

ROUGE-K: Do Your Summaries Have Keywords?

Sotaro Takeshita¹, Simone Paolo Ponzetto¹, Kai Eckert²

¹Data and Web Science Group, University of Mannheim, Germany

²Mannheim University of Applied Sciences, Mannheim, Germany

{sotaro.takeshita, ponzetto}@uni-mannheim.de

k.eckert@hs-mannheim.de

Abstract

Keywords, that is, content-relevant words in summaries play an important role in efficient information conveyance, making it critical to assess if system-generated summaries contain such informative words during evaluation. However, existing evaluation metrics for extreme summarization models do not pay explicit attention to keywords in summaries, leaving developers ignorant of their presence. To address this issue, we present a keyword-oriented evaluation metric, dubbed ROUGE-K, which provides a quantitative answer to the question of – *How well do summaries include keywords?* Through the lens of this keyword-aware metric, we surprisingly find that a current strong baseline model often misses essential information in their summaries. Our analysis reveals that human annotators indeed find the summaries with more keywords to be more relevant to the source documents. This is an important yet previously overlooked aspect in evaluating summarization systems. Finally, to enhance keyword inclusion, we propose four approaches for incorporating word importance into a transformer-based model and experimentally show that it enables guiding models to include more keywords while keeping the overall quality.¹

1 Introduction

Summarization systems compress long documents into shorter ones to convey important information more effectively to readers (Rush et al., 2015; Chopra et al., 2016). To convey all essential information correctly, it is crucial for summarization systems to include important, i.e., summary-relevant keywords. In our analysis, human annotators find that summaries with more keywords, words that are relevant for the summary (see Section 3), capture important information better than the ones with fewer keywords. However, existing

¹Our code: <https://github.com/sobamchan/rougek>.

Reference	R-1	BS
A novel, hybrid deep learning approach provides the best solution to a limited-data problem (which is important to the conservation of the Hawaiian language)		
Hypothesis 1: We propose two methods to solve the transliteration problem automatically, given that there were not enough data to train an end-to-end deep learning model.	27.45	0.8718
Hypothesis 2: We propose two methods to solve the Hawaiian orthography transliteration problem automatically using finite state transducers and a hybrid neural network.	26.09	0.8692

Table 1: An example where ROUGE and BERTScore (BS) can lead to misinterpretations. Although the incorrect generation (not including the word “Hawaiian”) in the first hypothesis is more critical than the one in the second summary (“neural network” instead of “deep learning”) to convey correct information, both metrics assign a higher score to the former summary.

evaluation metrics do not explicitly take such word importance into account. Table 1 shows an example. Two commonly used metrics, namely ROUGE (Lin, 2004) and BERTScore (Zhang et al., 2020), assign lower scores to the second hypothesis even though it contains an essential word that another summary misses. This discrepancy, namely that ROUGE assigns a lower score to a summary that contains more keywords and annotators find relevant, happens in 16.7% of the cases in our analysis.

In this paper, to shed light on this problem, we propose ROUGE-K, an extension of ROUGE which considers only those n-grams in the summaries that match a set of pre-defined keywords. We propose a simple heuristic that exploits the common structure of summarization datasets to extract keywords automatically, making it possible for our metric to scale in size and domain without additional annotation effort. Correlation analysis reveals that there is only a weak strength of dependence between our new metric and existing ones as well as summary lengths, thus demonstrating

that our approach can complement, rather than replace, previous metrics. Through a manual evaluation, we find that human annotators show substantially higher agreement with ROUGE-K than with ROUGE and BERTScore on relevance, in other words, how well summaries include important information, which is one of four commonly assessed aspects in manual evaluations of summaries (Fabri et al., 2021). This shows that while one still would use traditional ROUGE to assess the overall qualities, our metric can provide a better index for evaluating the relevance of summaries. As a showcase of this new metric, we evaluate both extractive (Liu and Lapata, 2019) and abstractive (Lewis et al., 2020; Dou et al., 2021; Saito et al., 2020) state-of-the-art models on two extreme summarization datasets from different domains, namely XSum (Narayan et al., 2018) and SciTLDR (Cachola et al., 2020), as well as a more traditional, non-extreme dataset, ScisummNet (Yasunaga et al., 2019). Besides news text (XSum), we choose summarization of scientific publications (ScisummNet and SciTLDR), since this is a domain where keyword inclusion within summaries plays a crucial role. Surprisingly, the results reveal that these strong baseline models often fail to include essential words in their summaries, and that ROUGE-K enables us to better distinguish systems’ performance than alternative metrics. We also apply our ROUGE-K to the evaluations of recent large language models (LLMs) and show how our metric better accounts for the powerful capabilities of LLM-based summarizers when compared to traditional ROUGE metrics. Finally, As a first attempt to address the limitations on summary keyword inclusion, we introduce four ways to incorporate a lightweight word importance feature into existing transformer-based models. Experiments show that our methods can guide models to include more keywords without any additional annotations and negative effects on overall summarization quality. Our contributions are the following ones:

- We introduce **a new keyword-oriented evaluation metric, dubbed ROUGE-K**, which complements existing metrics by focusing on keywords.
- **We validate our metric:** a) against human judgments of summary relevance, b) by quantifying its correlation to existing metrics and summary lengths, and c) its ability to distinguish performance among different systems.
- Our experiments on three different datasets for

summarization of scientific and news articles reveal that current **state-of-the-art models often fail to include important words** in summaries.

- We present experiments with **four approaches to incorporate word-importance scores** into BART and show that it can help to improve keyword inclusion without hurting the overall summarization qualities.

2 The need for another kind of ROUGE

ROUGE (Lin, 2004) is a long-running *de facto* standard evaluation metric for summarization systems. It is very popular due to its high correlation with human evaluation while keeping its simplicity and interpretability. However, several works report on its limitations (Akter et al., 2022; Fabbri et al., 2021), one of which is that it only takes the word surface into account and disregards semantics (Ng and Abrecht, 2015). Because it considers all n-gram matches to be equally important, ROUGE fails to detect salient words that underpin a summary’s quality.

As an example, Table 1 shows two generated summaries of the same article from the SciTLDR dataset (Cachola et al., 2020) as well as their scores computed by ROUGE and BERTScore (Zhang et al., 2020), a pre-training language model based metric. Both metrics assign a higher score to the first summary even though it misses an important keyword that the second summary contains. In the case of ROUGE, this is because it favors the longer but nonessential n-gram overlaps in the first summary. This limitation of evaluation metrics can mislead the development of summarization systems towards including more of longer but less important words in summaries than truly essential keywords. When multiple reference summaries are available, ROUGE can assign higher scores to words potentially more important than others by counting n-grams that appear several times across references, which indirectly considers word importance. However, most commonly used datasets contain only one reference summary (Hermann et al., 2015; Narayan et al., 2018). In addition, because of its implicit nature, when a generated summary has a different textual style (even if the semantics of the summary did not change) from its reference summary, the ROUGE score can easily deflate.

SciTLDR (Cachola et al., 2020)

We show that autoregressive models can generate high fidelity images .
We introduce a new inductive bias that integrates tree structures in recurrent neural networks.

XSum (Narayan et al., 2018)

Opec , the oil producers’ group is back in the driving seat.
Lenovo and Acer have both unveiled smartphones with much larger than normal batteries .

Table 2: Sample reference summaries with highlights on keywords extracted by our heuristic.

3 ROUGE-K

We present ROUGE-K, an extension of ROUGE that exclusively focuses on essential words in summaries. Its core idea is simple: ROUGE-K assesses the proportion of keywords from the reference summary that are included in the candidate summary. We compute ROUGE-K as:

$$R-K = \frac{\text{Count}(kws \cap n\text{-grams})}{\text{Count}(kws)}$$

where *kws* is a set of pre-defined keywords and *n-grams* is a target hypothesis. This provides a direct understanding of how well system summaries contain essential pieces of information. ROUGE-K is essentially a recall-oriented metric since it computes *coverage* of keywords. While it is possible to complement this formula with another one to compute precision, this would give the proportion of keywords in the candidate summary. However, this metric would indicate how good the system is at extracting keywords, not its summarization capabilities, i.e., one could have a summary consisting only of keywords but only marginally overlapping with the reference summary.

Keyword extraction. An essential prerequisite of ROUGE-K is the availability of keywords. Ideally, we would like these keywords to be available for *any* summarization corpus to enable the wide applicability of our metric. A solution is thus to extract keywords from reference summaries heuristically. Nan et al. (2021), for instance, use words detected by a named entity recognition (NER) model to evaluate entity-level factual consistency in summaries. However, (1) not all keywords are named entities, (2) NER models accurate enough to be used for evaluation are not available for all domains (e.g., scholarly documents), (3) the accuracy of NER models for documents in summarization datasets is unknown.

	R-1	R-2	R-L	BS	R-K
SciTLDR	61.11	58.89	60.00	57.78	72.22
XSum	63.73	59.80	56.86	62.75	70.59

Table 3: Agreement ratios (%) of each metric and human annotator on summary relevance, computed as the proportion of documents for which a given metric gives the highest score to the summary judged as most relevant from humans.

In this paper, we present a simple and interpretable way to extract keywords. We define keywords as *the n-grams used in multiple reference summaries*, assuming that words used in multiple human-written summaries for the same document repeatedly should be included in system summaries as well. First, we tokenize and lowercase the reference summaries, extract n-grams, and then remove stopwords from the extracted n-grams. Next, we compare n-grams from multiple references and extract those that appear in multiple references. To capture multi-word keywords, the extraction process starts from 10-grams to unigrams. *When there is only one reference summary available, the corresponding title is used as a proxy reference* which is known to contain key information (Koto et al., 2022; Cachola et al., 2020). Table 2 shows examples of keywords extracted by our heuristic. We benchmark our heuristic against TF-IDF (Salton and Buckley, 1988) and TextRank (Mihalcea and Tarau, 2004). To this end, we take the first 100 samples of the SciTLDR development data and for each summary, we extract the same amount of keywords as the one from our method (i.e., we keep the recall level fixed). We then quantify for each method a) the average number of wrong keywords per summary and b) the overall false discovery rate FDR (for both, lower is better) – our hunch is that for humans, it is easier to judge whether something is not a keyword (i.e., a word is unquestionably not being essential to convey the information), as opposed to being one. In both cases, our heuristic achieves the best performance: 0.64 vs. 0.85 and 0.94 on average wrong extractions per summary and 0.13 vs. 0.16 and 0.21 FDR when compared against TF-IDF and TextRank, respectively.

Agreement with human judgments. We now perform a manual evaluation to test how well ROUGE-K aligns with human judgements on rating the relevance of summaries (we follow Fabbri et al. (2021) and define ‘relevance’ as the *selection*

of important content from the source. We focus for manual evaluation on relevance only (as opposed to, e.g., Fabbri et al. (2021) considering three other aspects) because the purpose of ROUGE-K is to quantify how well summaries include essential words, and thus preserve important, i.e., relevant content, as opposed to, e.g., ROUGE taking into account style aspects.

Our dataset consists of pairs of summaries generated using different instances of the same model (BART), trained on each of SciTLDR and XSum with different random seeds. To avoid ties, we select a sample of 92 and 100, respectively from SciTLDR and XSum, summary pairs where the two models assign a different ROUGE-K score to each summary. We then ask four annotators from our CS graduate course to compare the summaries and rank them (i.e., label the best one among the two). We finally compute how often each evaluation metric assigns higher scores to the summaries preferred by the annotators. Results are shown in Table 3. In line with Fabbri et al. (2021), R-1 shows higher agreement than R-2 and R-L, and BERTScore marks a marginally lower score than ROUGE-1. Finally, ROUGE-K shows much higher agreement, indicating its strong ability to detect human-preferable summarization models.

Benchmarking BART with ROUGE-K. As a showcase of ROUGE-K, we evaluate BART (Lewis et al., 2020), a strong transformer-based generative language model on three different datasets: SciTLDR (Cachola et al., 2020), XSum (Narayan et al., 2018) and ScisummNet (Yasunaga et al., 2019). These cover different summarization tasks – i.e., extreme (SciTLDR, XSum) vs. non-extreme (ScisummNet) – as well as different domains – i.e., scholarly documents (SciTLDR, ScisummNet) vs. news (XSum). Datasets details are shown in Table 4. BART models are fine-tuned on the training set and early stopping is performed using validation data, and finally evaluated on the test split using the traditional ROUGE metrics and our ROUGE-K. Table 5 shows the results. Each score is an average of ten and three different random seeds, respectively, for SciTLDR and XSum/ScisummNet (a larger number of seeds is used for SciTLDR to obtain stable scores on its relatively small test dataset). Although one would consider the scores achieved by BART on ROUGE-1/-2/-L to be high, it only reaches 41.36% and 56.14% on ROUGE-K. In other words, a strong baseline model fails to in-

clude half of the essential n-grams in its summaries, unveiling a critical limitation previously missed by standard metrics.

Correlation with summary lengths. Since ROUGE-K is recall-oriented, it can potentially favor longer summaries, i.e., as suggested by the overall absolute scores obtained by BART in Table 5 on non-extreme summarization with ScisummNet data. To quantify this, we compute Pearson correlations between the number of words in summaries generated by BART and different evaluation scores. As Table 6 shows, ROUGE-K scores have marginally higher correlations with summary lengths than other ROUGE (F1) metrics, although they all are relatively weak, ranging from -0.07 to 0.17 on SciTLDR and even lower for XSum. These results are different from those from Sun et al. (2019), arguably because SciTLDR and XSum are extreme summarization datasets. On non-extreme summarization (ScisummNet), the results align instead with previous findings. However, we observe the same level of moderate correlation with the summary length between vanilla ROUGE and ROUGE-K.

Correlation with existing metrics. To better understand the relationship between ROUGE-K and other existing metrics, we perform an additional correlation analysis (Table 7). R-1 (avg) computes a R-1 for each reference given a sample and takes the average while R-1 (max) takes only the largest score. R-1 (avg) and R-1 (max) are the same for XSum and ScisummNet because there is only one reference summary in this dataset. The results indicate only a moderate strength of association between ROUGE-K and existing metrics, thus providing evidence that our metric can partially complement other metrics.

4 Importance-guided summarization

We next propose four ways to incorporate a soft guiding signal into BART to enforce the inclusion of keywords into the generated summaries.

Re-weighted encoding (RwEnc). The first approach is to modify the representations within the model with TF-IDF scores. Concretely, we compute the attention matrix in transformer layers as:

$$\text{attention matrix} = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}} + T\right)V$$

Dataset	Documents		Summaries			Extracted keywords	
	# documents (train/val/test)	# words per doc on avg.	# words per summary on avg.	compress. ratio	# references on avg. (train/val/test)	# keywords on avg. (train/val/test)	average lengths (train/val/test)
SciTLDR	1,992 / 619 / 618	5,000	21.00	238.10	2.0 / 3.3 / 4.2	1.9 / 4.2 / 5.2	1.7 / 1.5 / 1.5
XSum	204K / 11K / 11K	431	23.26	18.53	2.0 / 2.0 / 2.0	2.9 / 2.9 / 2.9	1.5 / 1.5 / 1.5
ScisummNet	750 / 92 / 91	4,700	167.49	28.06	2.0 / 2.0 / 2.0	2.8 / 3.0 / 2.6	1.7 / 1.6 / 1.6

Table 4: Statistic of datasets and extracted keywords.

	R-1	R-2	R-L	R-K
SciTLDR	43.93	22.31	36.58	41.36
XSum	44.43	21.00	35.94	56.14
ScisummNet	50.75	47.80	49.73	68.95

Table 5: BART performance evaluated by ROUGE-1/2/-L and our ROUGE-K.

	R-1	R-2	R-L	R-K
SciTLDR	-0.102	-0.070	-0.154	0.167
XSum	-0.003	-0.037	-0.075	0.057
ScisummNet	0.356	0.435	0.392	0.402

Table 6: Pearson Correlation between the number of words in summaries and evaluation metrics.

The first term within the softmax function is from the original transformer (Vaswani et al., 2017) where Q and K are query and key matrices respectively, and we introduce the second term T which is a matrix of TF-IDF scores over the input text. This enhances the model to propagate higher values for the important words to the upper layers. We apply this modification to the 0-, 4-, 8-th encoder layers, empirically selected on the dev data.

Re-weighted generation (RwGen). The second solution operates in the token selection phase. At each generation step, BART computes a probabilistic distribution over its vocabulary for the next token to produce. We modify this distribution by summing TF-IDF scores so that the words with higher scores are more likely to be selected:

$$\text{score}(y_{w_t} | w_{<t}, X, T)' = (1 - \lambda) * \text{score}(y_{w_t} | w_{<t}, X) + \lambda * T$$

where score is a fine-tuned BART that takes previously generated words ($w_{<t}$) and the source document (X), and predicts scores which are further transformed to the probability for the next token at the time step t by a softmax function. We introduce the second term (T) which is a vector filled with TF-IDF values for the source document. λ is a

	R-1 (avg)	R-1 (max)	BS
SciTLDR	0.510	0.434	0.383
XSum	0.318	0.318	0.237
ScisummNet	0.288	0.288	0.413

Table 7: Pearson Correlation between ROUGE-K and ROUGE-1 average, ROUGE-1 max and BERTScore.

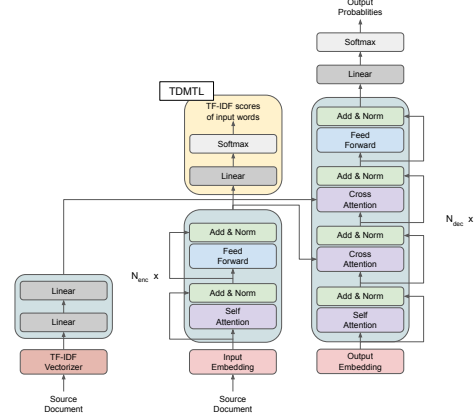


Figure 1: Overview of our TDSum model.

hyperparameter with which we control how much we shift the distribution from vanilla BART.

Multi-Task Learning with TF-IDF (TDMTL). Another solution is to modify the objective function to ask the model to predict TF-IDF scores in parallel with generating summaries. For this, we compute the mean squared error as loss for TF-IDF score prediction L_{tfidf} and the standard cross entropy loss for summarization L_{sum} . The final loss we minimize is the linear interpolation of the two task-specific losses: $(1 - \lambda)L_{sum} + \lambda L_{tfidf}$.

Injecting TF-IDF into the decoder (TDSum). Our last approach is inspired by Dou et al. (2021). Since their approach requires an explicit guidance signal (e.g., keywords), it uses additional models for keyword extraction leading to a drastic increase in computational costs. Instead, we propose to use light-weight TF-IDF scores as shown in Figure 1. TDSum equips two linear layers to process

TF-IDF scores for words in input documents and uses resulting word importance features in newly introduced cross-attention layers in each decoder layer to guide the model towards keyword-oriented summary generation. We train this model with the aforementioned TDMTL loss.

5 Experiments and results

5.1 Experimental setup

Datasets. We experiment on different domains and summarization tasks using SciTLDR (Cachola et al., 2020), XSum (Narayan et al., 2018) and ScisummNet (Yasunaga et al., 2019)

Baselines. We compare our models with three abstractive and one extractive summarizer:

- **BART** (Lewis et al., 2020) is a transformer-based generative language model, pre-trained with denoising objective function.
- **GSum** (Dou et al., 2021) is an extension of BART with additional parameters for processing textual guidance signals. Here, we input overlapping keywords, extracted as explained in Section 3.
- **MTL** (Saito et al., 2020) performs multitask training to predict keywords in source documents in addition to the summarization objective (we use our extracted keywords from Section 3).
- **PreSumm** (Liu and Lapata, 2019) is an extractive summarization model based on BERT (Devlin et al., 2019).

Hyperparameter tuning. We perform a grid search for each dataset and model using the development data and ROUGE-1 as a reference. We test for learning rate $\in \{1e-05, 2e-05, 3e-05\}$, gradient accumulation $\in \{4, 8\}$, number of beam search $\in \{2, 3\}$ and repetition penalty rate $\in \{0.8, 1.0\}$. We also explore $\lambda \in \{0.1, 0.2, 0.3\}$ for MTL and TDMTL and $\lambda \in \{30, 50\}$ for RwGen. During hyperparameter search, we use one random seed. The final reported results on the test data are the averaged performance over models fine-tuned with different random seeds. We use ten seeds for SciTLDR and ScisummNet and three for XSum. Our experiments are performed on RTX A6000 and utilize the implementation by Deutsch and Roth (2020) to compute ROUGE-1/2/L.

5.2 Results and discussion

We organize the discussion of our results around the following research questions:

- **RQ1:** Can models that incorporate TF-IDF scores increase the number of keywords in the summaries without degrading ROUGE scores?
- **RQ2:** Which kinds of keywords do models find hard to include in summaries?

RQ1: TF-IDF as guidance. We present our main results in Table 8. On the SciTLDR dataset, BART marginally outperforms two other baseline models on ROUGE-1/2/L. However, MTL performs the best on the ROUGE-K metric, thus showing its effectiveness of explicit training guidance. As reported by previous works (Cachola et al., 2020; Narayan et al., 2018), an extractive model considerably underperforms abstractive models on all the metrics in extreme summarization since it suffers from merging information across multiple input sentences into its outputs. Because keywords are also scattered over multiple sentences, it fails to include most keywords. Three out of four of our newly introduced TF-IDF-equipped models outperform vanilla BART on keyword inclusion, and TDSum significantly outperforms all the baselines on ROUGE-K while keeping its ROUGE scores on par with BART. TDMTL follows the same training procedure as MTL and learns to predict TF-IDF instead of keywords. While results still improve over BART, our results show that using hard signals (i.e., keywords) is preferable. RwGen is simple and fast to train, yet it includes more keywords than BART. On XSum, BART outperforms other baseline models on ROUGE-K. However, the MTL model exceeds other traditional ROUGE metrics showing that our metric can shed light on an aspect that other metrics cannot capture. Among our proposed methods, TDMTL performs well akin to the results on SciTLDR, outperforming BART on traditional ROUGE metrics, although the BART model still outperforms models with TF-IDF extensions on ROUGE-K.

We see similar trends for non-extreme summarization on ScisummNet, where our proposed models are on par (sometimes outperform) with baseline models on ROUGE metrics. All proposed methods outperform all baseline models on ROUGE-K, indicating that, even for longer summaries, TF-IDF can enhance keyword inclusion. One significant difference is that the extractive model (PreSumm) performs better than abstractive models on ROUGE-1/2. We speculate this is due to much longer output summaries (181.88 words on average for PreSumm vs. 48.01 for BART).

Model	SciTLDR				XSum				ScisummNet			
	R-1	R-2	R-L	R-K	R-1	R-2	R-L	R-K	R-1	R-2	R-L	R-K
BART	43.93	22.31	36.58	41.36	44.43	21.00	35.94	56.14	50.75	47.80	49.73	68.95
GSum	43.65	22.09	36.50	41.00	43.86	20.47	35.60	52.85	24.37	21.11	23.35	43.36
MTL	43.82	22.24	36.29	42.83	44.50	21.05	36.10	56.06	50.75	47.81	49.73	68.81
PreSumm	30.43	11.36	24.08	25.06	22.16	4.13	15.91	24.67	60.58	49.15	46.22	68.85
Std (1-4)	5.79	4.70	5.36	7.25	9.57	7.24	8.65	13.12	13.45	11.77	11.01	11.05
RwEnc	43.98	22.39	36.68	41.03	44.42	20.93	36.07	55.58	50.75	47.92	49.89	69.40
RwGen	43.96	22.35	36.59	41.60	44.57	21.04	36.09	56.03	50.38	47.78	49.54	69.19
TDMTL	44.08	22.48	36.75	41.85	44.50	21.05	36.10	56.06	50.64	47.75	49.67	69.56
TDSum	43.55	21.74	35.82	43.04	44.13	20.95	35.57	55.39	50.50	47.63	49.55	69.43
Std (all)	4.44	3.60	4.10	5.59	7.34	5.56	6.62	10.23	9.72	8.90	8.62	8.54

Table 8: Results on SciTLDR, XSum, and ScisummNet. Best results per metric are **bolded**. Scores with underline indicate that they significantly outperform all baseline models. We test for statistical significance using the Wilcoxon signed-rank test with $\alpha = 0.05$ (Dror et al., 2018).

Model	Dataset	IN-SRC	OUT-SRC
BART	SciTLDR	54.53	0.92
	XSum	73.78	30.34
	ScisummNet	75.21	8.17
TDSum	SciTLDR	56.75	1.37
	XSum	66.61	26.66
	ScisummNet	75.04	14.41

Table 9: ROUGE-K scores on keywords seen (IN-SRC) vs. unseen (OUT-SRC) in source documents.

RQ2: In search of missing keywords. We next focus on studying the relationship between a few specific characteristics of keywords and ROUGE-K scores. First, we look at *whether models can better include keywords if they appear in source documents*. We do this by splitting a list of pre-defined keywords into two lists, (1) an IN-SRC keyword list where all the keywords appear in the source documents, (2) an OUT-SRC keyword list where keywords cannot be found in the source documents, and then evaluate a model with ROUGE-K using each list. As Table 9 shows, on both datasets and models, ROUGE-K with OUT-SRC keywords is notably lower than IN-SRC ROUGE-K showing that when keywords are not in the source texts it is challenging for models to include them in summaries.

We next investigate whether there is a correlation between ROUGE-K and keyword length, that is, *whether longer keywords are more difficult to include*. Figure 2 shows that although there is one exceptional case ($N = 7$), ROUGE-K scores consistently decrease as keywords become longer, indicating the difficulty of including longer keywords in summaries. Another finding in this analysis is that while BART outperforms TDSum on XSum




Figure 2: ROUGE-K and keyword length.

when keywords with all lengths are used when the n-grams are longer ($N \geq 5$), TDSum starts to surpass BART on ROUGE-K.

LLMs on ROUGE-K To shed light on the summarization behaviour of recent large language models (LLMs), we evaluate an open-source model, namely instruction fine-tuned Llama 2 (Touvron et al., 2023) in two different sizes. The prompt used in our experiments is “Generate a summary of the following document in one sentence”. Due to our limited computational resources, we cut off inputs and outputs at 512 and 128 tokens, respectively, and also truncate all the sentences after the first one in the generated summary, if longer. Results are shown in Table 10. While they perform remarkably well on both ROUGE and our ROUGE-K given that inferences are performed in zero-shot manner, we observe that more than half of the keywords are missing, calling for better prompting strategies. On traditional ROUGE scores, which consider the

Dataset	Llama2 7B				Llama2 13B			
	R-1	R-2	R-L	R-K	R-1	R-2	R-L	R-K
SciTLDR	35.9	13.9	26.7	44.1	36.3	14.6	27.8	44.4
XSum	21.8	5.6	15.5	34.7	22.0	5.6	15.9	35.9
SciNet	46.1	24.1	32.5	64.6	46.7	25.3	33.0	69.5

Table 10: Results with Llama2 7B and 13B.

Generated summary (a)
Input: Deep convolutional neural networks (CNNs) are known to be robust against label noise on extensive datasets. However, at the same time, CNNs are [...] (https://openreview.net/forum?id=H1xmqiAqFm)
TLDR: The authors challenge the CNNs robustness to label noise using ImageNet 1k tree of WordNet.
BART: We investigate the behavior of CNNs under class-dependently simulated label noise , which is generated based on the conceptual distance between classes of a large dataset.
TDSum: We show that CNNs are more robust to class-dependent label noise than class-independent label noise, which is generated based on the conceptual distance between classes of a large dataset.

Generated summary (b)
Input: We explore ways of incorporating bilingual dictionaries to enable semi-supervised neural machine [...] (https://arxiv.org/abs/2004.02071)
TLDR: We use bilingual dictionaries for data augmentation for neural machine translation .
BART: We propose a simple data augmentation technique to address both this shortcoming.
TDSum: We propose a simple data augmentation technique to enable semi-supervised neural machine translation .

Table 11: Summaries generated by our models.

words related to writing style, Llama performs noticeably worse than fine-tuned models because fine-tuning can help models learn reference styles from the dataset. However, results are comparable on ROUGE-K. This indicates that ROUGE-K can better account for the high quality of LLMs, despite different styles between generated and reference summaries, which has been noticed when evaluating zero-shot models (Goyal et al., 2022).

Model distinguishability. Most if not all recent summarization papers perform evaluations using multiple ROUGE metrics, yet the gap between systems is very small, making it hard to distinguish models’ performance. Inspired by work from Xiao et al. (2022) on characterizing the distinguishability of datasets, we compute the standard deviation of scores from our models (cf. Section 5.1) for each metric, to see the distinguishability of ROUGE variants (larger deviation means higher distinguishability). As shown in Table 8, ROUGE-K achieves the

highest standard deviation among other ROUGE metrics for two extreme summarization datasets, i.e., it differentiates models better when summaries are required to be very short. We highlight this by means of the sample output shown in Table 11. In (a), BART fails to include one of the keywords ‘robust’ which is necessary to convey the purpose of the paper. In (b), the summary by BART does not mention the task the paper worked on (in this case, neural machine translation) while TDSum successfully includes it.

6 Related work

In the context of factual consistency evaluation (Kryscinski et al., 2020; Scialom et al., 2021; Fabri et al., 2022), Nan et al. (2021) propose to use a NER model to detect hallucinated named entities in summaries. While their approach also focuses on specific words in summaries as our ROUGE-K, it is limited because (1) not all critical information consists of named entities, (2) strong NER models are not available in many domains, and (3) NER performance is unknown for summarization datasets. Ng and Abrecht (2015) and Zhang et al. (2020) propose to use vector representations to compute semantic similarity between reference and candidate summaries. Eyal et al. (2019) instead propose to use a question-answering system to assess the summary quality. While these methods can exploit semantic knowledge stored in parameters in large models, as a side-effect, they introduce ‘blackboxness’ that hinders transparent model development. In contrast, we take a ‘bottom-up’ approach by proposing to focus on keyword availability.

7 Conclusion

In this paper, we proposed ROUGE-K, an extension of ROUGE to quantify how summary-relevant keywords are included in summaries. Using ROUGE-K, we showed human annotators prefer summaries with more keywords and how models often miss several essential keywords in their output. In a variety of experiments using the baseline provided by a large pre-trained language model (BART) we showed how ROUGE-K only moderately correlates with ROUGE and BERTScore, thus indicating that it can complement them, and correlates with the length of the generated summaries on a par with ROUGE F1 and BERTScore, despite being a recall-oriented metric. Finally, we proposed four ways to guide BART to include more keywords in its sum-

maries. We plan in future work to further test our metric’s applicability across different domains and languages, e.g., by relying on WikiLingua (Ladhak et al., 2020) and X-SciTLDR (Takeshita et al., 2022).

8 Limitations

This work has the following limitations: (1) Our new evaluation metric, ROUGE-K, uses a heuristic to extract keywords automatically. Although it enables to obtain better and more comprehensive keywords compared to other existing methods, some nonessential words are still included thus can bring some noise into the evaluation. (2) ROUGE-K does not take the context of keywords into consideration which leaves the possibility open that generated summaries with keywords still convey the meaning of keywords wrongly. (3) Like traditional ROUGE scores, ROUGE-K is based on hard string match, which cannot compensate for the semantics of, e.g., (near-)synonyms and paraphrases.

Acknowledgements

The work presented in this paper is funded by the German Research Foundation (DFG) under the VADIS (PO 1900/5-1; EC 477/7-1) project, and also supported by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1597-1 FUGG.

References

- Mousumi Akter, Naman Bansal, and Shubhra Kanti Karmaker. 2022. [Revisiting automatic evaluation of extractive summarization task: Can we do better than ROUGE?](#) In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1547–1560, Dublin, Ireland. Association for Computational Linguistics.
- Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel Weld. 2020. [TLDR: Extreme summarization of scientific documents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4766–4777, Online. Association for Computational Linguistics.
- 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: Human Language Technologies*, pages 93–98, San Diego, California. Association for Computational Linguistics.
- Daniel Deutsch and Dan Roth. 2020. [SacreROUGE: An open-source library for using and developing summarization evaluation metrics](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 120–125, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2021. [GSum: A general framework for guided neural abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4830–4842, Online. Association for Computational Linguistics.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. [The hitchhiker’s guide to testing statistical significance in natural language processing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Matan Eyal, Tal Baumel, and Michael Elhadad. 2019. [Question answering as an automatic evaluation metric for news article summarization](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3938–3948, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [QAFactEval: Improved QA-based factual consistency evaluation for summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2587–2601, Seattle, United States. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating Summarization Evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. [News Summarization and Evaluation in the Era of GPT-3](#). ArXiv:2209.12356 [cs].
- 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*, volume 28. Curran Associates, Inc.
- Fajri Koto, Timothy Baldwin, and Jey Han Lau. 2022. **LipKey: A large-scale news dataset for absent keyphrases generation and abstractive summarization.** In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3427–3437, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. **Evaluating the factual consistency of abstractive text summarization.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. **WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization.** In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries.** In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. **Text summarization with pretrained encoders.** 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 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. **TextRank: Bringing order into text.** In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Feng Nan, Ramesh Nallapati, Zhiguo Wang, Cicero Nogueira dos Santos, Henghui Zhu, Dejiao Zhang, Kathleen McKeown, and Bing Xiang. 2021. **Entity-level factual consistency of abstractive text summarization.** In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2727–2733, Online. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. **Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Jun-Ping Ng and Viktoria Abrecht. 2015. **Better summarization evaluation with word embeddings for ROUGE.** In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1925–1930, Lisbon, Portugal. Association for Computational Linguistics.
- 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, Lisbon, Portugal. Association for Computational Linguistics.
- Itsumi Saito, Kyosuke Nishida, Kosuke Nishida, and Junji Tomita. 2020. **Abstractive summarization with combination of pre-trained sequence-to-sequence and saliency models.** *CoRR*, abs/2003.13028.
- Gerard Salton and Chris Buckley. 1988. **Term-weighting approaches in automatic text retrieval.** *Inf. Process. Manag.*, 24(5):513–523.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. **QuestEval: Summarization asks for fact-based evaluation.** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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*, pages 21–29, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sotaro Takeshita, Tommaso Green, Niklas Friedrich, Kai Eckert, and Simone Paolo Ponzetto. 2022. **X-SCITLDR: cross-lingual extreme summarization of scholarly documents.** In *JCDL ’22: The ACM/IEEE Joint Conference on Digital Libraries in 2022, Cologne, Germany, June 20 - 24, 2022*, page 4. ACM.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,

Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008, Long Beach, CA, USA. Curran Associates, Inc.

Yang Xiao, Jinlan Fu, See-Kiong Ng, and Pengfei Liu. 2022. [Are All the Datasets in Benchmark Necessary? A Pilot Study of Dataset Evaluation for Text Classification](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2930–2941, Seattle, United States. Association for Computational Linguistics.

Michihiro Yasunaga, Jungo Kasai, Rui Zhang, Alexander R Fabbri, Irene Li, Dan Friedman, and Dragomir R Radev. 2019. [ScisummNet: A large annotated corpus and content-impact models for scientific paper summarization with citation networks](#). *Proc. Conf. AAAI Artif. Intell.*, 33:7386–7393. Publisher: Association for the Advancement of Artificial Intelligence (AAAI).

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.