009. The learning-to-learn strategies of adolescent students with disabilities: Highlighting, note taking, planning, and writing expository texts. *Assessment for Effective Intervention*, 34(3):147–161.
- Katja Filippova, Enrique Alfonseca, Carlos A Colmenares, Lukasz Kaiser, and Oriol Vinyals. 2015. Sentence compression by deletion with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 360–368.
- Katja Filippova and Yasemin Altun. 2013. Overcoming the lack of parallel data in sentence compression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1481–1491.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.
- FA Grootjen, GE Kachergis, et al. 2018. Automatic text summarization as a text extraction strategy for effective automated highlighting.
- Yanzhang He and Eric Fosler-Lussier. 2015. Segmental conditional random fields with deep neural networks as acoustic models for first-pass word recognition. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2016. Tying word vectors and word classifiers: A loss framework for language modeling. *arXiv preprint arXiv:1611.01462*.
- Abhyuday N Jagannatha and Hong Yu. 2016. Structured prediction models for rnn based sequence labeling in clinical text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2016, page 856. NIH Public Access.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Teun A van Dijk, Walter Kintsch and TA Van Dijk. 1978. Cognitive psychology and discourse: Recalling and summarizing stories. *Current trends in textlinguistics*, page 61.
- Lingpeng Kong, Chris Dyer, and Noah A Smith. 2015. Segmental recurrent neural networks. *arXiv preprint arXiv:1511.06018*.
- Ely Kozminsky. 1977. Altering comprehension: The effect of biasing titles on text comprehension. *Memory & Cognition*, 5(4):482–490.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *Proceedings of the 18th International Conference on Machine Learning 2001 (ICML 2001)*.
- Kenton Lee, Shimi Salant, Tom Kwiatkowski, Ankur Parikh, Dipanjan Das, and Jonathan Berant. 2016.

- Learning recurrent span representations for extractive question answering. *arXiv preprint arXiv:1611.01436*.
- Liyuan Liu, Jingbo Shang, Frank Xu, Xiang Ren, Huan Gui, Jian Peng, and Jiawei Han. 2017. Empower sequence labeling with task-aware neural language model. *arXiv preprint arXiv:1709.04109*.
- Robert F Lorch Jr, Julie Lemarié, and Russell A Grant. 2011. Three information functions of headings: A test of the sara theory of signaling. *Discourse processes*, 48(3):139–160.
- Liang Lu, Lingpeng Kong, Chris Dyer, and Noah A Smith. 2017. Multitask learning with ctc and segmental crf for speech recognition. *arXiv preprint arXiv:1702.06378*.
- Liang Lu, Lingpeng Kong, Chris Dyer, Noah A Smith, and Steve Renals. 2016. Segmental recurrent neural networks for end-to-end speech recognition. *arXiv preprint arXiv:1603.00223*.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*, pages 147–155. Association for Computational Linguistics.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. 2018. On the convergence of adam and beyond. In *International Conference on Learning Representations*.
- Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 193–203.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389.
- Sunita Sarawagi and William W Cohen. 2005. Semi-markov conditional random fields for information extraction. In *Advances in neural information processing systems*, pages 1185–1192.
- Diane L Schallert. 1975. Improving memory for prose: The relationship between depth of processing and context. *Center for the Study of Reading Technical Report; no. 005*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Edward E Smith and David A Swinney. 1992. The role of schemas in reading text: A real-time examination. *Discourse Processes*, 15(3):303–316.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Jennifer Wiley and Keith Rayner. 2000. Effects of titles on the processing of text and lexically ambiguous words: Evidence from eye movements. *Memory & Cognition*, 28(6):1011–1021.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. Challenges in data-to-document generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2263.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

- Bishan Yang and Claire Cardie. 2012. Extracting opinion expressions with semi-markov conditional random fields. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1335–1345. Association for Computational linguistics.
- Zhi-Xiu Ye and Zhen-Hua Ling. 2018. Hybrid semi-markov crf for neural sequence labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, Melbourne, Australia. ACL.
- Qingyu Zhou, Nan Yang, Furu Wei, Shaohan Huang, Ming Zhou, and Tiejun Zhao. 2018. Neural document summarization by jointly learning to score and select sentences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 654–663.