

# A Large-Scale Multi-Length Headline Corpus for Analyzing Length-Constrained Headline Generation Model Evaluation

Yuta Hitomi<sup>1</sup> Yuya Taguchi<sup>1</sup> Hideaki Tamori<sup>1</sup> Ko Kikuta<sup>\*2</sup> Jiro Nishitoba<sup>2</sup>  
Naoaki Okazaki<sup>3</sup> Kentaro Inui<sup>4</sup> Manabu Okumura<sup>3</sup>

<sup>1</sup> The Asahi Shimbun Company, <sup>2</sup> Retrieva, Inc., <sup>3</sup> Tokyo Institute of Technology

<sup>4</sup> Tohoku University, RIKEN Center Advanced Intelligence Project

{hitomi-y1, taguchi-y2, tamori-h}@asahi.com, kikutakou@gmail.com

jiro.nishitoba@retrieva.jp, okazaki@c.titech.ac.jp

inui@ecei.tohoku.ac.jp, oku@pi.titech.ac.jp

## Abstract

Browsing news articles on multiple devices is now possible. The lengths of news article headlines have precise upper bounds, dictated by the size of the display of the relevant device or interface. Therefore, controlling the length of headlines is essential when applying the task of headline generation to news production. However, because there is no corpus of headlines of multiple lengths for a given article, previous research on controlling output length in headline generation has not discussed whether the system outputs could be adequately evaluated without multiple references of different lengths. In this paper, we introduce two corpora, which are Japanese News Corpus (JNC) and JApAnese MUlti-Length Headline Corpus (JAMUL), to confirm the validity of previous evaluation settings. The JNC provides common supervision data for headline generation. The JAMUL is a large-scale evaluation dataset for headlines of three different lengths composed by professional editors. We report new findings on these corpora; for example, although the longest length reference summary can appropriately evaluate the existing methods controlling output length, this evaluation setting has several problems.

## 1 Introduction

The news media publish newspapers in print form and in electronic form. In the electric form, articles might be read on various types of devices using any application; thus, news media companies have an increasing need to produce multiple headlines for the same news article based on what would be most appropriate and most compelling on an array of devices. All devices and applications used for viewing articles have strict upper bounds regarding the number of characters

---

**Article:** トヨタ自動車は18日、エンジン車だけの車種を2025年ごろまでにゼロにすると発表した。...ハイブリッド車 (HV)やプラグインハイブリッド車 (PHV)、燃料電池車 (FCV)も加えた「電動車」を、すべての車種に設定する。...

On the 18th, Toyota announced that it would set the model of only engine cars to zero by about 2025... They set “electric vehicle” which is Hybrid Vehicle (HV), Plug-in Hybrid Vehicle (PHV), and Fuel Cell Vehicle (FCV) to all models...

---

**Headline for print media:** トヨタ、全車種に電動車25年ごろまでに HVやFCV含め

All Toyota models will contain electric vehicles including HV and FCV by about 2025.

---

**Multi-length headlines for digital media:**

9 chars	全車種に「電動車」
(10char-ref)	“Electric cars” for all models
13 chars	トヨタ、全車種に「電動車」
(13char-ref)	“Electric cars” for all Toyota’s models
	トヨタ、エンジン車だけの車種ゼロへ 2025年ごろ
24 chars	Toyota sets the number of models with only engine cars to zero by about 2025.
(26char-ref)	

Table 1: An example of four headlines for the same article that were created by professional editors. In this example, ‘電動車’(Electric cars) and ‘全’(all) are represented by red letters and are not included in the 24-character headline. These tokens cannot be evaluated by 24-character headlines. The blue tokens are not included in 9- and 13-character headlines. These tokens should not be included in shorter headlines.

allowed because of limitations in the space where the headline appears. The technology of automatic headline generation has the potential to contribute greatly to this domain, and the problems of news headline generation have motivated a wide range of studies (Wang et al., 2018; Chen et al., 2018; Kiyono et al., 2018; Zhou et al., 2018; Cao et al., 2018; Wang et al., 2019).

Table 1 shows sample headlines in three different lengths written by professional editors of a media company for the same news article: The length of the first headline for the digital media is restricted to 10 characters, the second to 13 charac-

---

This work was done at Retrieva, Inc. within Project.

ters, and the third to 26 characters. From a practical perspective, headlines must be generated under a rigid length constraint.

The first study to consider the length of system outputs in the encoder-decoder framework was [Rush et al. \(2015\)](#). This study controlled the length of an output sequence by reducing the score of the end-of-sentence token to  $-\infty$  until the method generated the desired number of words. Subsequently, [Kikuchi et al. \(2016\)](#) and [Fan et al. \(2018\)](#) proposed mechanisms for length control; however, these studies produced summaries of 30, 50, and 75 bytes, and the studies evaluated the summaries by using the reference summaries of a single length (approximately 75 bytes long) in DUC 2004<sup>1</sup>. In addition, [Takase and Okazaki \(2019\)](#) proposed the mechanism for length control and evaluated their method with part of the test set which is consisted by summaries satisfying some length constraints in Annotated English Gigaword corpus (AEG) ([Napolet et al., 2012](#)).

Thus, some questions can be posed: (1) Can previous evaluation settings adequately evaluate system outputs in headline generation task? (2) What type of problem should we solve in this task according to the target length? (3) How well do systems solve the problems? In this study, we present novel corpora to investigate these research questions. The contributions of this study are threefold.

1. We release the **J**apanese **N**ews **C**orpus (JNC)<sup>2</sup>, which includes 1.93 million pairs of headlines and the lead three sentences of Japanese news articles. We expect this corpus to provide common supervision data for headline generation.
2. We build the **J**apanese **M**ulti-**L**ength **H**eadline **C**orpus (JAMUL)<sup>2</sup> for the evaluation of headlines of different lengths. In this novel dataset, each news article is associated with multiple headlines of three different lengths.
3. We report new findings for the JAMUL; for example, although the longer reference seems to be able to evaluate the short system output, we also found a problem with this evaluation setting. Additionally, we clarified that the existing methods could not capture

<sup>1</sup><https://duc.nist.gov/duc2004/>

<sup>2</sup>[https://cl.asahi.com/api\\_data/jnc-jamul-en.html](https://cl.asahi.com/api_data/jnc-jamul-en.html)




Figure 1: Length distributions of headlines in (a) the JNC and (b) the JAMUL.

what must be changed according to the specified length.

## 2 JNC and JAMUL

### 2.1 Headlines Composed by a Media Company

Before describing the JNC and the JAMUL in detail, we explain the process where a Japanese media company composes headlines for a news article. We describe the process of The Asahi Shimbun Company, a Japanese newspaper company. First, reporters write an article and submit it to the editorial department. The editorial department picks up some articles to publish in the newspaper and writes a headline for the article dedicated to print media. We call these headlines *print headlines* or *length-insensitive headlines* hereafter.

In addition to print headlines, digital media editors, who are typically not the same editors for print, pick up the articles they want to distribute on digital media from the articles dedicated to digital media and compose three different headlines. The first headline for the digital signage and audio media has a limit of up to 10 characters. This type of headline is appended to the beginning of a concise summary of the article so that readers can under-

stand the news at first glance. The second type of headline is produced for mobile phones with small LCDs and small areas on the website sidebar (e.g., the access ranking); the upper limit of the number of characters is 13. The third type of headline is produced for desktop computer news websites, and the upper limit of the number of characters is 26. This limit is derived from the layout of the news site. We refer to the three types of headlines as *10char-ref*, *13char-ref*, and *26char-ref* (refer to Table 1 for examples). We collectively call these headlines *length-sensitive headlines*.

Table 1 presents an example of headlines written for an article by the professional editors. We extract the JNC and the JAMUL from the news production process of The Asahi Shimbun Company; therefore, they can be considered representative of contemporary editorial practice.

## 2.2 JNC

The JNC is a collection of 1,932,398 pairs of the three lead sentences<sup>2</sup> of articles and their print headlines published by The Asahi Shimbun Company from 2007 to 2016. Figure 1 (a) depicts the distribution of the headline lengths in the JNC. The headline lengths in the JNC are diverse because almost all headlines are not restricted by length limitations.

The JNC is useful for training headline generation models because it has many training instances. Furthermore, the corpus is suitable for training a model for variable-length headline generation because of the variety of the headline lengths.

## 2.3 JAMUL

The JAMUL is a corpus containing 1,489 full-text news articles and their length-insensitive headlines for print media and length-sensitive headlines of 10 characters, 13 characters, and 26 characters for digital media. We just extract these articles and headlines from the company database. All the articles and headlines were published by The Asahi Shimbun Company between September 2017 and March 2018. The volume of the news articles may be insufficient for training a headline generation model. However, as Figure 1 (b) shows, the JAMUL includes length-sensitive headlines that strictly preserve the length requirements. This novel characteristic of the JAMUL is suitable for the test set to adapt headline generation to business practice. In this paper, we use the first three

System	Reference	Precision	Recall
JAMUL Article	Paper headline	11.26	72.91
	10char-ref	5.42	85.02
	13char-ref	7.10	85.74
	26char-ref	13.26	77.38
AEG Article	AEG headline	17.74	62.03

Table 2: ROUGE-1 precision and recall scores when comparing article and length-insensitive/sensitive headlines.

System	Reference	Precision	Recall
Print headline	10char-ref	27.95	71.28
	13char-ref	37.70	74.48
	26char-ref	61.78	60.56

Table 3: ROUGE-1 precision and recall scores when comparing length-insensitive and -sensitive headlines.

sentences of the article as input because of fitting in the JNC corpus. But we publish the full edition of the article. No overlap of articles between the JNC and the JAMUL was confirmed.

## 2.4 Comparing of Headlines with Article Bodies

What type of operation which includes extractive and abstractive operations did the editors perform to create length-sensitive and length-insensitive headlines in the JAMUL? To clarify this question, we analyzed the proportions of the number of extractive and abstractive operations. Specifically, we reported ROUGE-1 precision and recall scores (Lin, 2004) in Table 2, assuming that articles are “system” summaries and that 10char-ref, 13char-ref, and 26char-ref headlines are “reference” summaries. Notably, we removed blank spaces, which were the most common token in longer headlines. The relatively high recall score indicates that the operations most often required to generate headlines are extractive, and the abstractive operation is about 15–27% of the total. In addition, we explored the proportions of each operation in AEG<sup>3</sup>, which has been used in many studies for English headline generation tasks. From the relatively high recall score, we observed that Japanese headlines tend to be more extractive.

## 2.5 Comparing of Length-sensitive Headlines with Print headlines

How similar are the headlines used for training (length-insensitive) and for evaluation (length-

<sup>3</sup>We obtained this dataset by applying the pre-processing script at <https://github.com/facebookarchive/NAMAS>

**Article:** 米フェイスブック(FB)は1日、2017年7 9月期決算を発表し、モバイル広告の伸びなどで売上高、純利益ともに四半期として過去最高を記録した。...FB上で偽ニュースの拡散防止など、安全確保のための要員を約2万人に倍増させることを明らかにしている。

On the 1st U.S. Facebook (FB) announced financial results from July to September in 2017 and achieved a record quarterly amount of sales and net income thanks to its mobile advertising growth and other factors. ... FB revealed it had doubled the number of personnel engaged in preventing fake news from spreading to about 20,000 to provide safety on FB.

**Headline for print media:** フェイスブック、四半期で最高益 モバイル広告好調  
Facebook achieved a record quarterly profit thanks to its mobile advertising business.

<b>Multi-length headlines for digital media:</b>	
7 chars	米FBが最高益
(10char-ref)	U.S. FB achieved record profit.
12 chars	米フェイスブックが最高益
(13char-ref)	U.S. Facebook achieved record profit. 米フェイスブックが最高益 偽ニュース対策で要員倍増も
26 chars (26char-ref)	U.S. Facebook achieved record profit and doubled its personnel for counter-measures against fake news.

Table 4: A typical example for comparing length-insensitive and -sensitive headlines

sensitive)? We estimated the appropriateness of length-insensitive headlines as a “seed” for producing length-sensitive headlines. More concretely, we report ROUGE-1 precision and recall scores in Table 3, assuming that length-insensitive headlines are “system” summaries, and that 10char-ref, 13char-ref and 26char-ref headlines are “reference” summaries. The relatively high recall scores indicate that the training and evaluation data were not very distant. Additionally, we found that the editors used a moderate number of words that did not appear in print headlines when composing length-sensitive headlines. Table 4 is an example of the typical differences between the length-insensitive and length-sensitive headlines. Comparing the 26-character headline with the print headline, the choices of contents are different; for example, the print headline reports the reason for the record profit, but the 26-character headline describes the topic regarding the increasing number of personnel. Next, comparing the 7-character (10char-ref) headline with the print headline, we observe that the choices of words are different; the print headline uses “Facebook,” which is changed to “FB” in the 7-character headline.

System	Reference	P	R
26char-ref	10char-ref	30.22	81.77
26char-ref	13char-ref	43.59	90.75
First 10 chars in 26char-ref	10char-ref	44.30	29.42
First 13 chars in 26char-ref	13char-ref	67.01	44.34

Table 5: Difference between 26char-ref headlines and shorter headlines. P and R denote ROUGE-1 precision and recall scores, respectively.

## 2.6 Comparing of Length-sensitive Headlines

How similar is the composition of headlines for a news article of different lengths? How good are 26char-ref headlines as “seeds” for generating 10char-ref or 13char-ref headlines? Is the simple strategy of trimming 26char-ref headlines to 10 or 13 characters sufficient? To answer these questions, we computed word-level precision and recall scores, assuming that 26char-ref headlines are “system” summaries and that 10char-ref and 13char-ref headlines are “reference” summaries.

The first and second rows of Table 5 represent the situation when we used 26char-ref headlines as they are, and without preserving the length constraint. Although this setting was unrealistic, we could estimate how similar the contents were among the length-sensitive headlines. The high recall scores indicate that 26char-ref headlines mostly cover the words included in the 10char-ref and 13char-ref headlines. From this result, we consider the contents among length-sensitive headlines are comparably similar. The third and fourth rows of Table 5 correspond to the strategy where we generated headlines in 10 and 13 characters from the first 10 and 13 characters of the 26char-ref headlines. This strategy achieved moderate success for generating headlines in 13 characters, but did not work well for headlines in 10 characters. In other words, we observed large differences between the 10char-ref and the first 10 characters of the 26char-ref headlines. Therefore, we need to re-generate the shorter headline, as well as the longest one.

In sum, we found similarities in headlines of different lengths in the JAMUL. However, the simple strategy of trimming a longer headline into a shorter headline is insufficient. Table 1 is an example of the typical differences among length-sensitive headlines. There is a little overlap between longer and shorter headlines because the 9- and 13-character headlines extract the shorter phrases which have the nearly same meaning as

the 24-character headline. Focusing on “車種” (models), the words are in the latter half of the 24-character headline, and we could confirm that important keywords are not always included at the beginning of the headlines.

### 3 Comparing of Headline Generation Methods on the JAMUL

In this section, we explore a question about evaluation: How reliable is the two conventional evaluation method; the first is the method that uses a single length summary for measuring the quality of summaries of different lengths; the second is the method that uses the specified length headlines which are extracted from the dataset constructed by length-insensitive headlines, for measuring the quality of each length summary. To answer this question, we generated multiple summaries of different lengths by using the existing methods, and measured the correlation between the performance values computed by the conventional evaluation methods and those computed on the JAMUL.

#### 3.1 Headline-generation Methods with the Mechanism to Control the Output Length

In this study, we explored four methods for headline generation that can control the output length. The first two methods, *LenEmb* and *LenInit*, were proposed by Kikuchi et al. (2016).

*LenEmb* provides the decoder with output length information in the form of the length embedding. *LenInit* controls the output length by multiplying the initial state of the decoder’s memory cell of long short-term memory (Hochreiter and Schmidhuber, 1997) by the desired length.

Fan et al. (2018) also proposed a length-controllable method for a convolutional sequence-to-sequence (ConvS2S) model (Gehring et al., 2017). Their method added special tokens indicating the range of the output length at the beginning of an input sequence. In our experiment, we used a special token to specify an output length<sup>4</sup> and called this method *SP-token*.

We also considered the *LC* method (Liu et al., 2018), which extends ConvS2S and multiplies the initial state of the residual connection (He et al., 2016) by the desired number of output tokens. In the experiment, we set the desired number of characters instead of that of tokens.

<sup>4</sup>Fan et al. (2018) also included special tokens for entities, but we did not use them in the experiments.

	Seq2Seq	ConvS2S	Transformer
Num of Layer	2	8	6
Dropout Rate	0.3	0.1	0.3
Grad Clipping	[-5.0, 5.0]	[-0.1, 0.1]	-
Learning Rate	0.001	0.2	0.001
Optimizer	Adam	NAG	Adam

Table 6: Parameters of each encoder-decoder model.

In addition to these four methods, we combined *SP-token* not only on ConvS2S but also on Seq2Seq and Transformer (Vaswani et al., 2017). Eventually, we examined six combinations in total: (1) Seq2Seq + *LenEmb*, (2) Seq2Seq + *LenInit*, (3) Seq2Seq + *SP-token*, (4) ConvS2S + *SP-token*, (5) ConvS2S + *LC*, and (6) Transformer + *SP-token*.

#### 3.2 Datasets and the Evaluation Protocol

We trained the six methods for headline generation on the JNC. We removed instances that were duplicated or unsuitable for training a headline generation model<sup>5</sup>. The filtering step obtained 1,554,558 pairs of newspaper articles and headlines. We randomly selected 98% of the instances (1,523,468 pairs) as a training set, selected 1% of the instances (15,545 pairs) as a validation set, and used the remainder (15,545 pairs) as a test set. As for the JNC test set, we additionally extracted the length-restricted test set to evaluate the length control methods. Specifically, we extracted 303, 1,436, and 4,289 headlines which are 8-10, 11-13, and 24-26 characters long, respectively, and corresponding articles from the JNC, similar to the 10char-ref, 13char-ref, and 26char-ref in the JAMUL. We call them the *length-restricted JNC test set*. We also filtered the JAMUL by setting lower bounds for the length of the headlines<sup>6</sup>. We set each lower bound to 8 for 10char-ref, to 11 for 13char-ref, and to 24 for 26char-ref. Finally, we achieved 1,288 instances for the JAMUL test set.

We used SentencePiece<sup>6</sup> (Kudo and Richardson, 2018) for tokenization. We set the merge operation to 8,000. Finally, we obtained 10,868 tokens for the source side and 9,556 tokens for the target side. When training a model, we set the length of each reference headline to the model. When generating headlines in the evaluation, we set the output lengths to 10, 13, and 26 char-

<sup>5</sup>The filtering script is available at: <https://github.com/asahi-research/Gingo>

<sup>6</sup><https://github.com/google/sentencepiece>

Models	10 characters			13 characters			26 characters		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
(1) Seq2Seq + <i>LenEmb</i>	34.62	15.02	33.54	39.31	18.14	37.07	43.65	19.73	36.03
(2) Seq2Seq + <i>LenInit</i>	36.14	16.88	35.15	41.39	19.55	38.83	46.30	21.33	37.79
(3) Seq2Seq + <i>SP-token</i>	38.01	17.16	36.72	42.03	19.83	39.56	46.62	21.47	38.29
(4) ConvS2S + <i>SP-token</i>	39.20	18.66	38.07	42.63	20.00	40.38	47.41	21.47	38.04
(5) ConvS2S + <i>LC</i>	34.93	15.52	34.02	38.80	17.21	36.68	42.59	19.34	35.41
(6) Transformer + <i>SP-token</i>	41.21	19.16	39.91	44.84	21.62	42.09	49.66	23.69	40.65

Table 7: ROUGE scores of each model on the JAMUL. The specified lengths are 10, 13, and 26 characters. R-1, R-2, and R-L represent ROUGE-1, ROUGE-2, and ROUGE-L, respectively. Note that (1) to (6) in Table 8 and after that represent models (1) to (6) of this table.

Models	10 characters			13 characters		
	R-1	R-2	R-L	R-1	R-2	R-L
(1)	21.58 <sub>[+0]</sub>	8.73 <sub>[+1]</sub>	19.56 <sub>[+0]</sub>	27.47 <sub>[-1]</sub>	11.82 <sub>[+0]</sub>	24.20 <sub>[-1]</sub>
(2)	22.17 <sub>[+0]</sub>	9.14 <sub>[+1]</sub>	20.12 <sub>[+0]</sub>	28.93 <sub>[+1]</sub>	12.98 <sub>[+2]</sub>	25.35 <sub>[+1]</sub>
(3)	22.80 <sub>[+1]</sub>	9.51 <sub>[+1]</sub>	20.59 <sub>[+1]</sub>	28.78 <sub>[-1]</sub>	12.64 <sub>[-1]</sub>	25.14 <sub>[-1]</sub>
(4)	22.53 <sub>[-1]</sub>	9.04 <sub>[-2]</sub>	20.47 <sub>[-1]</sub>	29.23 <sub>[+0]</sub>	12.82 <sub>[-1]</sub>	25.58 <sub>[+0]</sub>
(5)	21.60 <sub>[+0]</sub>	8.55 <sub>[-1]</sub>	19.63 <sub>[+0]</sub>	27.63 <sub>[+1]</sub>	11.72 <sub>[+0]</sub>	24.34 <sub>[+1]</sub>
(6)	24.36 <sub>[+0]</sub>	10.32 <sub>[+0]</sub>	21.97 <sub>[+0]</sub>	30.84 <sub>[+0]</sub>	13.85 <sub>[+0]</sub>	26.98 <sub>[+0]</sub>
$\tau$	0.867	0.600	0.867	0.733	0.733	0.733

Table 8: ROUGE scores of the system outputs in 10 and 13 characters evaluated by 26char-ref headlines as the references. Note that [] denote the change of the rank when we compare the rank in Table 7; + and - denote the rank up and rank down, respectively.

acters; each output was evaluated by the reference that had the same length in the JAMUL and JNC test sets. To generate summaries, we follow standard practice in disallowing repetition of the same trigram (Paulus et al., 2018). We evaluated all models by using three variants of ROUGE recall metric<sup>7</sup>: ROUGE-1, ROUGE-2, and ROUGE-L. Headlines exceeding the length limits were trimmed for a fair evaluation.

### 3.3 Implementation

We employed OpenNMT<sup>8</sup> (Klein et al., 2017) for Seq2Seq, and fairseq<sup>9</sup> (Ott et al., 2019) for ConvS2S and Transformer. We extended the implementations to realize *LenEmb*, *LenInit*, and *LC*. We set the dimensions for the token and length embeddings to 512, those for hidden states to 512, and the beam width to 5. These parameters are common in all the models. Table 6 summarizes other parameters specific to each sequence-to-sequence model. We used Nesterov’s accelerated gradient method (NAG) (Sutskever et al., 2013) with a momentum of 0.99 in ConvS2S. In Transformer, we set the number of attention heads to 8, the dimensions for the feed-forward network to 2,048, Adam’s  $\beta$  to 0.98, the warm-up steps to

<sup>7</sup>We used MeCab (Kudo et al., 2004) to tokenize the system outputs.

<sup>8</sup><https://github.com/OpenNMT/OpenNMT-py>

<sup>9</sup><https://github.com/pytorch/fairseq>

4,000, and label smoothing to 0.1.

### 3.4 Evaluation of Multi-length Headlines Generated by Methods on the JAMUL

Table 7 presents ROUGE scores of each method on the JAMUL test set. Transformer + *SP-token* was the clear winner in all length and evaluation metrics. Additionally, the three methods with *SP-token* outperformed the others.

What if we do not have multiple headlines of different lengths to evaluate the methods? To answer this question, we followed the evaluation setup of the previous studies on DUC 2004: The reference summaries of 75 bytes were used even when evaluating summaries of 30 and 50 bytes. Table 8 reports ROUGE scores for the system outputs in 10 and 13 characters evaluated based on the 26char-ref headlines. This evaluation setup reduced the performance differences between the methods. Although Transformer + *SP-token* remained the clear winner, the ranking in ROUGE scores of the other methods are flipped. We also computed rank correlation coefficients (Kendall’s  $\tau$ ) to assess the discrepancy in the ranking among the methods presented in Table 7 and Table 8. The last row of Table 8 reveals that the rank correlation is not perfect (lower than one) but moderate. We understand that  $\tau$  is maintained high to some extent because of two reasons: (1) Most of the

Models	10 characters			13 characters			26 characters		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
(1)	37.95 <sub>[+0]</sub>	19.97 <sub>[+1]</sub>	36.54 <sub>[+0]</sub>	41.54 <sub>[+1]</sub>	21.60 <sub>[+1]</sub>	39.41 <sub>[+1]</sub>	33.97 <sub>[+0]</sub>	13.54 <sub>[-1]</sub>	27.66 <sub>[+0]</sub>
(2)	38.35 <sub>[-1]</sub>	20.71 <sub>[+1]</sub>	37.04 <sub>[-1]</sub>	41.52 <sub>[-1]</sub>	21.13 <sub>[-1]</sub>	39.05 <sub>[-1]</sub>	35.50 <sub>[+0]</sub>	14.18 <sub>[+1]</sub>	28.64 <sub>[+0]</sub>
(3)	41.35 <sub>[+1]</sub>	21.95 <sub>[+1]</sub>	39.93 <sub>[+1]</sub>	43.62 <sub>[+1]</sub>	22.51 <sub>[+0]</sub>	40.89 <sub>[+1]</sub>	35.66 <sub>[+0]</sub>	14.14 <sub>[-1]</sub>	28.80 <sub>[+0]</sub>
(4)	39.33 <sub>[-1]</sub>	20.47 <sub>[-2]</sub>	37.46 <sub>[-2]</sub>	40.38 <sub>[-4]</sub>	20.47 <sub>[-4]</sub>	37.53 <sub>[-4]</sub>	36.84 <sub>[+0]</sub>	14.42 <sub>[+0]</sub>	29.21 <sub>[+0]</sub>
(5)	39.21 <sub>[+2]</sub>	19.59 <sub>[-1]</sub>	37.69 <sub>[+2]</sub>	42.99 <sub>[+3]</sub>	22.94 <sub>[+4]</sub>	40.53 <sub>[+3]</sub>	33.36 <sub>[+0]</sub>	13.55 <sub>[+1]</sub>	27.59 <sub>[+0]</sub>
(6)	45.14 <sub>[+0]</sub>	26.35 <sub>[+0]</sub>	43.70 <sub>[+0]</sub>	47.09 <sub>[+0]</sub>	25.80 <sub>[+0]</sub>	43.95 <sub>[+0]</sub>	37.86 <sub>[+0]</sub>	15.68 <sub>[+0]</sub>	30.46 <sub>[+0]</sub>
$\tau$	0.733	0.600	0.600	0.067	-0.067	0.067	1.000	0.733	1.000

Table 9: ROUGE scores of each model on the length-restricted JNC test set.

swaps occur in adjacent ranks, and (2) there are not many samples of the method in this analysis. Then, we focus on the samples flipped in this evaluation setup, as shown in Table 8. We observed that the rankings are swapped even if there is a difference of 1.0 or more (these differences are well observed in state-of-the-art competition) in Table 7. For example, we can show this issue between (3) and (4) in 10 characters of R-1 score, (2) and (4) in 10 characters of R-2 score, and (3) and (4) in 10 characters setting of R-L score in Table 7.

What if we do not have strict length headlines to evaluate the methods? To answer this question, we followed the evaluation setup of previous studies : The reference summaries of each target length were extracted from the JNC test set. Table 9 reports ROUGE scores for the system outputs in 10, 13, and 26 characters evaluated in the length-restricted JNC test set. In this evaluation setup, Transformer + *SP-token* remained the clear winner, but the performances of the other methods were flipped. We also calculated  $\tau$  in this case. The last row of Table 9 reveals that the 13 character setting considerably lost correlation with the JAMUL evaluation result. We guess this inconsistency is brought about by using the different content test sets among each length setting. Depending on the target length which is specified by the evaluation setting, this evaluation protocol might not evaluate the methods adequately.

## 4 Analysis

### 4.1 Performance of Word Selection According to the Output Length

How well do the existing methods change the word selection depending on the output length? As shown in the first and second rows of Table 5, the 10char-ref and 13char-ref headlines contain words that are not included in the 26char-ref headlines. In other words, the selection of words in the generated headline should be changed in response

to the length restriction. To confirm this question, we computed ROUGE-1 recall scores for the system outputs generated by each method, assuming that the groups of the words included in the 10char-ref or 13char-ref but not in the 26char-ref headlines are the “reference” summaries. For instance, the red words in Table 1 are the “reference” summaries in this experiment.

We report this result in Figure 2. The low recall score indicates that each system cannot select the words tailored to the length constraints. However, (4) ConvS2S + *SP-token* shows the highest performance. Taking this into account, Transformer + *SP-token* improves the important words in all lengths setting, but ConvS2S + *SP-token* may improve the change of words according to the target length.

### 4.2 Performance of Managing Extractive and Abstractive Tasks

In Table 2, we reported the proportion of the number of extractive and abstractive operations in the JAMUL. We analyze how the existing methods can reflect extractive and abstractive operations in generating summaries.

First, to observe extractive operations, we computed ROUGE-1 recall scores for the system outputs generated by each system, measuring the number of overlapping words between an article as “system” summaries and 10char-ref, 13char-ref, and 26char-ref headlines as “reference” summaries. The group of bars which contained *ext* in its name in Figure 3 reports the result. The relatively high recall score indicates that the length control method succeeds in managing extractive operations.

Next, we examine whether the length control methods could perform abstractive operations. We adopted the words included in 10char-ref, 13char-ref, or 26char-ref headlines but not included in an article as “reference” summaries, and computed the ROUGE-1 recall scores for the system outputs.




Figure 2: ROUGE-1 recall scores when comparing system outputs and the groups of the words included in 10char-ref or 13char-ref but not included in 26char-ref.




Figure 3: ROUGE-1 recall scores when comparing system outputs and the groups of the words which are extractive and abstractive.

The group of bars which contained *abs* in its name in Figure 3 reports the result. Regarding the outputs targeting 26 characters, the recall scores of around 30 point imply that each model can manage abstractive operations to some extent. In contrast, the low recall scores for the outputs targeting 10 and 13 characters revealed that all length control methods could not perform well on abstractive operations under the severe length constraint.

### 4.3 How Do Length Control Mechanisms Work?

We wondered whether a method that could control the output length would produce similar headlines even for different lengths for the same news article. To confirm this suspicion, we reported ROUGE-1 recall scores in Figure 4 with three different configurations: (a) evaluating the first 13 characters of headlines generated to be 26 characters long on 13char-ref headlines (blue); (b) evaluating headlines generated to be 13 characters long on 13char-ref headlines (pink); and setting (c) is the same as setting (a) but evaluated on the headline generated to be 13 characters long (green).

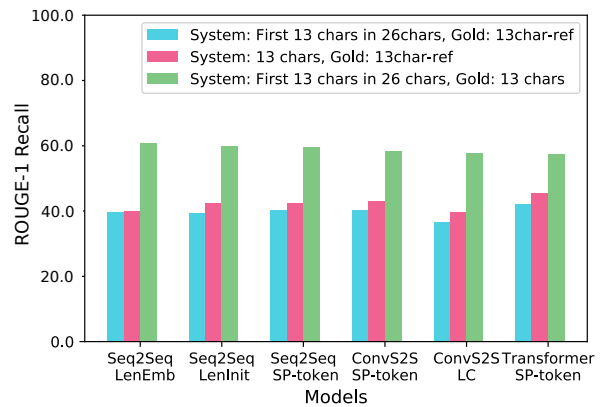


Figure 4: ROUGE-1 recall scores to assess the similarity of headlines generated for different lengths.

Setting (a) corresponds to the strategy where we trimmed headlines of different lengths to 13 characters long. This setting was worse than setting (b), where a method tailored headlines to the desired length. However, the difference in ROUGE scores between (a) and (b) was not very large, indicating that the existing methods do not drastically change the content for 13 characters long and 26 characters long. This tendency was also verified by setting (c), which assessed how much the first 13 characters of headlines generated to be 26 characters long covered the content of those generated to be 13 characters long. These facts suggest that we should explore in further research a method not only trained by generic supervision data (print headlines) but also tuned for the desired length.

## 5 Related Work

Rush et al. (2015) created the first approach to neural abstractive summarization. They generated a headline from the first sentence of a news article in the AEG (Napoles et al., 2012), which contains an enormous number of pairs of headlines and articles. After their study, a number of researchers addressed this task: For example, Chopra et al. (2016) used the encoder-decoder framework (Sutskever et al., 2014; Bahdanau et al., 2015) and Nallapati et al. (2016) incorporated additional features into the model, such as parts-of-speech tags and named entities. Suzuki and Nagata (2017) proposed word-frequency estimation to reduce the repeated phrases being generated. Zhou et al. (2017) proposed a gating mechanism (*sGate*) to ensure that important information is selected at each decoding step.

Furthermore, attempts to control the output



length in neural abstractive summarization have been gradually increasing. Shi et al. (2016) reported that hidden states in recurrent neural networks in the encoder-decoder framework could implicitly model the length of the output sequences. Kikuchi et al. (2016) was the first to propose the idea of controlling the output length in the encoder-decoder framework. Their approach inserts length information for the output length into the decoder. Additionally, Fan et al. (2018) reported that output lengths could be controlled by embedding special tokens given to an input sequence. These two studies used DUC 2004 (Over et al., 2007), which comprises only 75-byte summaries, to evaluate the outputs in multiple lengths. Liu et al. (2018) also proposed a method for controlling the number of output tokens in the ConvS2S model. In Transformer (Vaswani et al., 2017), Takase and Okazaki (2019) proposed two length control methods by extending positional embedding. They additionally evaluated the system outputs by using the reconstructed test set of AEG which consists of the fixed length headlines. Makino et al. (2019) proposed a global optimization method under a length constraint.

Sun et al. (2019) examined how to compare summarizers by considering the length bias of generated summaries in the test set that includes various length summaries. However, no previous work built a dataset for evaluating headlines of multiple lengths or reported an in-depth perspective on this task during the process of new production in the real world. However, a single length reference that could appropriately evaluate multiple length summaries in multiple document summarization was reported (Shapira et al., 2018). In that study, the authors confirmed the correlation coefficient of ROUGE scores between the scores using a single length reference and multiple (gold) length references in the evaluation. The present research differed in that we examined what kind of problems occurred and studied headline generation domain, which requires stricter keyword selection.

## 6 Conclusion

In this paper, we presented two new corpora: The JNC contains a large number of pairs of news articles and their headlines, and the JAMUL includes headlines of three different lengths (10, 13, and 26 characters long) written by professional editors.

This study is the first to analyze the characteristics of multiple headlines of different lengths, and to evaluate existing approaches for length control based on the reference headlines composed for different lengths. We found that Transformer model with a special length token (*SP-token*) outperformed the other methods on the JAMUL. Additionally, although we confirmed that single length (the longest) references could adequately evaluate multiple length system outputs, the rankings were swapped even if there is a difference of 1.0 or more in the ROUGE scores. We also confirmed the length-restricted test set, which was extracted from length-insensitive headlines and corresponding articles, could not adequately evaluate the multiple length system outputs depending on the specified length. In the analysis, the existing methods could not take into account the word selection according to length constraint. We also found it difficult to evaluate methods to controlling output length, because headlines of different lengths are written based on different goals, and because the training data does not necessarily reflect the goal of the headlines of a specific length. In the future, we plan to explore an approach to adapt a model trained on print headlines to those which dedicated to a different length.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *International Conference on Learning Representations (ICLR 2015)*.
- Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 2018. [Retrieve, rerank and rewrite: Soft template based neural summarization](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, pages 152–161.
- Wenhu Chen, Guanlin Li, Shuo Ren, Shujie Liu, Zhirui Zhang, Mu Li, and Ming Zhou. 2018. [Generative bridging network for neural sequence prediction](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2018)*, pages 1706–1715.
- Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2016)*, pages 93–98.
- Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In *Pro-*

- ceedings of the 2nd Workshop on Neural Machine Translation and Generation (WNMT 2018), pages 45–54.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *Proceedings of the 34th International Conference on Machine Learning, (ICML 2017)*, pages 1243–1252.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, (CVPR 2016)*, pages 770–778.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. [Controlling output length in neural encoder-decoders](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*, pages 1328–1338.
- Shun Kiyono, Sho Takase, Jun Suzuki, Naoaki Okazaki, Kentaro Inui, and Masaaki Nagata. 2018. [Unsupervised token-wise alignment to improve interpretation of encoder-decoder models](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP (BlackboxNLP 2018)*, pages 74–81.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations (ACL 2017)*, pages 67–72.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP 2018)*, pages 66–71.
- Taku Kudo, Kaoru Yamamoto, and Yuji Matsumoto. 2004. [Applying conditional random fields to japanese morphological analysis](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP 2004)*, pages 230–237.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81.
- Yizhu Liu, Zhiyi Luo, and Kenny Zhu. 2018. [Controlling length in abstractive summarization using a convolutional neural network](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*, pages 4110–4119.
- Takuya Makino, Tomoya Iwakura, Hiroya Takamura, and Manabu Okumura. 2019. [Global optimization under length constraint for neural text summarization](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, (ACL 2019)*, pages 1039–1048.
- Ramesh Nallapati, Bowen Zhou, Cícero Nogueira dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, (CoNLL 2016)*, pages 280–290.
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. [Annotated Gigaword](#). In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction (AKBC-WEKEX 2012)*, pages 95–100.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Demonstrations (NAACL-HLT 2019)*, pages 48–53.
- Paul Over, Hoa Dang, and Donna Harman. 2007. [DUC in context](#). *Information Processing and Management*, 43(6):1506–1520.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *6th International Conference on Learning Representations (ICLR 2018)*.
- 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 (EMNLP 2015)*, pages 379–389.
- Ori Shapira, David Gabay, Hadar Ronen, Judit Bar-Ilan, Yael Amsterdamer, Ani Nenkova, and Ido Dagan. 2018. [Evaluating multiple system summary lengths: A case study](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*, pages 774–778.
- Xing Shi, Kevin Knight, and Deniz Yuret. 2016. [Why neural translations are the right length](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*, pages 2278–2282.
- Simeng Sun, Ori Shapira, Ido Dagan, and Ani Nenkova. 2019. [How to compare summarizers without target length? pitfalls, solutions and re-examination of the neural summarization literature](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation (NeuralGen 2019)*, pages 21–29.

- Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. 2013. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, pages 1139–1147.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS 2014)*, pages 3104–3112.
- Jun Suzuki and Masaaki Nagata. 2017. [Cutting-off redundant repeating generations for neural abstractive summarization](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*, pages 291–297.
- Sho Takase and Naoaki Okazaki. 2019. [Positional encoding to control output sequence length](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2019)*, pages 3999–4004.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NIPS 2017)*, pages 6000–6010.
- Kai Wang, Xiaojun Quan, and Rui Wang. 2019. [Biset: Bi-directional selective encoding with template for abstractive summarization](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics (ACL 2019)*, pages 2153–2162.
- Li Wang, Junlin Yao, Yunzhe Tao, Li Zhong, Wei Liu, and Qiang Du. 2018. [A reinforced topic-aware convolutional sequence-to-sequence model for abstractive text summarization](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, (IJCAI 2018)*, pages 4453–4460.
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2017. [Selective encoding for abstractive sentence summarization](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 1095–1104.
- Qingyu Zhou, Nan Yang, Furu Wei, and Ming Zhou. 2018. [Sequential copying networks](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI 2018)*, pages 4987–4995.