

Evaluating Small Language Models for News Summarization: Implications and Factors Influencing Performance

Borui Xu

Shandong University
boruixu@mail.sdu.edu.cn

Yao Chen*

National University of Singapore
yaochen@nus.edu.sg

Zeyi Wen

HKUST(GZ) & HKUST
wenzeyi@ust.hk

Weiguo Liu*

Shandong University
weiguo.liu@sdu.edu.cn

Bingsheng He

National University of Singapore
hebs@comp.nus.edu.sg

Abstract

The increasing demand for efficient summarization tools in resource-constrained environments highlights the need for effective solutions. While large language models (LLMs) deliver superior summarization quality, their high computational resource requirements limit practical use applications. In contrast, small language models (SLMs) present a more accessible alternative, capable of real-time summarization on edge devices. However, their summarization capabilities and comparative performance against LLMs remain underexplored. This paper addresses this gap by presenting a comprehensive evaluation of 19 SLMs for news summarization across 2,000 news samples, focusing on relevance, coherence, factual consistency, and summary length. Our findings reveal significant variations in SLM performance, with top-performing models such as Phi3-Mini and Llama3.2-3B-Ins achieving results comparable to those of 70B LLMs while generating more concise summaries. Notably, SLMs are better suited for simple prompts, as overly complex prompts may lead to a decline in summary quality. Additionally, our analysis indicates that instruction tuning does not consistently enhance the news summarization capabilities of SLMs. This research not only contributes to the understanding of SLMs but also provides practical insights for researchers seeking efficient summarization solutions that balance performance and resource use.

1 Introduction

Automatic text summarization is a fundamental challenge in Natural Language Processing (NLP) that plays a crucial role in various applications, including search engine optimization (Matošević, 2019), financial forecasting (Li et al., 2023), and public sentiment analysis (Chakraborty et al., 2019). The ability to distill large volumes of information into concise summaries is particularly

*Corresponding authors




Figure 1: Variation of SLMs in size and BertScore for news summarization.

important in today’s fast-paced digital environment, where users often seek quick insights from news articles. Recent advancements in large language models (LLMs) such as GPT-3 (Brown et al., 2020), OPT (Zhang et al., 2022), and Llama2 (Touvron et al., 2023) have significantly enhanced the quality and coherence of generated summaries compared to traditional models (Luhn, 1958; Dong et al., 2018; Zhang et al., 2018; Raffel et al., 2020; Genest and Lapalme, 2012). However, the hundreds of billions of parameters in these LLMs require substantial computational and storage resources. For example, deploying the Llama2-70B model at FP16 precision typically necessitates two Nvidia A100 GPUs with 80GB of memory each. This means that only organizations with significant computing power can offer LLM services, raising concerns about service stability and data privacy.

As a result, general-purpose small language models (SLMs) have emerged as a potential alternative, as shown in Figure 1. They share the same decoder-only architecture as LLMs but have fewer than 4B parameters¹ (Microsoft Azure, 2024; Yang et al., 2024; Abdin et al., 2024; Meta AI, 2024). Same as

¹Currently, there is no clear definition of SLMs, but they generally have fewer than 4 billion parameters.

the LLMs, they generate specified content based on prompts. SLMs are designed to run efficiently on edge devices like smartphones and personal computers, making them more accessible. This accessibility makes SLMs an attractive option for applications requiring quick, stable, and low-cost solutions that also protect user privacy, such as summarizing news articles for mobile users or enhancing data processing privacy in sensitive environments.

Despite the promising capabilities of SLMs, their specific performance in the domain of news summarization remains underexplored. Previous studies, such as Tiny Titans (Fu et al., 2024), have evaluated specifically fine-tuned SLMs in meeting summarization. The adopted ROUGE metric (Lin, 2004) only reflects the quality in terms of relevance. In addition, broader evaluations of SLMs across various tasks have been evaluated but omitted news summarization (Lu et al., 2024), highlighting the gap in understanding SLM in this area.

This paper aims to address these gaps by conducting a comprehensive evaluation of pre-trained general-purpose SLMs in news summarization, a process that abstracts news articles into shorter versions while preserving key information. We seek to answer the following research questions:

First, **How do different SLMs compare to each other, and how do they compare to LLMs?** We benchmark 19 SLMs on 2,000 news samples². Our findings reveal that the news summarization capabilities of different SLMs vary significantly as shown in Figure 1. And the best-performing SLMs can generate news summaries comparable in quality to those produced by LLMs while producing shorter summaries. Notably, Phi3-Mini and Llama3.2-3B-Ins emerge as the top performers.

Second, **What is the impact of prompt complexity on the performance of SLMs?** We compare the impact of prompts with varying levels of detail on the quality of the summaries. Our findings indicate that prompt engineering has a limited impact on enhancing summarization abilities, with even complex prompts potentially degrading model performance. Thus, simple prompts suffice in practical applications.

Third, **How does instruction tuning affect the summarization capabilities of different SLMs?** We find that the effects of instruction tuning on SLMs vary significantly. While the Llama3.2 se-

²The codes and results are available at https://github.com/Xtra-Computing/SLM_Summary_Benchmark

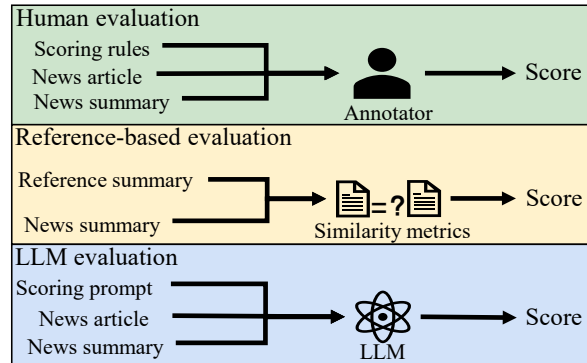


Figure 2: Comparison of text summarization evaluation.

ries models show substantial improvements after instruction tuning, which aligns with existing works (Zhang et al., 2024b), models like Qwen2 and InternLM2 exhibit minimal changes.

The rest of the paper is organized as follows: Section 2 covers the background and related work, Section 3 introduces reference summary generation, Section 4 illustrates the benchmark design, Section 5 presents the benchmarking results of SLMs, Section 6 analyzes the influencing factors on SLM summarization, and Section 7 concludes our findings.

2 Background and Related Work

2.1 Small Language Model

Small language models (SLMs) are characterized by their reduced parameter count and computational requirements compared to larger LLMs. These models can be general-purpose or specialized, and this paper focuses on the evaluation of pre-trained general-purpose SLMs, which will be referred to simply as SLMs throughout the remainder of the paper. Structurally, SLMs and LLMs share a common architecture, consisting of stacked decoder-only layers from the transformer framework (Vaswani et al., 2017), generating outputs in an autoregressive manner. SLMs typically have fewer decoder layers, attention heads, and smaller hidden dimensions, and they are trained on smaller, high-quality datasets (Lu et al., 2024). While there is no clear definition for the parameter count that qualifies a model as an SLM, models that can operate on consumer-grade devices are generally considered to fall within this category. This paper focuses on models smaller than 8GB in FP16 precision, corresponding to fewer than 4B parameters.

2.2 Text Summarization and Its Evaluation

Text summarization is the process of compressing a large volume of text into a shorter version while retaining its main ideas and essential information. Traditional text summarization models usually take a single news article as input and output a summary. However, SLMs require both a prompt and an article to generate a summary.

The evaluation methods of text summarization can be categorized into three categories, as shown in Figure 2: human evaluation, reference-based evaluation, and LLM evaluation.

Human evaluation is an intuitive and effective method, involving the hiring of annotators to score the generated text. For instance, Huang et al. (2020) employed human evaluation to evaluate 10 representative summarization models, revealing that extractive summary methods often performed better in human evaluation. The Summeval benchmark (Fabbri et al., 2021) utilized human evaluation methods to evaluate 23 traditional models and found Pegasus (Zhang et al., 2020a) performed best. Pu et al. (2023a) used human evaluation to assess LLMs and found that they were more favored by evaluators. Zhang et al. (2024b) also used human evaluation to evaluate the LLMs, discovering that instruction tuning had a greater impact on performance than model size. INSTRUSUM (Liu et al., 2024) evaluated instruction controllable summarization of LLMs by annotators and found them still unsatisfactory in summarizing text based on complex instructions.

Reference-based evaluation is a commonly used method. It scores generated text by calculating its similarity to reference texts using various metrics, and many such metrics have been proposed (Lin, 2004; Papineni et al., 2002; Sellam et al., 2020; Zhang et al., 2020b; Yuan et al., 2021). In addition to human evaluation, Huang et al. (2020) and Fabbri et al. (2021) also used many reference-based metrics to evaluate models and found metrics like Rouge and BertScore are very close to human evaluations. Maynez et al. (2023) used the Rouge metric to measure text similarity. Tiny Titans (Fu et al., 2024) assessed the performance of small models in meeting summarization and found FLAN-T5 performed best.

LLM evaluation is explored by recent studies (Goyal et al., 2022; Chiang and Lee, 2023; Gao et al., 2023b; Liu et al., 2024; Wang et al., 2023). This approach involves guiding LLMs with

prompts and providing both the summary and the original article for scoring. However, Shen et al. (2023) indicated that LLMs may exhibit biases when evaluating different models, warranting further research to ensure fairness.

In conclusion, existing evaluations primarily evaluate the news summarization capabilities of traditional models or LLMs, leaving a gap in the evaluation of SLMs.

3 LLM-Augmented Reference-Based Evaluation

As the reference-based method offers higher efficiency, better reproducibility, and lower cost compared to human and LLM evaluation (Fabbri et al., 2021; Zhang et al., 2024b), we use the reference-based method for large-scale SLM evaluation. However, some existing research (Tejaswin et al., 2021; Fabbri et al., 2021) indicates that the poor quality of heuristic reference summaries in current datasets weakens their correlation with human preferences. Moreover, Zhang et al. (2024b) find that using high-quality summaries as references can significantly improve alignment with human preferences. Thus, we use LLM-generated summaries as references in our SLM evaluation instead of relying on original dataset references.

LLM-generated reference summaries offer two key advantages. First, they have higher quality than original heuristic summaries; Zhang et al. (2024b) find LLMs averaged 4.5 on a five-point scale, while original references scored only 3.6. Additionally, Pu et al. (2023a) find that LLM-generated summaries are preferred to those written by freelance authors by up to 84% of the time. Second, LLMs produce high-quality summaries quickly, generating hundreds in a few hours, which benefits large-scale news evaluations. In contrast, a freelance writer typically takes about 15 minutes for a single summary (Zhang et al., 2024b).

We also verify whether the LLM-generated reference summaries can improve the effectiveness of reference-based evaluation on three human evaluation datasets (Zhang et al., 2024b; Fabbri et al., 2021) in Table 1. These datasets include summaries generated by various models as well as human ratings. We compare the correlation of three popular similarity metrics with human scoring using different references. BertScore and BLEURT calculate similarity based on fine-tuned language models, while RougeL measures similarity through textual

<p>Prompt 1</p> <p>News: {news} Summarize the news in two sentences. Summary:</p>
<p>Prompt 2</p> <p>You are a helpful summary assistant. You can help users summarize news in two sentences. News: {news} Summarize the news in two sentences. Summary:</p>
<p>Prompt 3</p> <p>News: {news} Summarize the news in two sentences. Your summary should:</p> <ol style="list-style-type: none"> 1. Capture the main points of the article 2. Be concise and informative 3. Use clear and simple language 4. Avoid unnecessary details or opinions <p>Summary:</p>

Figure 3: The prompt templates for the language model to generate summaries.

overlap. We use prompt 2 in Figure 3 as the input of the LLM. For all datasets, we consistently use two sentences to summarize the article. LLM-generated summaries significantly improve consistency than original references. BertScore performs best in both the relevance and coherence aspects.

4 Benchmark Design

In this section, we first introduce the dataset for news summarization evaluation, followed by the SLM details, evaluation metrics, and reference summary generation.

4.1 News Article Selection

All experiments are conducted on the following four datasets: CNN/DM (Hermann et al., 2015), XSum (Narayan et al., 2018), Newsroom (Grusky et al., 2018), and BBC2024. The first three datasets have been widely used for verifying model text summarization (Goyal et al., 2022; Maynez et al., 2023; Zhang et al., 2024b). However, their older data may overlap with the training datasets of some models. To address this, we create BBC2024, featuring news articles from January to March 2024 on the BBC website³. Based on prior research, 500 samples per dataset are sufficient to differentiate the models (Maynez et al., 2023). We randomly select 500 samples from each of the four test sets and construct 2000 samples in total. Each sample contains 500 to 1500 tokens according to the Qwen1.5-72B-Chat tokenizer.

³<http://dracos.co.uk/made/bbc-news-archive/>

4.2 Model Selection

We select 21 popular models with parameter sizes not exceeding 4 billion for benchmarking, including 19 SLMs and 2 models specifically designed for text summarization, Pegasus-Large and Brio, which have been fine-tuned on text summarization datasets, detailed in Table 2.

All models are sourced from the Hugging Face⁴ and tested based on the lm-evaluation-harness tool (Gao et al., 2023a). For general-purpose SLMs, we use the prompt from the huggingface hallucination leaderboard⁵ (Prompt 1 in Figure 3). To ensure reproducibility, we employ the greedy decoding strategy, generating until a newline character or <EOS> token is reached. Outputs are post-processed to remove prompt words and incomplete sentences. Given the 2048-token limit of some models, we only evaluate the zero-shot approach.

4.3 Evaluation Metric

We evaluate the summary quality in relevance, coherence, factual consistency, and text compression using BertScore (Zhang et al., 2020b), HHEM-2.1-Open (Vectara, 2024), and summary length.

BertScore is a robust semantic similarity metric designed to assess the quality of the generated text and achieves the best correlation in Table 1. It compares the embeddings of candidate and reference sentences through contextualized word representations from BERT (Devlin et al., 2019). It can measure how well the summary captures key information (relevance) and maintains logical flow (coherence). We use the F1 score of BertScore⁶, ranging from 0 to 100, where higher values indicate better quality.

HHEM-2.1-Open is a hallucination detection model which outperforms GPT-3.5-Turbo and even GPT-4. It can evaluate whether news summaries are factually consistent with the original article. It is based on a fine-tuned T5 model to flag hallucinations with a score. When the score falls below 0.5, it indicates a summary is inconsistent with the source. We report the percentage of summaries deemed factually consistent; higher values indicate better performance.

Since BertScore is insensitive to summary length, we also track the average summary length

⁴<https://huggingface.co>

⁵<https://huggingface.co/spaces/hallucinations-leaderboard/leaderboard>

⁶<https://huggingface.co/microsoft/deberta-xlarge-mnli>

Metric	Reference type	Relevance evaluation				Coherence evaluation			
		Bench -CNN/DM	Bench -XSum	SummEval	Average	Bench -CNN/DM	Bench -XSum	SummEval	Average
RougeL	Original	0.6732	0.2647	0.3714	0.4364	0.4248	0.6765	0.1238	0.4084
	Qwen1.5-72B	0.7778	0.7059	0.4476	0.6438	0.5556	0.5000	0.2381	0.4312
	llama2-70B	0.7778	0.6324	0.4095	0.6066	0.5294	0.3676	0.2000	0.3657
BertScore	Original	0.6209	0.2206	0.4476	0.4297	0.4771	0.6912	0.4286	0.5323
	Qwen1.5-72B	0.7516	0.6324	0.8095	0.7312	0.5817	0.6029	0.6000	0.5949
	llama2-70B	0.7516	0.6618	0.8286	0.7413	0.5294	0.6029	0.6190	0.5838
BLEURT	Original	0.5425	0.2647	0.181	0.3294	0.3203	0.6176	0.2381	0.3929
	Qwen1.5-72B	0.5817	0.7206	0.4476	0.5833	0.4902	0.4853	0.5429	0.5061
	llama2-70B	0.5556	0.7059	0.5429	0.6015	0.5425	0.4706	0.6762	0.5631

Table 1: System-level Kendall’s tau correlation coefficients between reference-based and human evaluations under different reference summary settings. "Original" is the original heuristic reference from the dataset; others are generated by LLMs. References generated by LLMs improve the effectiveness of reference-based evaluations.

Model	Parameters	Instruction Tuning	Reference
Brio *	406M	✓	(Liu et al., 2022)
LiteLlama	460M	×	(Han, 2024)
Qwen2-0.5B	500M	×	(Yang et al., 2024)
Qwen2-0.5B-Ins	500M	✓	(Yang et al., 2024)
Bloom-560M	560M	×	(Le Scao et al., 2022)
Pegasus-Large*	568M	✓	(Zhang et al., 2020a)
TinyLlama	1.1B	✓	(Zhang et al., 2024a)
Llama3.2-1B	1.2B	×	(Meta AI, 2024)
Llama3.2-1B-Ins	1.2B	✓	(Meta AI, 2024)
GPT-Neo-1.3B	1.3B	×	(Black et al., 2021)
Qwen2-1.5B	1.5B	×	(Yang et al., 2024)
Qwen2-1.5B-Ins	1.5B	✓	(Yang et al., 2024)
InternLM2-1.8B	1.8B	×	(Cai et al., 2024)
InternLM2-1.8B-Chat	1.8B	✓	(Cai et al., 2024)
MiniCPM	2.4B	✓	(Hu et al., 2024)
Gemma-1.1	2.5B	✓	(Team et al., 2024)
GPT-Neo-2.7B	2.7B	×	(Black et al., 2021)
Phi-2	2.7B	×	(Blog, 2023)
Llama3.2-3B	3.2B	×	(Meta AI, 2024)
Llama3.2-3B-Ins	3.2B	✓	(Meta AI, 2024)
Phi-3-Mini	3.8B	✓	(Abdin et al., 2024)

Table 2: Models for news summarization benchmark. Models marked with * have been fine-tuned on text summarization datasets.

to assess text compression. With similar BertScore results, shorter summaries indicate better text compression ability.

4.4 Reference Summary Generation

To mitigate the impact of occasional low-quality reference summaries on evaluation results, as well as the tendency for models within the same series to score higher due to possible similar output content, we utilize two LLMs, Qwen1.5-72B-Chat and Llama2-70B-Chat, to generate two sets of reference summaries. We then average the scores during the SLM evaluation. We use Prompt 2 from Figure 3 as the prompt template and apply a greedy strategy to generate summaries based on the vLLM inference framework (Kwon et al., 2023).

News URL: <https://www.bbc.com/sport/football/68395537>

Reference summary by Llama2-72B-Chat:	
Manchester City came from behind to beat Manchester United 3-1 in a thrilling Premier League match. Phil Foden scored a brace, including a stunning equalizer, and Erling Haaland added a third goal for City in stoppage time. Marcus Rashford had given United an early lead with a sensational strike, but City dominated the game and deservedly took all three points. The win moves City to within a point of leaders Liverpool, while United remain in sixth place.	
LiteLlama	BertScore: 40.81
<u>Manchester City's goal was a result of a poor decision by the referee. The referee was not aware of the goal and the referee was not aware of the decision. The referee was not aware of the ...</u>	
Qwen2-0.5B	BertScore: 66.58
Manchester City's Phil Foden scored twice in the last minute to help them beat Manchester United at Etihad Stadium, with Manchester United's Jadon Sancho missing a chance to equalize.	
Llama3.2-3B-Ins	BertScore: 70.10
Manchester City came from behind to beat Manchester United 3-1 at Etihad Stadium, with Phil Foden scoring a stunning equaliser and s in stoppage time. Pep Guardiola hailed Foden as the "best player in the Premier League right now" after his impressive display, which saw him score a second goal and create chances for his teammates.	

Figure 4: Example summaries from SLMs and LLMs. The bold part is the same as the reference, the underline indicates irrelevant content, and the red indicates incorrect content. Summaries with BertScore above 70, such as those from Llama3.2-3B-Ins, demonstrate similar quality to LLMs.

5 Evaluation Results

5.1 Relevance and Coherence Evaluation

Table 3 shows the average BertScore results of SLMs. All models demonstrate consistent performance across datasets. Due to the significant score differences, we categorize the models into three approximate ranges, using two specialized summarization models, Pegasus-Large and Brio, as reference points.

The first range includes scores below 60. **Models scoring below 60 struggle to effectively sum-**

Score Range	Model	XSum		Newsroom		CNN/DM		BBC2024		Average
		Qwen1.5	Llama2	Qwen1.5	Llama2	Qwen1.5	Llama2	Qwen1.5	Llama2	
< 60	LiteLlama	38.68	39.24	38.02	38.70	40.58	41.38	38.05	38.86	39.19
	Bloom-560M	48.23	48.43	42.73	43.28	49.70	50.04	47.56	48.01	47.25
	GPT-Neo-1.3B	59.67	59.55	52.54	53.11	58.37	58.81	58.16	58.18	57.30
	GPT-Neo-2.7B	60.11	60.15	53.65	54.14	59.15	59.97	59.29	59.50	58.25
60 ~ 70	Llama3.2-1B	62.44	62.33	55.25	55.86	60.32	60.67	62.17	62.33	60.17
	Pegasus-Large	61.39	61.44	61.26	61.64	61.94	62.26	63.12	63.47	62.07
	Llama3.2-3B	66.14	65.84	59.93	60.68	64.19	64.51	66.62	66.57	64.31
	MiniCPM	65.59	65.77	62.48	63.13	66.82	67.51	65.54	65.96	65.35
	TinyLlama	68.07	69.16	64.70	65.63	68.66	69.23	66.93	68.07	67.56
	InternLM2-1.8B	68.64	69.20	64.27	65.00	68.81	69.49	67.88	68.74	67.75
	InternLM2-1.8B-Chat	67.77	68.82	64.99	65.57	68.78	69.38	68.32	69.62	67.91
	Qwen2-0.5B-Ins	69.30	70.01	65.85	66.65	69.31	70.06	67.54	68.64	68.42
	Qwen2-0.5B	69.57	70.71	65.94	66.92	69.50	70.48	68.23	69.67	68.88
	Brio	70.67	70.71	68.43	68.41	69.86	70.34	70.39	70.21	69.88
> 70	Phi-2	70.95	72.03	68.15	69.14	71.19	72.21	70.38	71.50	70.69
	Gemma-1.1	70.80	72.27	69.42	70.30	71.88	72.83	70.60	71.67	71.22
	Qwen2-1.5B	72.03	72.90	69.19	69.91	71.33	72.42	71.08	72.18	71.38
	Qwen2-1.5B-Ins	72.21	73.06	69.49	69.86	71.64	71.95	71.01	71.97	71.40
	Llama3.2-1B-Ins	72.24	73.27	70.54	70.98	72.43	73.61	71.68	72.55	72.16
	Phi-3-Mini	74.67	75.14	72.42	72.32	74.86	74.88	74.08	73.72	74.01
	Llama3.2-3B-Ins	74.40	75.33	72.81	73.41	74.54	75.29	74.51	74.95	74.41

Table 3: BertScore of SLMs on four text summarization datasets. Qwen1.5-72B-Chat and Llama2-70B-Chat are used to generate reference summaries.

marize articles. LiteLlama, Bloom, and GPT-Neo series fall into this range. In terms of relevance, these models often miss key points; in terms of coherence, these models sometimes produce repetitive outputs, leading to overly long summaries. Figure 4 shows an example from LiteLlama, it scores only 40.81, deviating significantly from the news.

The second range covers scores between 60 and 70. **Models in this range produce useful summaries but occasionally lack key points.** Models like TinyLlama and Qwen-0.5B series fall into this range. In relevance, these models generally relate to the original content but may omit crucial details. In coherence, the sentences can be well-organized. Fig 4 shows examples from Qwen2-0.5B. Although its summary is relevant and coherent to the news article, it omits the final score.

The third range includes scores above 70. **Models in this range produce summaries comparable to LLMs, with occasional inconsistencies.** They effectively capture key points and maintain coherent structure. Phi3-Mini and Llama3.2-3B-Ins stand out in this group. As shown in Figure 4, Llama3.2-3B-Ins generates a concise, accurate summary, correctly capturing the match outcome even though the source news lacks a direct score description.

5.2 Factual Consistency Evaluation

During summary generation, SLMs may produce false information due to hallucinations. Therefore,




Figure 5: Average summary length comparison. SLMs with high BertScore generate 50-70 word summaries.

we report the factual consistency rate in Table 4. Models with high BertScore generally have better factual consistency as shown in Figure 4. SLMs that perform well in fact consistency include Phi3-Mini (96.7%) and Qwen2-1.5B-Ins (96.5%), maintaining over 95% consistency across all datasets. The traditional model, Pegasus-large, outperforms SLMs in fact consistency, which achieves an exceptional 99.9% average. This is because it often copies sentences from the original text, ensuring alignment but sometimes missing key information.

5.3 Summary Length Evaluation

Figure 5 shows the average summary lengths on 4 datasets produced by SLMs, which vary significantly across models. For scores below 60, summaries typically exceed 100 words due to redundant content and poor summarization. In the 60–70 range, length varies widely; e.g., the

Model	XSUM	Newsroom	CNN/DM	BBC2024	Average
LiteLama	60.6	63.4	68.8	63.6	64.1
Bloom-560M	83.6	77.4	87.0	83.4	82.9
GPT-Neo-1.3B	91.2	76.8	92.4	92.0	88.1
GPT-Neo-2.7B	83.2	67.0	87.6	84.0	80.5
Llama3.2-1B	97.6	85.8	97.6	95.8	94.2
Pegasus-Large	99.6	100.0	100.0	100.0	99.9
Llama3.2-3B	98.8	85.6	98.8	97.4	95.2
MiniCPM	85.8	72.8	93.6	89.6	85.5
TinyLlama	95.4	92.4	98.4	98.0	96.1
InternLM2-1.8B	91.8	87.4	96.0	94.8	92.5
InternLM2-1.8B-Chat	88.6	84.6	93.8	90.8	89.5
Qwen2-0.5B-Ins	87.2	86.0	93.2	89.6	89.0
Qwen2-0.5B	93.0	92.8	96.8	95.2	94.5
Brio	94.8	89.0	97.8	95.4	94.3
Phi2	95.2	93.4	97.8	96.8	95.8
Gemma-1.1	96.6	93.6	97.8	94.6	95.7
Qwen2-1.5B	95.6	96.2	97.4	96.6	96.5
Qwen2-1.5B-Ins	91.8	91.2	94.4	92.4	92.5
Llama3.2-1B-Ins	92.2	88.0	95.4	94.0	92.4
Phi3-Mini	96.0	94.8	98.2	97.6	96.7
Llama3.2-3B-Ins	93.2	93.2	96.6	95.6	94.7

Table 4: Factual consistency rate (%) on news datasets.

Llama3.2 series includes too many unnecessary details, while TinyLlama omits key information, leading to shorter summaries. Models with scores above 70 generate more consistent summaries, averaging around 50 to 70 words in length. For the top-performing models, Phi3-Mini and Llama3.2 series, Phi3-Mini generates more concise summaries while maintaining similar relevance.

5.4 Comparison with LLMs

To further evaluate the SLM performance in news summarization, we compare the well-performing SLMs with LLMs, including ChatGLM3 (GLM et al., 2024), Mistral-7B-Ins (Jiang et al., 2023), Llama3-70B-Ins (AI@Meta, 2024), and Qwen2 series, all in instruction-tuned versions. As the models show similar performance across various reference summaries and datasets, we use Llama2-70B-Chat to generate references and present the average results on BBC2024 in Table 5.

We highlight the top two metrics within each category. As can be seen, the summarization quality produced by the SLMs is on par with that of the LLMs. The difference in BertScore and factual consistency is minimal, with SLMs occasionally performing better. Notably, when BertScore scores are similar, smaller models tend to generate shorter summaries that facilitate quicker reading and comprehension of news outlines.

5.5 Human Evaluation

To further validate the effectiveness of our metrics, we performed a small-scale human evaluation, referencing existing studies (Fabbri et al., 2021). Specifically, six highly educated annotators evaluated the relevance of news summaries generated

Model	BertScore	Factual consistency	Summary length
Qwen2-1.5B-Ins	71.97	92.4	41
Phi3-Mini	73.72	97.6	51
Llama3.2-3B-Ins	74.95	95.6	70
ChatGLM3	73.41	96.0	59
Mistral-7B-Ins	74.32	93.0	60
Qwen2-7B-Ins	74.71	95.6	59
Llama3-70B-Ins	76.06	95.8	71
Qwen2-72B-Ins	73.78	98.4	86

Table 5: Average metric comparison of SLMs and LLMs on BBC2024 dataset.

Model	Prompt 1	Prompt 2	Prompt 3
Qwen2-0.5B-Ins	68.64	60.31	62.46
Qwen2-1.5B-Ins	71.97	69.74	72.30
Llama3.2-1B-Ins	72.55	72.16	67.53
Phi3-Mini	73.72	73.27	73.83
Llama3.2-3B-Ins	74.95	74.94	75.12

Table 6: BertScore comparison with different prompts.

by five SLMs based on the original news articles (scored on a 1-5 scale, with higher scores indicating greater relevance). We randomly selected 20 news instances from the BBC2024 dataset as our evaluation set. As shown in Table 7, the average scores of each model demonstrate a strong alignment between BertScore and human evaluations, achieving a Kendall correlation coefficient of 1.

5.6 Overall Evaluation and Model Selection

SLMs exhibit considerable variability in text summarization performance, with larger and newer models generally showing stronger capabilities. Among them, Phi3-Mini and Llama3.2-3B-Ins perform the best, matching the 70B LLM in relevance, coherence, and factual consistency, while producing shorter summaries. All of these comparisons suggest that SLMs can effectively replace LLMs on edge devices for news summarization.

Although we evaluate 19 SLMs, there is considerable variation in their parameter sizes. To optimize deployment on edge devices, we provide model selection recommendations based on size. For models under 1B parameters, Brio and Qwen2-0.5B are the top choices overall. General-purpose language models in this range offer no clear advantage over those specialized in text summarization. For models between 1B and 2B parameters, Llama3.2-1B-Ins demonstrates a clear edge. For models above 2B parameters, Llama3.2-3B-Ins and Phi3-Mini outperform others across all criteria.

Model	Ann. 1	Ann. 2	Ann. 3	Ann. 4	Ann. 5	Ann. 6	Average	BertScore
Bloom-560M	1.95	1.25	1.43	1.60	1.00	1.00	1.37	46.30
Llama3.2-3B	3.15	2.10	2.25	2.50	1.35	1.15	2.08	66.36
Qwen2-0.5B	4.05	2.15	3.28	2.70	3.60	3.25	3.17	69.85
Phi3-Mini	4.45	4.60	4.00	3.45	4.60	4.25	4.20	73.14
Llama3.2-3B-Ins	4.65	4.45	4.35	3.65	4.00	4.10	4.23	75.15

Table 7: Humance evaluation in relevance on 5 SLMs.

News URL: <https://www.bbc.com/sport/football/68395537>

Llama3.2-3B-Ins + Prompt 1	BertScore: 70.10
Manchester City came from behind to beat Manchester United 3-1 at Etihad Stadium, with Phil Foden scoring a stunning equaliser and Erling Haaland sealing the win in stoppage time. Pep Guardiola hailed Foden as the "best player in the Premier League right now" after his impressive display, which saw him score a second goal and create chances for his teammates.	
Llama3.2-3B-Ins + Prompt 2	BertScore: 70.44
Manchester City came from behind to beat Manchester United 3-1 at Etihad Stadium, with Phil Foden scoring a stunning equalizer and Erling Haaland sealing the win in stoppage time. Pep Guardiola hailed Foden as the "best player in the Premier League right now" after his impressive performance, which included a sublime goal and a second that saw him burst away from a static Casemiro.	
Llama3.2-3B-Ins + Prompt 3	BertScore: 69.93
Manchester City came from behind to beat Manchester United 3-1 at Etihad Stadium, with Phil Foden scoring a stunning equaliser and Erling Haaland sealing the win in stoppage time. Pep Guardiola hailed Foden as the "best player in the Premier League right now" after his impressive display, which saw him score twice and create chances for his teammates.	

Figure 6: Examples summaries generated by Llama3.2-3B-Ins with different prompts. The bold parts show the same content, suggesting that the prompt design has a limited improvement in summary quality.

6 Influencing Factor Analysis

6.1 Prompt Design

Prompt engineering has demonstrated significant power in many tasks, helping to improve the output quality of LLMs (Zhao et al., 2021; Sahoo et al., 2024). Therefore, we use different prompt templates shown in Figure 3 to analyze the impact of prompt engineering on BertScore scores, factual consistency, and summary length. The instructions become more detailed from Prompt 1 to Prompt 3.

Table 6 compares the BertScore. The small score differences among prompts suggest that detailed descriptions do not significantly improve summary relevance and coherence and may even negatively affect some models. This suggests that **simple instructions are more suitable for SLMs, and overly complex prompts may degrade the performance**. For instance, the BertScore of Qwen2-0.5B-Ins drops significantly. The decline is likely due to prompt noise, where the model struggles to balance multiple instructions, resulting in less

Model	Prompt 1	Prompt 2	Prompt 3
Qwen2-0.5B-Ins	89.6	89.8	78.2
Qwen2-1.5B-Ins	92.4	90.8	89.0
Llama3.2-1B-Ins	94.0	94.8	91.6
Phi3-Mini	97.6	96.4	97.2
Llama3.2-3B-Ins	95.6	96.2	94.6

Table 8: Factual consistency rate (%) comparison with different prompts.

Model	Prompt 1	Prompt 2	Prompt 3
Qwen2-0.5B-Ins	65	150	151
Qwen2-1.5B-Ins	41	99	53
Llama3.2-1B-Ins	63	60	58
Phi3-Mini	51	57	65
Llama3.2-3B-Ins	70	67	61

Table 9: Summary length with different prompts. All prompts require models to summarize in two sentences.

coherent summaries. Additionally, Figure 6 shows an example of Llama3.2-3B-Ins with different prompts. All summaries capture the key information, with only slight differences in details.

Table 8 provides a comparison of factual consistency rates for various models using three different prompts. Many models, such as Qwen2-0.5B-Ins, exhibit a significant drop in factual consistency when using prompt 3. This further demonstrates that models with very small parameters are not well-suited for complex prompt design. The summary length variations across different prompts in Table 9 also indicate the limitations of prompt design. Some models fail to generate summaries of appropriate length given complex prompts.

Overall, the improvement of prompt design on the quality of SLM summaries is limited. For SLMs that already perform well with simple prompts, they can capture the key points without the need for complex prompts. In contrast, SLMs that perform poorly may struggle with the logical relationships in complex prompts, leading to lengthy summaries. For practical applications, we recommend using simple and clear prompts when

deploying SLMs for news summarization.

6.2 Instruction Tuning

Instruction tuning plays a crucial role in training LLMs by enhancing their ability to follow specific instructions. And some existing studies claim that instruction-tuned language models have stronger summarization capabilities (Goyal et al., 2022; Zhang et al., 2024b). However, in our extensive evaluations, with the exception of the Llama3.2 series, models perform very similarly before and after instruction tuning. In fact, models without instruction tuning often exhibit higher factual consistency. In the sample analysis, we find that the Llama3.2 models without instruction tuning include too many details, while instruction tuning helps it produce more general summaries. However, this trend is not observed in other models like Qwen2 and InternLM2 series, which is different from the conclusions of previous work (Zhang et al., 2024b). We leave the deep study of instruction tuning on summarization ability as a future research direction.

7 Conclusion

This paper presents the comprehensive evaluation of SLMs for news summarization, comparing 19 models across diverse datasets. Our results demonstrate that top-performing models like Phi3-Mini and Llama3.2-3B-Ins can match the performance of larger 70B LLMs while producing shorter, more concise summaries. These findings highlight the potential of SLMs for real-world applications, particularly in resource-constrained environments. Further exploration shows that simple prompts are more effective for SLM summarization, while instruction tuning provides inconsistent benefits, necessitating further research.

Limitations

Although using LLM-generated summaries improved the reliability of reference-based summarization evaluation methods, it may introduce bias. For instance, if the LLM fails to accurately summarize the news, an SLM that produces similar content might receive an undeservedly high score. To mitigate this potential bias, we use summaries generated by multiple LLMs as references and calculate the average scores. Furthermore, although BertScore demonstrates a high level of consistency with human evaluations in coherence, it is not specifically designed to evaluate coherence. In

future work, we plan to explore the use of more accurate metrics for coherence evaluation. Due to the input length limits of most SLMs, we only considered news articles with a maximum of 1,500 tokens. In the future, we will explore how to effectively use SLMs to summarize much longer news articles. Additionally, we do not consider quantized models in this study. In future work, we will also incorporate an analysis of the performance of quantized models into our evaluation.

Acknowledgments

This work is funded by the National Key R&D Program of China (No. 2019YFA0709400), the NSFC Project (No. 62306256), the Guangzhou Science and Technology Development Projects (No. 2023A03J0143 and No. 2024A04J4458), and the Guangzhou Municipality Big Data Intelligence Key Lab (NO. 2023A03J0012). This work is also supported by the National Research Foundation, Singapore under its Industry Alignment Fund – Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. Borui’s work was done when he was a visiting student at the National University of Singapore.

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- AI@Meta. 2024. [Llama 3 model card](#).
- Stella Biderman, Kieran Bicheno, and Leo Gao. 2022. Datasheet for the pile. *arXiv preprint arXiv:2201.07311*.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#). If you use this software, please cite it using these metadata.
- Microsoft Research Blog. 2023. [Phi-2: The surprising power of small language models](#). Accessed: 2024-06-27.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda

- Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaying Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yinling Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingdong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. 2024. [Internlm2 technical report](#). Preprint, arXiv:2403.17297.
- Roshni Chakraborty, Maitry Bhavsar, Sourav Kumar Dandapat, and Joydeep Chandra. 2019. [Tweet summarization of news articles: An objective ordering-based perspective](#). *IEEE Transactions on Computational Social Systems*, 6(4):761–777.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Together Computer. 2023. [Redpajama: an open dataset for training large language models](#).
- 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.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. [Bandit-Sum: Extractive summarization as a contextual bandit](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3739–3748, Brussels, Belgium. 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.
- Xue-Yong Fu, Md Tahmid Rahman Laskar, Elena Khasanova, Cheng Chen, and Shashi Tn. 2024. [Tiny titans: Can smaller large language models punch above their weight in the real world for meeting summarization?](#) In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 387–394, Mexico City, Mexico. Association for Computational Linguistics.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The Pile: An 800GB dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023a. [A framework for few-shot language model evaluation](#).
- Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. 2023b. Human-like summarization evaluation with chatgpt. *arXiv preprint arXiv:2304.02554*.
- Pierre-Etienne Genest and Guy Lapalme. 2012. [Fully abstractive approach to guided summarization](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 354–358, Jeju Island, Korea. Association for Computational Linguistics.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jidai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li,

- Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *Preprint*, arXiv:2406.12793.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Xiaotian Han. 2024. [Litellama-460m-1t](#). Accessed: 2024-06-27.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, page 1693–1701, Cambridge, MA, USA. MIT Press.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, et al. 2024. [Minicpm: Unveiling the potential of small language models with scalable training strategies](#). *arXiv preprint arXiv:2404.06395*.
- Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020. [What have we achieved on text summarization?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 446–469, Online. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Teven Le Scao, Thomas Wang, Daniel Hesslow, Stas Bekman, M Saiful Bari, Stella Biderman, Hady Elsahar, Niklas Muennighoff, Jason Phang, Ofir Press, Colin Raffel, Victor Sanh, Sheng Shen, Lintang Sutawika, Jaesung Tae, Zheng Xin Yong, Julien Launay, and Iz Beltagy. 2022. [What language model to train if you have one million GPU hours?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 765–782, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Haozhou Li, Qinke Peng, Xu Mou, Ying Wang, Zeyuan Zeng, and Muhammad Fiaz Bashir. 2023. [Abstractive financial news summarization via transformer-bilstm encoder and graph attention-based decoder](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:3190–3205.
- 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.
- Yixin Liu, Alexander Fabbri, Jiawen Chen, Yilun Zhao, Simeng Han, Shafiq Joty, Pengfei Liu, Dragomir Radev, Chien-Sheng Wu, and Arman Cohan. 2024. [Benchmarking generation and evaluation capabilities of large language models for instruction controllable summarization](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4481–4501, Mexico City, Mexico. Association for Computational Linguistics.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. [BRIO: Bringing order to abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.
- Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, Nicholas D. Lane, and Mengwei Xu. 2024. [Small language models: Survey, measurements, and insights](#). *Preprint*, arXiv:2409.15790.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*, 2(2):159–165.
- Goran Matošević. 2019. Text summarization techniques for meta description generation in process of search engine optimization. In *Artificial Intelligence and Algorithms in Intelligent Systems*, pages 165–173, Cham. Springer International Publishing.
- Joshua Maynez, Priyanka Agrawal, and Sebastian Gehrmann. 2023. [Benchmarking large language model capabilities for conditional generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9194–9213, Toronto, Canada. Association for Computational Linguistics.
- Meta AI. 2024. [Llama 3.2: Revolutionizing edge ai and vision with open, customizable models](#). Accessed: 2024-09-28.
- Microsoft Azure. 2024. [Introducing phi-3: Redefining what's possible with slms](#).

- 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.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023a. Summarization is (almost) dead. *arXiv preprint arXiv:2309.09558*.
- Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023b. Summarization is (almost) dead. *arXiv preprint arXiv:2309.09558*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. [A systematic survey of prompt engineering in large language models: Techniques and applications](#). *Preprint*, arXiv:2402.07927.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Chenhui Shen, Liying Cheng, Xuan-Phi Nguyen, Yang You, and Lidong Bing. 2023. [Large language models are not yet human-level evaluators for abstractive summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4215–4233, Singapore. Association for Computational Linguistics.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Priyam Tejaswin, Dhruv Naik, and Pengfei Liu. 2021. [How well do you know your summarization datasets?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3436–3449, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esionu, 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](#). *Preprint*, arXiv:2307.09288.
- 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. Curran Associates, Inc.
- Vectara. 2024. [Hallucination evaluation model](#). https://huggingface.co/vectara/hallucination_evaluation_model. Accessed: 2024-10-02.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. [Is chatgpt a good nlg evaluator? a preliminary study](#). *arXiv preprint arXiv:2303.04048*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. Pegasus: pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20*. JMLR.org.

Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024a. *Tinyllama: An open-source small language model*. Preprint, arXiv:2401.02385.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. *Opt: Open pre-trained transformer language models*. Preprint, arXiv:2205.01068.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.

Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. 2024b. Benchmarking large language models for news summarization. *Transactions of the Association for Computational Linguistics*, 12:39–57.

Xingxing Zhang, Mirella Lapata, Furu Wei, and Ming Zhou. 2018. *Neural latent extractive document summarization*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 779–784, Brussels, Belgium. Association for Computational Linguistics.

Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*, pages 12697–12706. PMLR.

A Challenges in Existing Summary Evaluation Method

In this section, we will supplement the detailed analysis of existing evaluation methods and explain why we use the reference-based evaluation method.

A.1 Human evaluation

Human evaluation remains the most intuitive and effective method for text summarization evaluation (Fabbri et al., 2021; Pu et al., 2023b; Liu et al., 2024; Zhang et al., 2024b). Annotators generally score generated summaries on aspects such as relevance and coherence (scale 1-5) based on the source text, and the final score is the average across all samples. Despite its reliability, human evaluation faces three main challenges for large-scale use: (1) Time-consuming—even evaluating 100 samples for each model can take weeks (Zhang et al.,




Figure 7: Pairwise system-level Kendall’s tau correlation on relevance for different annotators in SummEval benchmark. "CS" is the abbreviation for crowd-sourced. Darker colors indicate better consistency.

2024b); (2) Low reproducibility—scores may vary due to different annotators; (3) High cost—only expert annotators are suitable, but hiring them is costly (Fabbri et al., 2021). As shown in Figure 7, experts demonstrate significantly higher agreement than crowd-sourced annotators.

A.2 Reference-based Evaluation

Reference-based methods are among the earliest for evaluating text summarization (Papineni et al., 2002). They are more cost-effective and faster than human evaluation, using statistical metrics or small models to compute the similarity between generated and reference summaries. Higher similarity indicates better quality. There are two main categories: one measures textual overlap (e.g., ROUGE (Lin, 2004), BLEU (Papineni et al., 2002)); the other uses language models like BERT (Devlin et al., 2019) to compute semantic similarity (e.g., BertScore (Zhang et al., 2020b), BLEURT (Sellam et al., 2020)). However, the reference quality affects the evaluation reliability heavily (Tejaswin et al., 2021; Zhang et al., 2024b). Heuristic methods used to create references in many datasets often lead to poor-quality summaries. For instance, XSum (Narayan et al., 2018) uses only the first sentence of a news article as the reference, resulting in 54% of summaries missing key information (Tejaswin et al., 2021). Moreover, since different datasets use varied reference generation methods, evaluation reliability fluctuates. Table 1 illustrates the impact of reference quality on reference-based methods. The correlation with human evaluation varies significantly

across datasets when using the original reference for each dataset. It highlights the need for a unified, high-quality reference generation method.

A.3 LLM evaluation

Recent studies (Goyal et al., 2022; Chiang and Lee, 2023; Gao et al., 2023b; Liu et al., 2024; Wang et al., 2023) have explored using LLMs for human-like summarization evaluation to achieve better alignment with human judgment. This approach involves guiding LLMs with prompts and providing both the summary and the original article for scoring. However, while LLM-based evaluation is reproducible and promising, it shares cost and efficiency challenges with human evaluation. Using commercial LLM APIs for large datasets can be expensive, and deploying open-source LLMs (e.g., 70B models) locally requires costly GPUs like the Nvidia A100 for reasonable speed. Additionally, studies (Shen et al., 2023) indicate that LLMs may exhibit biases when evaluating different models, warranting further research to ensure fairness.

Overall, existing evaluation methods face issues related to financial cost, reproducibility, and effectiveness. Among them, reference-based methods are the most cost-effective and efficient. Therefore, we consider the LLM-augmented reference-based evaluation method as our metrics.

B Model Details

In this section, we describe all benchmarked SLMs in our experiments.

LiteLlama (Han, 2024): LiteLlama is an open-source reproduction of Llama2 (Touvron et al., 2023). It has 460M parameters and is trained on 1T tokens from the RedPajama dataset (Computer, 2023).

Bloom-560M (Le Scao et al., 2022): Bloom-560M is part of the BLOOM (BigScience Large Open-science Open-access Multilingual) family of models developed by the BigScience project. It has 560M parameters and is trained on 1.5T pre-processed text.

TinyLlama (Zhang et al., 2024a): TinyLlama has the same architecture and tokenizer as Llama 2. It has 1.1B parameters and is trained on 3T tokens. We use the chat finetuned version.

GPT-Neo series (Black et al., 2021): GPT-Neo Series are open-source language models developed by EleutherAI to reproduce GPT-3 architecture. It has two versions with 1.3B and 2.7B parameters.

They are trained on 380B and 420B tokens, respectively, on the Pile dataset (Gao et al., 2020; Biderman et al., 2022).

Qwen2 series (Yang et al., 2024): Qwen2 comprises a series of language models pre-trained on multilingual datasets. And some models use instruction tuning to align with human preferences. We select models with no more than 7B parameters for benchmarking, including 0.5B, 0.5B-Ins, 1.5B, 1.5B-Ins, 7B, and 7B-Ins. Models with the "Ins" suffix indicate the instruction tuning version.

InterLM2 series (Cai et al., 2024): The InternLM2 series includes models with various parameter sizes. They all support long contexts of up to 200,000 characters. We select the 1.8B version and the 1.8B-Chat version for evaluation.

Gemma-1.1 (Team et al., 2024): Gemma-1.1 is developed by Google and is trained using a novel RLHF method. We select the 2B instruction tuning version for benchmarking.

MiniCPM (Hu et al., 2024): MiniCPM is an end-size language model with 2.7B parameters. It employs various post-training methods to align with human preferences. We select the supervised fine-tuning (SFT) version for benchmarking.

Llama3.2 series (Meta AI, 2024): In the Llama 3.2 series, Meta introduced two SLMs, 1B and 3B, designed for edge devices. While Meta highlighted their capability to summarize social media messages, their performance in news summarization has yet to be evaluated.

Phi series (Blog, 2023; Abdin et al., 2024): Phi series are some SLMs developed by Microsoft. They are trained on high-quality datasets. We select Phi-2(2.7B parameters) and Phi-3-Mini(3.8B parameters, 4K context length) for evaluation.

Brio (Liu et al., 2022): Brio is one of the state-of-the-art specialized language models designed for text summarization, which leverages reinforcement learning to optimize the selection of sentences, ensuring high-quality, informative, and coherent summaries. It only has 406M parameters, and we select the version trained on the CNN/DM dataset.

Pegasus-Large (Zhang et al., 2020a): Pegasus-Large also is a specialized language model designed for abstractive text summarization, utilizing gap-sentence generation pre-training to achieve exceptional performance on summarization tasks. It has 568M parameters.