

Reference and Document Aware Semantic Evaluation Methods for Korean Language Summarization

Dongyub Lee¹, Myeongcheol Shin², Taesun Whang³, Seungwoo Cho²,
Byeongil Ko², Daniel Lee², Eunggyun Kim², Jaechoon Jo^{4*}

¹Kakao Corp.

jude.lee@kakaocorp.com

²Kakao Enterprise Corp.

{index.sh, john.w, kobi.k, daniel.e, json.ng}@kakaenterprise.com

³Korea University, South Korea

{taesunwhang}@korea.ac.kr

⁴Hanshin University, South Korea

{jaechoon}@hs.ac.kr

Abstract

Text summarization refers to the process that generates a shorter form of text from the source document preserving salient information. Many existing works for text summarization are generally evaluated by using recall-oriented understudy for gisting evaluation (ROUGE) scores. However, as ROUGE scores are computed based on n-gram overlap, they do not reflect semantic meaning correspondences between generated and reference summaries. Because Korean is an agglutinative language that combines various morphemes into a word that express several meanings, ROUGE is not suitable for Korean summarization. In this paper, we propose evaluation metrics that reflect semantic meanings of a reference summary and the original document, *Reference and Document Aware Semantic Score (RDASS)*. We then propose a method for improving the correlation of the metrics with human judgment. Evaluation results show that the correlation with human judgment is significantly higher for our evaluation metrics than for ROUGE scores.

1 Introduction

The task of text summarization is to generate a reference summary that conveys all the salient information of an original document. There are two strategies for this type of summarization (i.e., extractive and abstractive summarization). With the extractive approach, the most noticeable key sentences are extracted from the source and compiled into a reference (Zhong et al., 2019; Wang et al., 2019; Xiao and Carenini, 2019). The second approach is abstractive, with which a paraphrased summary is generated from the source (Zhang et al., 2018; Guo et al., 2018; Wenbo et al., 2019). The generated summary may not contain the same words that appear in the source document. Therefore, measuring factual alignment between the generated summary and source document is important (Kryscinski et al., 2019).

Most summarization models are evaluated using recall-oriented understudy for gisting evaluation (ROUGE) (Lin, 2004a), which measures n-gram overlaps between generated and reference summaries. ROUGE has proven to have a high correlation with manual evaluation methods, such as pyramid (Nenkova et al., 2007) and TAC AESOP (Owczarzak and Dang, 2011). However, Louis (2013) showed that the correlation significantly decreased when only one reference summary was provided. Additionally, considering the process by which a person manually summarizes a document, ROUGE is limited, because it does not reflect semantic meanings between generated and reference summaries. For example, when a person summarizes a document, they tend to use words that are implicit while not always using the explicit words from the original document. As the ROUGE score is computed based on an n-gram overlap, the score can be low even if two words have the same semantic meaning. Table 1 shows an example of the ROUGE limitation when applied to a Korean summarization. This tendency is particularly prevalent in Korean, which is an agglutinative language that combines various morphemes into a

* corresponding author

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

word to express several meanings and grammatical functions, unlike English. In this process, complex morphological variations can occur. Therefore, leveraging ROUGE scores produces inaccurate results.

Article
<p>‘슬기로운 의사생활’이 또다시 최고 시청률을 경신하며 고공행진을 이어갔다. 26일 방송된 tvN 2020 목요 스페셜 ‘슬기로운 의사생활’ 3회는 케이블, IPTV, 위성을 통합한 유료플랫폼에서 가구 평균 8.6%, 최고 10%의 시청률을 기록했다. 3주 연속 시청률 상승세다.</p> <p>The tv program “sage doctor life” set an all time record for its highest viewer ratings. On the 26th, TVN 2020 Thursday Special “Sage Doctor Life,” aired on the 26th, recorded an average household rating of 8.6% and a maximum of 10% on a paid platform incorporating cable, IPTV, and satellite. The ratings have been rising for three consecutive weeks.</p>
Reference Summary
<p>‘슬기로운 의사생활’ 최고 시청률 10% 돌파... 3회 연속 상승세</p> <p>The tv program “Sage doctor life” breaks its all time high 10% viewer ratings for 3 consecutive episodes.</p>
Wrong Candidate
<p>‘슬기로운 의사생활’ 최저 시청률 10% 돌파... 3회 연속 하락세</p> <p>The tv program “Sage doctor life” reached the lowest viewing rate of 10% for 3 times in a row.</p> <p>Rouge scores with Reference Summary (R-1/R-2/R-L): 0.78/ 0.63/ 0.78</p> <p>Ours (RDASS): 0.44</p>
Correct Candidate
<p>‘슬기로운 의사생활’ 최고 시청률 경신... 3주 연속 상승세</p> <p>The tv program “Sage doctor life” set an all time record for its highest viewer ratings for 3 consecutive episodes.</p> <p>Rouge scores with Reference Summary (R-1/R-2/R-L): 0.71/ 0.53/ 0.71</p> <p>Ours (RDASS): 0.56</p>

Table 1: An example showing the limitations of ROUGE in Korean summarization. The incorrectly generated summary has a high ROUGE score, but has the opposite semantic meaning. Text areas marked in blue and red serve as indicators for distinguishing the factualness of the semantic comparisons, as reflected by the our metrics shown.

To overcome this limitation, an evaluation method that considers the semantic information of both the generated and reference summary is required. It is important to examine the factuality between the generated summary and source document, because the generated summary may contain false information. Each person summarizes information in different manners, and it is difficult to agree, even after cross-checking (Kryscinski et al., 2019). Therefore, the source document should also be considered with generated and reference summary.

In this study, we propose metrics for evaluating a summarization model that consider both the source document and reference summary together with the generated summary (see Table 1). Our contributions can be summarized as follows:

- We propose the evaluation metrics that can be applied to a summarization model using deep semantic information.
- We propose methods to improve the correlation between the proposed evaluation metrics and human judgment.
- Via extensive evaluation, we demonstrate that the correlation with human judgment is significantly higher for our proposed evaluation metrics than for ROUGE scores.

2 Related Work

Evaluation methods of text summarization are divided into two strategies: manual and automatic. Manual evaluation is expensive and difficult (Nenkova and Passonneau, 2004; Passonneau et al., 2013). Several studies have been conducted to develop automatic methods that facilitate fast and low-cost evaluations. There are two types of automatic evaluation methods: extrinsic and intrinsic. An extrinsic automatic method evaluates a summarization model based on how it affects the completion of tasks comprising the judgment of document relevance (Dorr et al., 2004). The intrinsic automatic method evaluates quality via a property analysis or by calculating its similarity to a manually generated summary. Intrinsic methods include the pyramid method (Nenkova et al., 2007), the basic-elements method (Hovy et al., 2006),

and ROUGE (Lin, 2004b). The pyramid method inspects various human-made summaries and creates summary content units, each with a scoring weight. The basic-elements method is similar to the pyramid method. ROUGE evaluates the similarity of the lexical overlap between the candidate and reference summary.

As the ROUGE score is computed based on the n-gram overlap, it does not account for synonymous words or phrases. Many approaches have been proposed to overcome this limitation. ParaEval (Zhou et al., 2006), ROUGE-WE (Ng and Abrecht, 2015), ROUGE 2.0 (Ganesan, 2018), and ROUGE-G (ShafieiBavani et al., 2018) have been used to extend ROUGE to support synonymous constructs. ParaEval uses a matching method based on paraphrase tables. ROUGE-WE uses a lexical matching method with a semantic similarity measure and the cosine distances between tokens. ROUGE 2.0 uses WordNet as a synonym dictionary and computed token overlaps with all synonyms of matched words. ROUGE-G uses lexical and semantic matching from WordNet. These approaches have limitations because they require hand-crafted lexical and synonym dictionaries, which are particularly difficult to construct in Korean. Our research is similar to (Zhang et al., 2019), which utilized BERT to compute semantic score between the generated and reference sentence. However, (Zhang et al., 2019) does not consider the document, whereas our research considers the document to be characterized in the evaluation of summarization tasks. Overall, our research is different from previous approaches in that 1) We propose a method to evaluate generated summary by considering documents as well as reference summary. 2) In addition, our evaluation model is robust to out of vocabulary (OOV) words because it leverages a pre-trained neural network (SBERT) based on byte pair encoding (BPE) (Gage, 1994) tokenization method from unsupervised learning. Considering the fact that Korean is an agglutinative language, this feature is essential. 3) Finally, Our evaluation model can be further trained to capture more contextualized information both on reference summary and document.

Text summarization models can be divided into abstractive, extractive, and hybrid. Abstractive models reword phrases and create summaries having novel phrases constructed from the original document. Recent text summarization approaches have leveraged multi-task and multi-reward training (Jiang and Bansal, 2018; Paulus et al., 2017; Pasunuru and Bansal, 2018; Guo et al., 2018), attention-with-copying mechanisms (Tan et al., 2017; See et al., 2017; Cohan et al., 2018), and unsupervised training strategies (Schumann, 2018; Chu and Liu, 2018). The extractive method extracts the most-suitable sentences (or words) from the source document and copies them directly into the summary. Many researchers (Neto et al., 2002; Colmenares et al., 2015; Filippova and Altun, 2013) have utilized domain expertise to develop heuristics for refining summary texts. Recently, neural-based text summarization models have been proposed to train the model for predicting whether a span of text should be included in the summary (Nallapati et al., 2016; Narayan et al., 2017; Xu and Durrett, 2019; Liu et al., 2019a). Reinforcement learning-based summarization models have also been proposed to directly optimize models (Wu and Hu, 2018; Dong et al., 2018; Narayan et al., 2018b). The hybrid approach uses both abstractive and extractive methods. With this approach, the summarization process is divided into two phases: content selection and paraphrasing (Gehrmann et al., 2018; Hsu et al., 2018; Chen and Bansal, 2018; Liu et al., 2018).

3 Methodology

From Table 1, we can observe the importance of considering both the document and reference summary together for proper evaluation of the summarization model. In Subsection 3.1, we propose a method for evaluating the generated summary with the reference summary to reflect deep semantic meaning. Next, we propose a method for evaluating the generated summary with the original document and reference summary together. The reference-document-aware evaluation metric model can be further trained to capture more contextualized information from both on reference summary and document (Subsection 3.2).

3.1 Reference and Document Aware Semantic Evaluation

Let us define the generated summary from the summarization model as $y_p = [w_1, \dots, w_n]$ and reference summary as $y_r = [w_1, \dots, w_m]$, where w_i indicates each word. Then, each summary representation, v_p and v_r , can be constructed using sentence-embedding methods. Neural-based sentence-embedding

methods have been broadly studied. Conneau (2017) trained a siamese bidirectional long short-term memory model with a max-pooling strategy on the Stanford Natural Language Inference (SNLI) corpus (Bowman et al., 2015) and the MultiGenre Natural Language Inference (NLI) dataset (Williams et al., 2017). Cer (2018) proposed the universal sentence encoder to train a transformer on the SNLI dataset. Reimers (2019) recently proposed sentence-BERT(SBERT), which leverages a pre-trained BERT (Devlin et al., 2018), trained with a combination of the SNLI and multi-genre NLI, and showed state-of-the-art sentence embedding performance. SBERT is suitable for semantic similarity searches and showed faster inference speeds than previous state-of-the-art approaches, including BERT, RoBERTa (Liu et al., 2019b), and the universal sentence encoder.

We leverage a pre-trained SBERT to construct summary representations. Each word representation, e , is obtained from SBERT as

$$E = [e_{cls}, e_1, \dots, e_n, e_{sep}] = \text{SBERT}([CLS], w_1, \dots, w_n, [SEP]). \quad (1)$$

Subsequently, mean-pooling is performed to construct v_p as,

$$v_p(j) = \frac{\sum_{i=0}^n e_i[j]}{n} \quad (2)$$

where j represents an index of a word-embedding dimension, and n represents a length of E . v_r can also be obtained in the same manner.

The semantic similarity score, $s(p, r)$, between v_p and v_r can be obtained as follows:

$$s(p, r) = \cos(v_p, v_r) = \frac{v_p^T \cdot v_r}{\|v_p\| \|v_r\|}. \quad (3)$$

Recall that it is important to consider factual consistency of generated summary with the source document, and, given the same document, the method of summarizing important information varies from person to person (Owczarzak et al., 2012; Kryscinski et al., 2019). Therefore, the source document should also be considered with the generated summary when evaluating the summarization model.

Given a document, $D = [w_1, \dots, w_k]$, the document representation, v_d , can be obtained using Eqs. (1) and (2). Thus, the similarity score between v_p and v_d can be defined as,

$$s(p, d) = \cos(v_p, v_d) = \frac{v_p^T \cdot v_d}{\|v_p\| \|v_d\|}. \quad (4)$$

Given a reference and source document, the reference-document-aware semantic score (*RDASS*) of the generated summary is defined by averaging $s(p, r)$ and $s(p, d)$:

$$RDASS = \frac{s(p, r) + s(p, d)}{2}. \quad (5)$$

We also experimented with a sum, max and min operation between $s(p, r)$ and $s(p, d)$, but averaging the two scores reports highest correlation with human judgment.

3.2 Fine-tuning SBERT with the Abstractive Summarization Model

As SBERT is a trainable metric model, it can be further trained to capture more contextualized information about the reference summary and source document. We propose a fine-tuning method for SBERT that uses the abstractive summarization model.

Most neural approaches for abstractive summarization are based on an encoder-decoder architecture (See et al., 2017). Formally, given a document, $D = [w_1, \dots, w_k]$, the objective is to generate a summary, $y_p = [w_1, \dots, w_n]$, from a hidden representation, $h_p = [h_1, \dots, h_n]$. The hidden representation is the output vector of the decoder. We leverage the hidden representation of the decoder to fine-tune the SBERT.

Following (Reimers and Gurevych, 2019), we adopt a triplet objective to fine-tune the SBERT. Given an anchor h_p , a positive reference representation v_r^p , a negative representation v_r^n , and a Euclidean distance d , the triplet objective for generated and reference summaries $J(p, r)$ is then defined as

$$J(p, r) = \max(0, \epsilon + d(h_p, v_r^p) - d(h_p, v_r^n)), \quad (6)$$

where ϵ represents a margin that ensures h_p is closer to v_r^p than v_r^n . We set ϵ as 1. Similarly, the triplet objective for generated summary and document can be defined as

$$J(p, d) = \max(0, \epsilon + d(h_p, v_d^p) - d(h_p, v_d^n)). \quad (7)$$

Thus, the final objective for SBERT is to minimize the combined two triplet objectives as

$$J = J(p, r) + J(p, d). \quad (8)$$

The objective function SBERT J is jointly optimized with the abstractive summarization objective. Usually, the negative log-likelihood objective between the generated and reference summaries is used for abstractive summarization (See et al., 2017; Narayan et al., 2018a). We refer to the fine-tuned SBERT with abstractive summarization model as “FWA-SBERT.”

4 Experimental Setup

4.1 Dataset

We trained and evaluated our models using the Korean Daum/News dataset¹, comprising 10 topics, such as politics, economy, international, culture, information technology, and others. From this, we extracted 3-million news articles. The number of articles for training, validating, and testing was 2.98M, 0.01M, and 0.01M respectively. We refer to this dataset as Daum/News. We used Daum/News to fully understand the content of the article and conduct a proper evaluation. The dataset contains articles from 143 newspapers, each having different summary styles, and the effectiveness of the proposed methods is exemplified using it. Therefore, we expect that our research can be applied to different languages.

4.2 Summarization Model

We adopted abstractive summarization model of (Liu and Lapata, 2019)². Liu (2019) leveraged pre-trained BERT as an encoder and a six-layered transformer as a decoder, showing state-of-the-art results on Cable News Network/DailyMail (Hermann et al., 2015), New York Times (Sandhaus, 2008), and XSum (Narayan et al., 2018a) datasets. We set all environments according to (Liu and Lapata, 2019), except that we leveraged the pre-trained BERT trained on Korean dataset (Subsection 4.3) instead of english-bert-base-uncased. We trained the abstractive summarization model on Korean Daum/News dataset.

4.3 SBERT

To leverage SBERT, we first pre-trained BERT (bert-base-uncased) on Korean dataset, comprising 23M sentences and 1.6M documents, including Wiki, Sejong corpus, and web documents. Next, we trained SBERT with classification and regression objectives from NLI (Bowman et al., 2015; Williams et al., 2017) and the semantical textual similarity (STS) benchmark (STSb) (Cer et al., 2017). Because NLI and STSb datasets are in English, we leveraged the Korean NLI and STS dataset³ (Ham et al., 2020) which translated from Kakao Machine Translator⁴. Evaluation of the STS benchmark test dataset was conducted, showing an 80.52 Spearman’s rank correlation result. Subsequently, the pre-trained SBERT model was fine-tuned with the abstractive summarization model to capture more contextualized information of the reference summary and source document with a generated summary (Subsection 3.2). All training was conducted on the Kakao Brain Cloud with 4 Tesla V100 graphical processing units.

¹<https://media.daum.net/>

²<https://github.com/nlpyang/PreSumm>

³<https://github.com/kakaobrain/KorNLUDatasets>

⁴<https://translate.kakao.com>

4.4 Human Judgment

To demonstrate the effectiveness of the reference-document-aware semantic metric, we evaluated its correlation with human judgment. Following (Kryscinski et al., 2019), we asked annotators to score relevance, consistency, and fluency. **Relevance** represents the degree of appropriateness of the document, **consistency** represents the degree of factualness, and **fluency** represents the degree of the quality of generated summary. Additionally, **human avg** represents the average value of the scores for the three indicators. Given a document, reference summary, and generated summary, each annotator scored in the range of 1 to 5 points for the evaluation indicator (i.e., relevance, consistency, fluency). The human judgment was conducted by 6 judges having a PhD (3 judges) or a MS (3 judges) degree in computer science. The averaged human score of relevance was 3.8, consistency was 3.6, and fluency was 3.9 for 200 sampled summaries from Korean Daum/News test dataset.

5 Results

In this section, we first report the performance of the summarization model using the ROUGE and proposed evaluation metrics (Subsection 3.1). Next, we report how the proposed evaluation metrics correlated to human judgment. We also report the correlation of the proposed evaluation metrics to ROUGE to show that the proposed methods complement ROUGE. Finally, through qualitative evaluation, we demonstrate the limitations of ROUGE and the superiority of the proposed evaluation metrics.

5.1 Performance of the Summarization Model

Model	Proposed Evaluation Metrics			ROUGE		
	$s(p, r)$	$s(p, d)$	$RDASS$	1	2	L
Reference Summary	1.00	0.55	0.78	1.00	1.00	1.00
Lead-1	0.71	0.64	0.68	0.13	0.03	0.13
Lead-3	0.66	0.79	0.73	0.07	0.01	0.07
BERTSUMABS (Liu and Lapata, 2019)	0.83	0.46	0.65	0.35	0.15	0.35

Table 2: Performance of the summarization model on the DAUM/NEWS dataset.

The abstractive summarization model is based on the neural architecture of (Liu and Lapata, 2019). We trained the summarization model on the Daum/News dataset. To evaluate the summarization model, we used ROUGE and the proposed evaluation metrics. The fine-tuned FWA-SBERT was then used to evaluate the proposed semantic scores ($s(p, r)$, $s(p, d)$, and $RDASS$). Table 2 shows the performance of the summarization model with baseline methods (Reference Summary, Lead 1, and 3) on the Daum/News dataset.

We set the reference summary as upper-bound. In the case of the reference summary, the reporter tends to use implicit words when summarizing the document, so the $s(p, d)$ score is relatively low compared to the Lead baselines. However, because the $s(p, r)$ score is 1.00, the reference summary shows the highest $RDASS$ score. For Lead-1, $s(p, r)$ shows higher performance than $s(p, d)$, and for Lead-3, $s(p, d)$ shows higher performance than $s(p, r)$. The reason for this performance is that Lead-3 contains more sentences from the document, so the similarity with the reference summary $s(p, r)$ is low, but the similarity with the document $s(p, d)$ is increased. In the case of ROUGE performance of lead baselines, relatively low performance can be confirmed compared to other researches (Kryscinski et al., 2019) conducted in English dataset. The reason is that in the case of Korean, the same semantic meaning is expressed differently because of the nature of the language of the agglutinative language. A detailed example of this is described in Table 5 below. However, it can be seen that the $RDASS$ score of lead baselines is similar to that of the reference summary. Through this, we can confirm that the proposed evaluation method can reflect the semantic meaning of the reference summary and document well. In the case of the (Liu and Lapata, 2019), it shows higher similarity with the reference summary than the Lead baselines, but since it is based on the generation model, it does not extract the sentence from the document as the Lead baselines. As a result, it shows the relatively low $s(p, d)$ score. We describe how these results are correlated with human judgment in the next section.

5.2 Correlation with Human Judgment

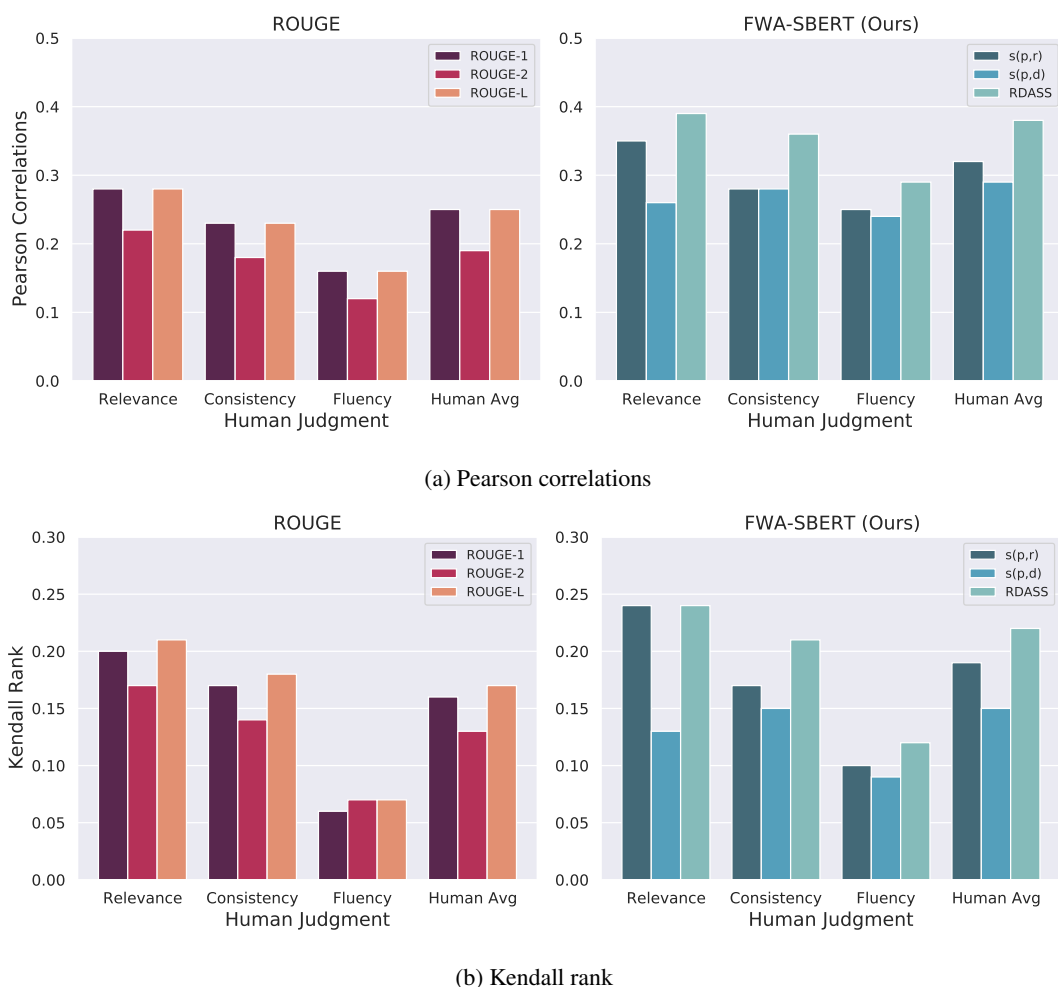


Figure 1: Pearson correlations and Kendall rank of the proposed evaluation metrics with human judgment.

Figures (1a) and (1b) show the Pearson correlation and Kendall rank, respectively, of the proposed evaluation metrics with human judgment on the 200 sampled summaries. Pearson correlation measure whether the two variables are linearly related, where 1 indicates positive linear correlation and -1 indicates negative linear correlation. And Kendall rank measure the rank correlation of the two variables, where 1 indicates two variables are similar and -1 indicates dissimilar. Both correlation measure methods are widely used in summarization task to analyze correlation with human judgment.

In the Pearson correlation matrix, the correlation with human judgment was significantly higher for the proposed evaluation metrics than for ROUGE scores. Additionally, in the Kendall rank matrix, the proposed evaluation metrics showed highest correlation with human judgment than did the ROUGE scores. Among the proposed evaluation metrics, $s(p,r)$ showed higher performance than $s(p,d)$ and RDASS showed the highest correlation with human judgment. These results indicate that the proposed evaluation metrics can reflect deep semantic meaning overcoming the limitations of ROUGE which based on n-gram overlap.

To demonstrate the effectiveness of fine-tuning SBERT with an abstractive summarization model, we set baseline methods depending on which sentence representation methods to use for the proposed methods (Subsection 3.1) as follows:

Multilingual Universal Sentence Encoder (MUSE): MUSE (Yang et al., 2019) is a multilingual sentence encoder that embeds text from 16 languages into a single semantic space using multi-task

Sentence Representation		Relevance		Consistency		Fluency		Human Avg	
		Pearson	Kendall	Pearson	Kendall	Pearson	Kendall	Pearson	Kendall
MUSE	$s(p, r)$	0.29	0.19	0.18	0.10	0.22	0.08	0.25	0.13
	$s(p, d)$	0.09	0.05	0.13	0.06	0.15	0.04	0.13	0.06
	RDASS	0.29	0.19	0.24	0.12	0.23	0.09	0.28	0.14
P-SBERT	$s(p, r)$	0.34	0.22	0.27	0.17	0.25	0.09	0.32	0.18
	$s(p, d)$	0.24	0.13	0.27	0.15	0.22	0.09	0.27	0.15
	RDASS	0.37	0.22	0.34	0.20	0.29	0.11	0.37	0.21
FWA-SBERT	$s(p, r)$	0.35	0.24	0.28	0.17	0.25	0.10	0.32	0.19
	$s(p, d)$	0.26	0.13	0.28	0.15	0.24	0.09	0.29	0.15
	RDASS	0.39	0.24	0.36	0.21	0.29	0.12	0.38	0.22

Table 3: Performance comparison depended upon which sentence representation was used.

learning. This model was trained on more than 1-billion question-answer pairs and showed competitive state-of-the-art results on semantic (Gillick et al., 2018), bitext retrieval (Ziemski et al., 2016), and retrieval question-answering (Yang et al., 2019).

Pre-trained SBERT: We only leveraged pre-trained SBERT without fine-tuning. We refer to this as “P-SBERT.”

Table 3 show the performance comparison depended upon which sentence representation was used. P-SBERT shows the high correlation coefficient with humans than MUSE. Overall, when the FWA-SBERT was used, it showed the closest correlation with human judgment.

Through quantitative evaluation, we demonstrated that the proposed evaluation metrics had a high correlation with human judgment and that the method of fine-tuning SBERT improved the performance of the proposed evaluation metrics.

We also experimented to understand how each evaluation metric was correlated to each other. As shown in Table 4, there was a high correlation among the ROUGE metrics. However, the proposed evaluation metrics had a relatively low correlation with ROUGE. This indicates that the proposed evaluation metrics reflected semantic meaning, in our case, that ROUGE could not. Thus, we believe it complements ROUGE metrics.

	ROUGE-1	ROUGE-2	ROUGE-L	$s(p, r)$	$s(p, d)$	RDASS
ROUGE-1	1.00	0.84	0.99	0.64	0.16	0.54
ROUGE-2		1.00	0.85	0.52	0.09	0.45
ROUGE-L			1.00	0.63	0.17	0.53
$s(p, r)$				1.00	0.32	0.77
$s(p, d)$					1.00	0.69
RDASS						1.00

Table 4: Pearson correlation of ROUGE and the proposed evaluation metrics.

5.3 Qualitative Analysis

In this section, through qualitative analysis, we demonstrate the effectiveness of our evaluation metrics. Table 5 shows ROUGE, RDASS and human evaluation results for the generated summaries for the two articles.

In article-1, the generated summary “*On the 30th birthday of Messi, he had a good time with his family*” has the same semantic meaning as the reference summary “*Messi’s 30th birthday with his wife and son*”. However, since the sentence having the same semantic meaning can be variously expressed in Korean, which has the characteristics of agglutinative language, the ROUGE score is low while human evaluation scores are high. Likewise, the generated summary “*Samsung Electronics launches new ‘qled tv’ in Brazil*” in article-2 has a same semantic meaning as the reference summary “*Samsung Electronics*

Article-1
<p>리오넬 메시(30·fc바르셀로나)가 자신의 서른 번째 생일을 가족과 함께 오붓하게 보냈다. 지난 24일 만 서른 살이 된 메시는 자신의 인스타그램에 집에서 가족들과 함께 보낸 생일상을 찍은 사진을 올렸다. 메시는 오랜 그의 여자친구이자, 이제 아내가 되는 안토넬라 로쿠조(29), 아들 티아고가 함께 다정하게 사진을 찍었다.</p> <p>Lionel Messi (30 fc Barcelona) spent his thirtieth birthday with his family. Messi, who turned thirty on the 24th, posted a picture of his birthday on Instagram with his family at home. Messi was tenderly photographed by his longtime girlfriend, Antonella Rokujo (29), and his son, Thiago.</p>
Reference Summary
<p>메시가 30번째 생일 함께한 이는 아내와 아들</p> <p>Messi's 30th birthday with his wife and son.</p>
Generated Summary
<p>메시 30번째 생일, 가족과 함께 오붓하게 보내</p> <p>On the 30th birthday of Messi, he had a good time with his family.</p> <p>Rouge(1/ 2/ L): 0.14/ 0.00/ 0.14</p> <p>RDASS: 0.81</p> <p>Human Evaluation (relevance/ consistency/ fluency): 4.4/ 4.2/ 4.2</p>
Article-2
<p>삼성전자는 19일(현지시간) 브라질 상파울루의 팔라시오 탕가라 호텔에서 'QLED TV 론칭 이벤트'를 열고 2017년형 QLED TV 라인업을 선보였다고 23일 밝혔다. 4월 중남미에서는 처음으로 멕시코에서 QLED TV를 출시한 뒤 파나마, 콜롬비아 등으로 확대하다 이번에 중남미 최대 시장인 브라질에 제품을 출시한 것이다. 브라질은 전체 중남미 TV 시장의 40%(금액 기준) 이상을 차지할 정도로 중요한 TV 시장이다. 올해 1~4월 브라질 TV 시장은 작년 같은 기간보다 13%(수량 기준) 성장했고, 특히 프리미엄 TV 시장인 UHD(초고화질) TV는 작년보다 50% 이상 시장이 커졌다. 특히 삼성전자는 브라질 UHD TV 시장에서 올해 1~4월 56%(수량 기준) 점유율로 압도적 1위를 차지했다.</p> <p>Samsung Electronics announced on the 23rd that it held a 'QLED TV launching event' at the Palacio Tangara Hotel in Sao Paulo, Brazil on the 19th (local time) and introduced the 2017 QLED TV lineup. In April, it launched the QLED TV in Mexico for the first time in Latin America, and then expanded to Panama, Colombia, etc. This time it launched the product in Brazil, the largest market in Latin America. Brazil is an important TV market, accounting for more than 40% of total Latin American TV market. In January-April this year, the Brazilian TV market grew 13% (in quantity) from the same period last year. In particular, the UHD (Ultra High Definition) TV, a premium TV market, was 50% larger than last year. In particular, Samsung Electronics took the dominant position in the UHD TV market in Brazil with 56% (based on quantity) in January-April this year.</p>
Reference Summary
<p>삼성전자, 중남미 최대 시장 브라질에 qled tv 론칭</p> <p>Samsung Electronics launches 'qled tv' in Brazil, the largest market in Latin America.</p>
Generated Summary
<p>삼성전자, 브라질서 'qled tv' 신제품 출시</p> <p>Samsung Electronics launches new 'qled tv' in Brazil.</p> <p>Rouge(1/ 2/ L): 0.14/ 0.00/ 0.14</p> <p>RDASS: 0.71</p> <p>Human Evaluation(relevance/ consistency/ fluency): 4.6/ 4.4/ 4.4</p>

Table 5: Example articles from the DAUM/News test dataset. ROUGE, RDASS and human evaluation results for the generated summaries are represented.

launches 'qled tv' in Brazil, the largest market in Latin America". The generated summary in both articles is correct, but the ROUGE score is low. On the other hand, the RDASS score indicates a higher score, and indicates that the generated summary is the correct answer.

6 Conclusion

In this paper, we pointed out the limitation of the widely used ROUGE evaluation metric when adopting Korean summarization. Since Korean is an agglutinative language, the generated summary having the same semantic meaning with reference summary can be variously expressed. Therefore, only leveraging ROUGE metric can produce inaccurate evaluation results. To overcome this limitation, we proposed RDASS (Reference and Document Aware Semantic Score) evaluation metric. The RDASS can reflect deep semantic relationships of a generated, reference summary, and document. Through extensive evaluations, we demonstrated that the correlation with human judgment is higher for the proposed evaluation metric (RDASS) than for ROUGE scores. In future work, we will demonstrate the effectiveness of the proposed method in English summarization dataset.

Acknowledgements

We thank GyoungEun Han, the member of the Kakao Enterprise NLP team, for giving linguistic advice in writing this paper. This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program (IITP-2020-2018-0-01405) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation).

References

- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. *arXiv preprint arXiv:1805.11080*.
- Eric Chu and Peter J Liu. 2018. Unsupervised neural multi-document abstractive summarization. *arXiv preprint arXiv:1810.05739*.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. *arXiv preprint arXiv:1804.05685*.
- Carlos A Colmenares, Marina Litvak, Amin Mantrach, and Fabrizio Silvestri. 2015. Heads: Headline generation as sequence prediction using an abstract feature-rich space.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. Banditsum: Extractive summarization as a contextual bandit. *arXiv preprint arXiv:1809.09672*.
- Bonnie Dorr, Christof Monz, Douglas Oard, David Zajic, and Richard Schwartz. 2004. Extrinsic evaluation of automatic metrics for summarization. Technical report, MARYLAND UNIV COLLEGE PARK INST FOR ADVANCED COMPUTER STUDIES.
- 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.
- Philip Gage. 1994. A new algorithm for data compression. *C Users Journal*, 12(2):23–38.
- Kavita Ganesan. 2018. Rouge 2.0: Updated and improved measures for evaluation of summarization tasks. *arXiv preprint arXiv:1803.01937*.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.
- Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. 2018. End-to-end retrieval in continuous space. *arXiv preprint arXiv:1811.08008*.
- Han Guo, Ramakanth Pasunuru, and Mohit Bansal. 2018. Soft layer-specific multi-task summarization with entailment and question generation. *arXiv preprint arXiv:1805.11004*.
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020. Kornli and korsts: New benchmark datasets for korean natural language understanding. *arXiv preprint arXiv:2004.03289*.

- 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.
- Eduard H Hovy, Chin-Yew Lin, Liang Zhou, and Junichi Fukumoto. 2006. Automated summarization evaluation with basic elements. In *LREC*, volume 6, pages 899–902. Citeseer.
- Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. *arXiv preprint arXiv:1805.06266*.
- Yichen Jiang and Mohit Bansal. 2018. Closed-book training to improve summarization encoder memory. *arXiv preprint arXiv:1809.04585*.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Neural text summarization: A critical evaluation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551.
- Chin-Yew Lin. 2004a. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July. Association for Computational Linguistics.
- Chin-Yew Lin. 2004b. Rouge: A package for automatic evaluation of summaries. In *Proceedings of Workshop on Text Summarization Branches Out, Post2Conference Workshop of ACL*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*.
- Jingyun Liu, Jackie CK Cheung, and Annie Louis. 2019a. What comes next? extractive summarization by next-sentence prediction. *arXiv preprint arXiv:1901.03859*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Annie Louis and Ani Nenkova. 2013. Automatically assessing machine summary content without a gold standard. *Computational Linguistics*, 39(2):267–300.
- Ramesh Nallapati, Bowen Zhou, and Mingbo Ma. 2016. Classify or select: Neural architectures for extractive document summarization. *arXiv preprint arXiv:1611.04244*.
- Shashi Narayan, Nikos Papasrantopoulos, Shay B Cohen, and Mirella Lapata. 2017. Neural extractive summarization with side information. *arXiv preprint arXiv:1704.04530*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018a. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018b. Ranking sentences for extractive summarization with reinforcement learning. *arXiv preprint arXiv:1802.08636*.
- Ani Nenkova and Rebecca J Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the human language technology conference of the north american chapter of the association for computational linguistics: Hlt-naacl 2004*, pages 145–152.
- Ani Nenkova, Rebecca Passonneau, and Kathleen McKeown. 2007. The pyramid method: Incorporating human content selection variation in summarization evaluation. *ACM Transactions on Speech and Language Processing (TSLP)*, 4(2):4–es.
- Joel Larocca Neto, Alex A Freitas, and Celso AA Kaestner. 2002. Automatic text summarization using a machine learning approach. In *Brazilian symposium on artificial intelligence*, pages 205–215. Springer.
- Jun-Ping Ng and Viktoria Abrecht. 2015. Better summarization evaluation with word embeddings for rouge. *arXiv preprint arXiv:1508.06034*.

- Karolina Owczarzak and Hoa Trang Dang. 2011. Overview of the tac 2011 summarization track: Guided task and aesop task. In *Proceedings of the Text Analysis Conference (TAC 2011)*, Gaithersburg, Maryland, USA, November.
- Karolina Owczarzak, Hoa Trang Dang, Peter A Rankel, and John M Conroy. 2012. Assessing the effect of inconsistent assessors on summarization evaluation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 359–362. Association for Computational Linguistics.
- Rebecca J Passonneau, Emily Chen, Weiwei Guo, and Dolores Perin. 2013. Automated pyramid scoring of summaries using distributional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 143–147.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Multi-reward reinforced summarization with saliency and entailment. *arXiv preprint arXiv:1804.06451*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Raphael Schumann. 2018. Unsupervised abstractive sentence summarization using length controlled variational autoencoder. *arXiv preprint arXiv:1809.05233*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Elaheh ShafieiBavani, Mohammad Ebrahimi, Raymond Wong, and Fang Chen. 2018. A graph-theoretic summary evaluation for rouge. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 762–767.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. Abstractive document summarization with a graph-based attentional neural model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1171–1181.
- Hong Wang, Xin Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019. Self-supervised learning for contextualized extractive summarization. *arXiv preprint arXiv:1906.04466*.
- Wang Wenbo, Gao Yang, Huang Heyan, and Zhou Yuxiang. 2019. Concept pointer network for abstractive summarization. *arXiv preprint arXiv:1910.08486*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Yuxiang Wu and Baotian Hu. 2018. Learning to extract coherent summary via deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Wen Xiao and Giuseppe Carenini. 2019. Extractive summarization of long documents by combining global and local context. *arXiv preprint arXiv:1909.08089*.
- Jiacheng Xu and Greg Durrett. 2019. Neural extractive text summarization with syntactic compression. *arXiv preprint arXiv:1902.00863*.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-Hsuan Sung, et al. 2019. Multilingual universal sentence encoder for semantic retrieval. *arXiv preprint arXiv:1907.04307*.
- Fang-Fang Zhang, Jin-ge Yao, and Rui Yan. 2018. On the abstractiveness of neural document summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 785–790.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Ming Zhong, Pengfei Liu, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2019. Searching for effective neural extractive summarization: What works and what’s next. *arXiv preprint arXiv:1907.03491*.

Liang Zhou, Chin-Yew Lin, Dragos Stefan Munteanu, and Eduard Hovy. 2006. Paraeval: Using paraphrases to evaluate summaries automatically. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 447–454. Association for Computational Linguistics.

Michał Ziemiński, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The united nations parallel corpus v1.0. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3530–3534.