

Salient Information Prompting to Steer Content in Prompt-based Abstractive Summarization

Lei Xu¹, Mohammed Asad Karim^{2†}, Saket Dingliwal¹, Aparna Elangovan¹

¹Amazon AWS AI Labs

²Carnegie Mellon University

{leixx, skdin, aeg}@amazon.com mkarim2@cs.cmu.edu

Abstract

Large language models (LLMs) can generate fluent summaries across domains using prompting techniques, reducing the need to train models for summarization applications. However, crafting effective prompts that guide LLMs to generate summaries with the appropriate level of detail and writing style remains a challenge. In this paper, we explore the use of salient information extracted from the source document to enhance summarization prompts. We show that adding keyphrases in prompts can improve ROUGE F1 and recall, making the generated summaries more similar to the reference and more complete. The number of keyphrases can control the precision-recall trade-off. Furthermore, our analysis reveals that incorporating phrase-level salient information is superior to word- or sentence-level. However, the impact on hallucination is not universally positive across LLMs. To conduct this analysis, we introduce Keyphrase Signal Extractor (SigExt), a lightweight model that can be finetuned to extract salient keyphrases. By using SigExt, we achieve consistent ROUGE improvements across datasets and open-weight and proprietary LLMs without any LLM customization. Our findings provide insights into leveraging salient information in building prompt-based summarization systems.

1 Introduction

Abstractive summarization aims to generate concise summaries that capture the most salient information from lengthy source documents. Prior work has shown that emphasizing keywords from source documents can enhance summarization performance on supervised finetuned (SFT) models (Gu et al., 2016). However, existing approaches (Nallapati et al., 2016; See et al., 2017; Liu et al., 2021) require extensive modifications to the architecture

and loss functions, hindering widespread adoption, especially for large language models (LLMs) with billions of parameters. Recent work (Li et al., 2023a) trains a separate network using reinforcement learning (RL) to generate keyphrases for LLM prompts, but training RL model is non-trivial due to convergence and stability issues (Wang et al., 2024). Emphasizing salient information in the prompt can help zero-shot LLMs generate more complete summaries, and steer LLMs to generate summaries that align with the desired use case. However, there is also a lack of analysis on how emphasizing salient information in prompts would affect the LLM behavior.

We first address the challenge of applying salient information to LLMs. We obtain keyphrases using a stand-alone keyphrase signal extractor called SigExt, and prompt the LLMs to consider these keyphrases when generating summaries. Unlike prior work relying on complex keyphrase generators optimized for specific LLMs, SigExt is LLM-agnostic, allowing leveraging salient information with large API-based models that cannot be finetuned. We demonstrate consistent improvement in ROUGE scores on 4 representative summarization datasets and 3 recent LLMs – Claude, Mistral (Jiang et al., 2023), and Falcon (Almazrouei et al., 2023) – highlighting the wide adaptability of our approach. Secondly, we conduct comprehensive experiments using SigExt to gain insights into how keyphrases in prompts affect different aspects of summary quality. We show that adding keyphrases improves ROUGE F1 and recall, making the generated summaries more similar to the reference and more complete. Adjusting the number of keyphrases influences the trade-off between precision and recall. Including additional keyphrases in the prompt tends to produce more detailed summaries, enhancing recall. Our findings indicate that using phrase-level salient information is more effective than word- or sentence-level approaches.

[†]Work done during an internship at AWS AI Labs.

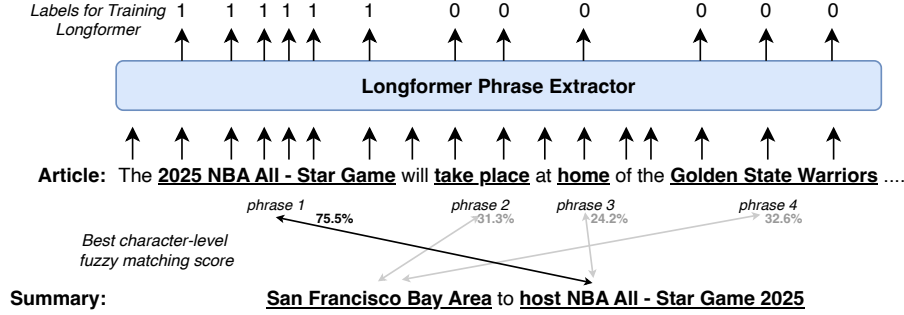


Figure 1: SigExt – a finetuned Longformer to extract keyphrases from an article. We construct labels by thresholding the character-level fuzzy matching score between phrases in the article and the summary.

However, for certain large language models like Mistral, adding keyphrases may lead to more hallucinations.

Our analysis offers guidance for applying similar strategies in real-world summarization applications. While incorporating salient information is an effective method for enhancing and controlling the completeness of summaries, and using phrase-level granularity proves more effective, the risk of introducing hallucinations must be carefully considered. This risk depends on the specific LLM being used, the method for gathering salient information, and the criticality of the application.

Our key contributions are as follows:

- 1) We present SigExt, a simple yet effective keyphrase extraction model using a finetuned Longformer (Beltagy et al., 2020). Once trained, SigExt is LLM-agnostic, enabling performance boost for different LLMs by adding extracted keyphrases in prompts without requiring LLM finetuning.
- 2) We provide a comprehensive analysis on the impact of adding salient information in prompts for summarization, including insights on summary length, reference alignment, completeness, and hallucination.
- 3) We demonstrate that SigExt has cross-domain generalization capability through a general-purpose version (GP-SigExt) pretrained on 7 datasets.

2 Method

In this section, we introduce SigExt – a keyphrase extractor designed for boosting summarization quality of prompt-based LLMs. Figure 1 gives an overview. SigExt tokenizes the source document into phrases (phrase tokenization is detailed in Section 2.1), and simultaneously predict whether each phrase is important. To train the model, we create target labels by identifying phrases appear in

both the source document and the summary, then optimizing the cross entropy loss. Compared to previous a keyphrase generator that uses RL (Li et al., 2023a), SigExt allows easier control of keyphrase numbers, faster training and inference, and better consistency across domains. We directly incorporate keyphrases in prompt, making it generalizable across LLMs. To handle longer input lengths while maintaining efficiency, we build SigExt using Longformer, so that training and inference can be done on a single GPU.

2.1 Phrase tokenization

Let $\mathbf{x} = x_1, \dots, x_n$ be a source document of n tokens, and $\mathbf{y} = y_1, \dots, y_m$ be the target summary of m tokens. The document is segmented into non-overlapping phrases by removing stopwords and punctuation. After this, we get a sequence of T non-overlapping phrases, denoted as $\text{Phrase}(\mathbf{x}) = [p_i = x_{l_i} \dots x_{r_i}]_{i=1 \dots T}$. Similarly, we get T' phrases from the summary denoted as $\text{Phrase}(\mathbf{y}) = [q_j = y_{l'_j} \dots y_{r'_j}]_{j=1 \dots T'}$.

2.2 Labels and learning objective

We label each input phrase by compute the fuzzy matching score

$$\text{fuzz}(a, b) = \frac{|\text{longest_common_sequence}(a, b)|}{\max(|a|, |b|)},$$

against all phrases in the summary. If the maximum score exceeds certain threshold ϵ , it is considered a keyphrase, formally

$$\text{label}(p_i) = \begin{cases} 1 & \max_{j \in 1 \dots T'} \text{fuzz}(p_i, q_j) \geq \epsilon, \\ 0 & \text{otherwise.} \end{cases}$$

We train a classification model to predict the label. Specifically, we use a Longformer and add a classification head on top of each token. We compute the cross entropy loss on tokens that belong to phrases,

while ignoring predictions on punctuation and stop-word tokens. We apply class balancing weight λ when the label of the token is 0.

2.3 Application of SigExt on summarization

We first finetune SigExt on the summarization dataset to get a task-specific keyphrase extractor. During inference, we use SigExt to extract keyphrases, then wrap the source article with a summarization prompt, and include keyphrases in the prompt. Here is an example prompt:

```
Here is an news article: <text> \nHere are a few keyphrases from the article: <key_phrases> \nPlease write an summary for the article. \nSummary:
```

To select keyphrases, we first score each phrase by calculating the average logits of its tokens. We then select the top- K deduplicated phrases according to their logits scores, removing any duplicates that exceed a fuzzy matching threshold ϵ and keeping the longer phrase in those cases. We replace `<key_phrases>` with comma separated keyphrases. This prompt then serves as the input to the LLM which produces the final summary.

2.4 Cross domain generalization

In order to generalize the keyphrase extractor model to new domains without fine-tuning for the target domain, we train a general purpose keyphrase extractor using a combination of 7 datasets. The datasets are XSUM (Narayan et al., 2018), Multi-News (Fabbri et al., 2019), Gigaword (Nallapati et al., 2017), Big-Patent (Sharma et al., 2019), AESLC (Zhang and Tetreault, 2019), BillSum (Kornilova and Eidelman, 2019), and WikiHow (Koupae and Wang, 2018). We call this general-purpose keyphrase signal extractor model GP-SigExt.

3 Experiments

Datasets: We select 4 representative datasets – SAMSum (Gliwa et al., 2019), CNN/Daily-Mail (Nallapati et al., 2016), ArXiv (Cohan et al., 2018), and MeetingBank (Hu et al., 2023) – to evaluate our method. These datasets cover short and long text, as well as regular document and conversation summarization. Dataset details are shown in Table 11 in Appendix. We truncate input text to 4,000 tokens to fit the context window of the Longformer model. We follow the convention to evaluate on 500 randomly sampled examples (Zhang et al., 2020). We report results averaged on 3 runs.

LLMs and Prompts: We evaluate SigExt on Claude-Instant, Mistral-7B-Instruct, and Falcon-40B-Instruct LLMs. We do not use Falcon on ArXiv and MeetingBank datasets due to its limited context window. We manually optimized the prompts for each model and task to achieve competitive zero-shot performance. All prompts are listed in Appendix A.

SigExt & GP-SigExt Parameters: We use Longformer-large (433M) for the keyphrase extractor. We set the fuzzy matching threshold $\epsilon = 70\%$, and the class balancing weight $\lambda = 0.1$. For SigExt, we sample 1000 examples from training set, we train SigExt starting with original Longformer-large checkpoint. For GP-SigExt, we sample 10000 examples from each of the 7 dataset mentioned in Sec. 2.4. We train SigExt and GP-SigExt for 10 epochs, and use validation set to pick the best checkpoint based on recall@20 (Metric defined in Sec. 3.7).

During prompting, we try $K = 10, 15, 20$ keyphrases for the CNN, SAMSum, and MeetingBank datasets, and $K = 30, 35, 40$ keyphrases for the ArXiv dataset. We pick the best number of keyphrases based on ROUGE scores on the validation set. We also conduct an ablation study on the effect of different numbers of keyphrases.

Baseline: We compare our methods with naive zero-shot prompting. We adapt a **2-pass** extract-then-abstract method (Zhang et al., 2023) to the three LLMs and use it as a baseline. This method uses the LLM to extract sentences from the source document in the first pass, then uses the second pass to revise the extracted sentences into an abstractive summary. We also compare with Directional Stimulus Prompting (Li et al., 2023b) which utilize reinforcement learning to select good keywords.

Evaluation Metrics: We compute ROUGE-1/-L F1 scores (abbreviated as **R1-f**, **RL-f**) to evaluate summary quality. We also report ROUGE-1 Recall (R1-r) to assess the completeness. We use AlignScore (Zha et al., 2023) to evaluate the faithfulness of the summary.

3.1 Main Results

Table 1 shows the ROUGE scores on all 4 datasets. The F1 scores are improved by using GP-SigExt without any fine-tuning on new datasets. By fine-tuning only the phrase extractor, SigExt further improves the score, showing that using a supervisedly learned keyphrase extractor can make the LLM generate summaries more similar to the ref-

Method	SAMSum			CNN/DailyMail			ArXiv			MeetingBank			Avg.
	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	Δ R1-f
Claude-Ins.	40.0	30.3	52.8	38.1	23.9	41.9	44.4	23.1	53.2	32.2	21.8	43.4	
+2-stage	40.3	31.0	46.9	39.2	24.6	48.3	44.0	22.9	50.4	30.8	20.7	43.8	-0.1
+GP-SigExt	40.0	30.0	57.3	40.2	24.9	47.5	44.7	23.2	53.5	36.3	25.7	53.1	1.6
+SigExt	41.6	30.9	59.5	42.0	26.6	48.6	45.2	23.5	53.7	42.3	31.9	60.5	4.1
Mistral-7B	40.5	31.7	48.2	38.9	24.8	42.6	43.1	24.6	41.6	34.4	25.2	50.3	
+2-stage	38.7	30.6	45.4	38.0	24.4	48.6	39.5	22.0	41.9	32.0	23.5	52.0	-2.2
+GP-SigExt	41.9	32.2	50.7	39.5	25.2	45.3	42.8	23.8	44.7	34.1	24.7	54.8	0.4
+SigExt	44.1	33.9	54.5	40.9	26.0	47.9	43.6	24.2	45.2	37.0	27.2	58.7	2.2
Falcon-40B	37.1	28.7	46.3	25.7	16.4	33.8							
+2-stage	36.1	28.1	54.1	34.2	22.1	53.2							3.8
+GP-SigExt	38.5	29.4	54.1	31.9	20.4	42.3							3.8
+SigExt	39.9	30.4	56.1	33.5	21.3	43.2							5.3
0-shot SOTA	38.8	30.6	-	36.0	22.3	-	34.6	18.3		36.4	26.8	-	

Table 1: Performance of SigExt & GP-SigExt on summarization using Claude Instant, Mistral-7B-Instruct, and Falcon-40B-Instruct. SigExt is trained with 1000 examples, while GP-SigExt is not fine-tuned on the dataset. We compare our methods with zero-shot prompting and 2-stage extract-then-abstract baselines. We show ROUGE-1 F-Measure (R1-f), ROUGE-L F-Measure (RL-f), and ROUGE-1 recall (R1-r). The LLMs are not fine-tuned. We directly copy zero-shot SOTA for SAMSum and CNN from Laskar et al. (2023), ArXiv from Xiao et al. (2022), and MeetingBank from Hu et al. (2023).

erence. On average, compared to the already strong zero-shot Claude Instant baseline, R1-F improves by 1.6% with GP-SigExt and 4.1% with SigExt. Similar improvements are also observed on Mistral and Falcon models. Besides F1 scores, adding keyphrases extracted by both SigExt and GP-SigExt into the prompts can significantly increase the R1-r score, showing that adding salient information can improve the completeness of the summary. Our method achieves a smaller gain on the ArXiv dataset compared to other datasets. We hypothesize that this is because paper abstracts have a standard format, and the keyphrases they should contain are thus more predictable. As a result, the zero-shot LLM can already identify and include these keyphrases in the output. For other datasets, where the summary is more subjective, our method can help the LLM incorporate proper information in the summary to better align with the reference.

Although the length of the summary slightly increase with the introduction of keyphrases, we do not achieve these improvements by excessively increasing the length of the summary. On average, the length of Claude Instant summaries increases by 4.7 words after adding keyphrases, whereas it increases by 13.6 words for Mistral and 12.3 words for Falcon.

We also compare the performance of SigExt with recent Directional Stimulus Prompting baseline on

ChatGPT(gpt-3.5-turbo) in Table 2. We show that SigExt can also boost ChatGPT zero-shot performance, and outperform the baseline.

Method	#examples	R1-f	RL-f
Vanilla	0	38.5	25.5
Directional Stimulus	4000	40.2	26.8
SigExt	4000	42.2	27.0

Table 2: Comparing SigExt with baselines using ChatGPT and CNN dataset.

3.2 Human Qualitative Check

To verify the quality of the notes, we follow Liu et al. (2023) and conduct a human evaluation, in which they annotated Atomic Content Units (ACUs) for several public datasets. Each ACU represents a fact that should appear in the summary. We select 50 documents from the CNN and SAMSum datasets, respectively, and ask human annotators to verify whether the given ACU appears in the summary. We report both the **raw ACU coverage** and **length-normalized ACU coverage**, as proposed by Liu et al. (2023). Table 3 shows that SigExt consistently outperforms the vanilla LLM in terms of ACU coverage.

3.3 Number of Keyphrases

We try different numbers of keyphrases in the prompt for each dataset, and show the ROUGE-1 Precision/Recall/F1 curves in Figure 2. The F1

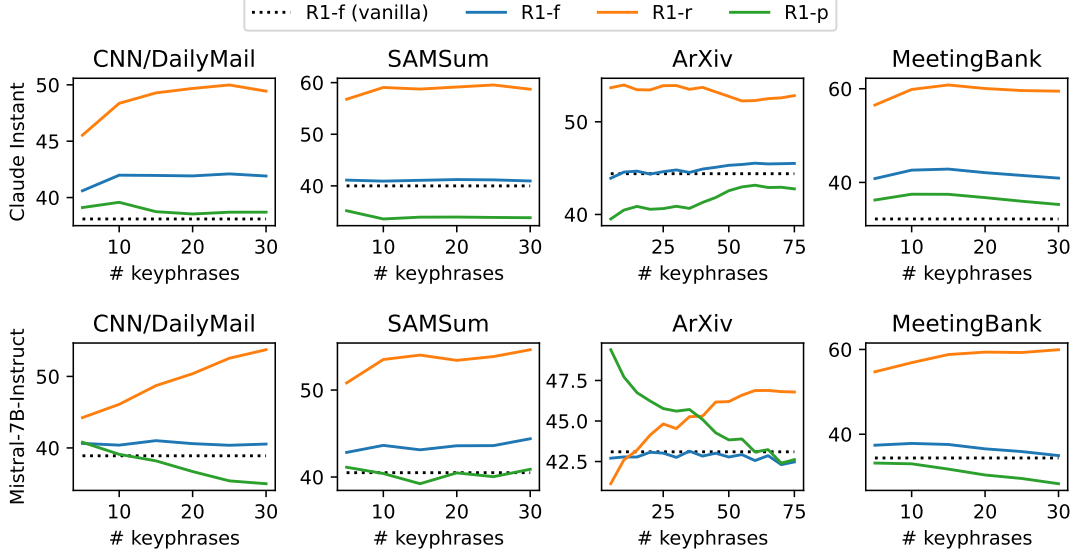


Figure 2: Effect of using different number of keyphrases on the precision-recall trade off.

	Raw ACU		Normalized ACU	
	Claude	+SigExt	Claude	+SigExt
CNN	43.8%	52.4%	40.7%	47.3%
SAMSum	53.6%	63.3%	38.4%	40.7%

Table 3: ACU coverage human evaluation on CNN and SAMSum using Claude Instant generated summaries.

scores of our model are stable when changing the number of keyphrases within a fairly wide range, showing that introducing keyphrases can consistently improve the summary quality.

As we increase the number of keyphrases, there is a clear trend of increasing recall and decreasing precision for the Mistral model. This is less evident for the Claude model. Since we add a length constraint explicitly in the prompt (e.g., "write a summary in 3 sentences"), the Claude model appears to follow these instructions better than the Mistral models. Mistral models tend to try to cover all the keywords provided in the prompt. Consequently, the recall increases significantly when increasing the number of keywords for the Mistral models.

3.4 Granularity of Salient Information

We also explore how different granularity of salient information can affect the summarization performance. We compare word-, phrase-, and sentence-level SigExt. The results are shown in Table 4. The phrase-level salient information can always achieve top or near-top performance, while the word-level and sentence-level approaches have much larger

variance. The word-level information performs poorly on the ArXiv dataset because for academic papers, there are many multi-word phrases that are important in the summary. If these are split, they are no longer helpful for summarization. In contrast, the sentence-level information is not so effective, especially on the MeetingBank dataset. When the dataset is highly abstractive, the important words are dispersed across the document, making it difficult to extract a few sentences to cover the content of the summary (See examples in Appendix Table 10).

Claude-Instant	R1-f	RL-f	R1-f	RL-f	
		SAMSum		CNN	
+SigExt (word)	41.4	30.9	42.0	26.2	
+SigExt (phrase)	41.6	30.9	42.0	26.6	
+SigExt (sent)	39.1	29.7	40.3	25.7	
		ArXiv		M.Bank	
+SigExt (word)	42.2	21.0	41.9	31.7	
+SigExt (phrase)	45.2	23.5	42.3	31.9	
+SigExt (sent)	44.8	23.8	36.2	25.8	

Table 4: Different granularity of salient information.

3.5 Summary Factuality

As shown in Table 5, the effect of adding keyphrases on the AlignScore is LLM and task-specific. For the Claude Instant and Falcon models, the AlignScore is typically improved by incorporating keyphrases. In contrast, the AlignScore always decreases for the Mistral model. These results suggest that keyphrases are not universally helpful for

improving the faithfulness of the generated summaries. Table 8 shows a few examples where hallucination is introduced in the summary due to the keyphrases. The failure pattern is if a keyphrase is negated in the document, Mistral model would ignore the negation.

	SamSum	CNN	ArXiv	M.Bank
Claude Ins.	85.8	83.8	53.7	73.1
+SigExt	88.0	82.3	60.0	74.7
Mistral-7B	88.9	88.8	56.9	79.1
+SigExt	84.7	87.0	49.5	77.1
Falcon-40B	81.6	67.7		
+SigExt	81.6	75.0		

Table 5: Summary factuality measured by AlignScore.

3.6 Introducing External Oracle Keyphrases

We also analyze how external keyphrases which do appear in the source document would affect the performance. We use oracle keyphrases that appears in the reference summary but do not appear in the source document as additional information in the prompt. The ROUGE-1 score and AlignScore are shown on Table 6. The ROUGE score increases significantly while the AlignScore falls. It indicates that introducing external keyphrases might hurt the factuality of the summary.

Claude-Ins.	R1-f	Align.	R1-f	Align.
	SamSum		CNN	
+SigExt	41.6	88.0	42.0	82.3
+Oracle	50.0	86.3	50.0	78.8
	ArXiv		M.Bank	
+SigExt	45.2	60.0	42.3	74.7
+Oracle	51.6	45.9	48.2	56.7

Table 6: Summary quality with oracle keyphrases.

3.7 Effectiveness of keyphrase extraction

In this part, we analyze the effectiveness of the Longformer keyphrase extractor. We define recall@ K metric to evaluate the keyphrase extraction performance. We define the recall@ K as the recall of *oracle keyphrases* in the top- K deduplicated keywords, where oracle keyphrases are constructed by finding the phrase in the source document with the highest fuzzy match score to each phrase in the target summary. We compare our method with two statistical methods, Rake (Rose et al., 2010) and TextRank (Mihalcea and Tarau,

2004). Recent work has proposed transformer-based keyphrase extraction models (Sun et al., 2020; Ding and Luo, 2021) that focus on generating noun phrases to better align with human annotation. However, in our setting, the oracle keyphrases are constructed heuristically and are not limited to noun phrases, making these models a poor fit for comparison. Therefore, we do not include them. The evaluation results are shown on Table 7. We show that GP-SigExt already outperforms statistical methods. And the fine-tuned SigExt achieves additional 5.9% and 3.7% improvements on two datasets respectively.

Method	SAMSum R@15	CNN R@15	M.Bank R@15	Arxiv R@35
Rake	68.3	11.9	17.1	14.2
TextRank	80.5	20.8	19.3	22.4
GP-SigExt (+32 ex.)	75.5	27.7	40.3	31.7
(+128 ex.)	81.5	29.7	47.3	32.0
	85.5	32.9	62.2	32.7
SigExt (1k ex.)	83.3	33.6	65.7	35.4

Table 7: Keyphrase extraction performance.

3.8 Case study

We show some examples in Appendix Table 9. We found the extracted keyphrases can help the LLM incorporate precise details in the summary, hence the summaries better align with the gold summary. In the first two examples, the keyphrases contain exact numbers and times, and the LLM was able to include them in the summary. In the third example, with SigExt, the summary covers more topics than the vanilla model. Since we instruct the LLM to “consider” these keyphrases, the LLM was able to skip or rephrase some to get more fluent results.

4 Related Work

Leveraging keyword in abstractive summarization has been explored in many works. Switching Generator-Pointer (Nallapati et al., 2016) and CopyNet (Gu et al., 2016) modify a recurrent neural network model (Chopra et al., 2016) to directly copy keywords from the source text. More recent work has adopted transformer architectures (Vaswani et al., 2017), which have become dominant in natural language processing. Liu et al. (2022) introduces a bias in the attention matrix to help transformer models focus on keywords. All these models need to be trained or finetuned on large-scale training data. While finetuned models typically

achieve higher ROUGE scores than prompting a pretrained model, prompt-based summarizers are preferred in some industrial use cases due to their flexibility and reduced need for data collection. Incorporating keyphrases in the prompt can effectively control the length and content coverage of the summary, something that fine-tuning methods cannot easily achieve. Therefore, we cannot compare with these methods using metrics like ROUGE.

Instruction finetuned LLMs (Chung et al., 2022; Touvron et al., 2023; Zhang et al., 2022) have shown strong performance on summarization purely via prompts, without finetuning data. Such models are often offered via APIs, enabling easier development and deployment of summarization applications. Keyphrases are still helpful for these large models, as Li et al. (2023a) show that a keyphrase generator trained with reinforcement learning can improve summarization performance.

There has been interest in 2-stage extractive-then-abstractive approaches (Su et al., 2020; Liu et al., 2021; Li et al., 2021; Su et al., 2022; Yang et al., 2023). These first extract keyphrases or sentences before abstractively summarizing them. These methods are trained end-to-end for domain-specific use cases, while our method can be pre-trained for general purpose zero-shot use cases. Practically, any keyword extractor, for example KeyBERT or LLMBERT (Grootendorst, 2020), can be used for the first stage to enhance the summarization in the second stage. The 2-stage methods could also be implemented as Chain-of-Thought (CoT) by generating intermediate hints and final results in the same prompt, such as Adams et al. (2023). In our experiments, we compare our method with a 2-stage prompting approach – first generating keywords using one prompt, then using those keywords for summarization in the second prompt. While slightly different from previous work, the 2-stage baseline effectively captures the use of intermediate reasoning steps of LLMs.

5 Conclusion

In this paper, we propose a lightweight approach to incorporate keyphrases into the prompt for LLM-based abstractive summarization. SigExt involves training a phrase extractor using supervised learning to identify salient keyphrases from the input text. These keyphrases are then injected into the prompt provided to the LLM for summary generation. We demonstrate that this approach can effec-

tively improve the ROUGE scores of the generated summaries, indicating a higher similarity to reference summaries. Introducing keyphrases in the prompt enhances the faithfulness of the summary by ensuring that important information is captured. Additionally, our approach offers control over the length and precision/recall trade-off of the summary. Notably, our pretrained keyphrase extractor – GP-SigExt – can improve summarization performance out-of-the-box without any finetuning, even in cases where training data is not available.

Limitations

Model Design: We use Longformer as the backbone model to build SigExt because it is lightweight and supports long context length. However, we do not evaluate the impact of using other similar-sized pre-trained language models. Additionally, we extract training labels using a fuzzy matching approach to make the model more generalizable, but more domain-specific approaches for keyphrase extraction may yield better performance.

Evaluation: As is common in summarization research, we rely primarily on automatic metrics and qualitative example checks to evaluate performance. These techniques have known limitations in assessing summary quality. Meanwhile, human evaluation has its own challenges. Therefore, how to best evaluate the quality of abstractive summarization models remain as an open question.

References

- Griffin Adams, Alex Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. [From sparse to dense: GPT-4 summarization with chain of density prompting](#). In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 68–74, Singapore. Association for Computational Linguistics.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#). *Preprint*, arXiv:2311.16867.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *Preprint*, arXiv:2004.05150.
- 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.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.
- Haoran Ding and Xiao Luo. 2021. Attentionrank: Un-supervised keyphrase extraction using self and cross attentions. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1919–1928.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Maarten Grootendorst. 2020. Keybert: Minimal keyword extraction with bert.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Yebowen Hu, Timothy Ganter, Hanieh Deilamsalehy, Franck Dernoncourt, Hassan Foroosh, and Fei Liu. 2023. MeetingBank: A benchmark dataset for meeting summarization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16409–16423, Toronto, Canada. 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, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. Mistral 7b. *Preprint*, arXiv:2310.06825.
- Anastassia Kornilova and Vladimir Eidelman. 2019. BillSum: A corpus for automatic summarization of US legislation. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 48–56, Hong Kong, China. Association for Computational Linguistics.
- Mahnaz Koupaee and William Yang Wang. 2018. Wikihow: A large scale text summarization dataset. *Preprint*, arXiv:1810.09305.
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Huang. 2023. A systematic study and comprehensive evaluation of ChatGPT on benchmark datasets. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 431–469, Toronto, Canada. Association for Computational Linguistics.
- Haoran Li, Arash Einolghozati, Srinivasan Iyer, Bhargavi Paranjape, Yashar Mehdad, Sonal Gupta, and Marjan Ghazvininejad. 2021. Ease: Extractive-abstractive summarization with explanations. *arXiv preprint arXiv:2105.06982*.
- Zekun Li, Baolin Peng, Pengcheng He, Michel Galley, Jianfeng Gao, and Xifeng Yan. 2023a. Guiding large language models via directional stimulus prompting. *arXiv preprint arXiv:2302.11520*.
- Zekun Li, Baolin Peng, Pengcheng He, Michel Galley, Jianfeng Gao, and Xifeng Yan. 2023b. Guiding large language models via directional stimulus prompting. *arXiv preprint arXiv:2302.11520*.
- Shuaiqi Liu, Jiannong Cao, Ruosong Yang, and Zhiyuan Wen. 2022. Key phrase aware transformer for abstractive summarization. *Information Processing & Management*, 59(3):102913.
- Yixin Liu, Alex Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, et al. 2023. Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4140–4170.
- Yizhu Liu, Qi Jia, and Kenny Zhu. 2021. Keyword-aware abstractive summarization by extracting set-level intermediate summaries. In *Proceedings of the Web Conference 2021*, pages 3042–3054.

- 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.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of the AAAI conference on artificial intelligence*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. 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.
- Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. Automatic keyword extraction from individual documents. *Text mining: applications and theory*, pages 1–20.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Eva Sharma, Chen Li, and Lu Wang. 2019. [BIG-PATENT: A large-scale dataset for abstractive and coherent summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2204–2213, Florence, Italy. Association for Computational Linguistics.
- Jing Su, Longxiang Zhang, Hamid Reza Hassanzadeh, and Thomas Schaaf. 2022. Extract and abstract with bart for clinical notes from doctor-patient conversations. *Proc. Interspeech 2022*, pages 2488–2492.
- Ming-Hsiang Su, Chung-Hsien Wu, and Hao-Tse Cheng. 2020. A two-stage transformer-based approach for variable-length abstractive summarization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2061–2072.
- Yi Sun, Hangping Qiu, Yu Zheng, Zhongwei Wang, and Chaoran Zhang. 2020. Sifrank: a new baseline for unsupervised keyphrase extraction based on pre-trained language model. *IEEE Access*, 8:10896–10906.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- 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.
- Xu Wang, Sen Wang, Xingxing Liang, Dawei Zhao, Jincui Huang, Xin Xu, Bin Dai, and Qiguang Miao. 2024. [Deep reinforcement learning: A survey](#). *IEEE Transactions on Neural Networks and Learning Systems*, 35(4):5064–5078.
- Wen Xiao, Iz Beltagy, Giuseppe Carenini, and Arman Cohan. 2022. [PRIMERA: Pyramid-based masked sentence pre-training for multi-document summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5245–5263, Dublin, Ireland. Association for Computational Linguistics.
- Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Ming Li, and David Lo. 2023. Apidocbooster: An extract-then-abstract framework leveraging large language models for augmenting api documentation. *arXiv preprint arXiv:2312.10934*.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. [Extractive summarization via ChatGPT for faithful summary generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3270–3278, Singapore. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2020. Pegasus: pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*, pages 11328–11339.
- Rui Zhang and Joel Tetreault. 2019. [This email could save your life: Introducing the task of email subject line generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 446–456, Florence, Italy. Association for Computational Linguistics.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Appendix

A All Prompts

Here we show all the prompts we used in the experiments. In prompt, <text> will be replaced with source documents, and <keywords> will be replaced with comma separated keyphrases extracted by SigExt. We conduct light prompt engineering to get a reasonably good zero-shot prompt.

A.1 Zero-shot Claude Instant Prompts

SAMSum

Here is a conversation:
<text>

Please write a very short 1 sentence summary.

SAMSum with SigExt

Here is a conversation:
<text>

Please write a very short 1 sentence summary. Consider include the following information: <keywords>

CNN/DailyMail

Here is a news article:
<text>

Please write a summary for the article in 2-3 sentences.

CNN/DailyMail with SigExt

Here is a news article:
<text>

Please write a summary for the article in 2-3 sentences. Consider include the following information: <keywords>.

ArXiv

Here is a research paper:
<text>

Please write a comprehensive paper abstract section.

ArXiv with SigExt

Here is a research paper:
<text>

Please write a comprehensive paper abstract section. Consider include the following information: <keywords>

MeetingBank

Here is a conversation:
<text>

Please write a summary in about 5 sentences.

MeetingBank with SigExt

Here is a conversation:
<text>

Please write a summary in about 5 sentences. Consider include the following information: <keywords>

A.2 Zero-shot Mistral Prompts

SAMSum

```
<s>[INST]Here is a conversation:  
<text>
```

```
Please write a short 1 sentence summary. [/INST]
```

SAMSum with SigExt

```
<s>[INST]Here is a conversation:  
<text>
```

```
Please write a short 1 sentence summary. Consider include the following information:  
<keywords>[/INST]
```

CNN/DailyMail

```
<s>[INST]Here is a news article:  
<text>
```

```
Please write a short summary for the article in 1-2 sentences. [/INST]
```

CNN/DailyMail with SigExt

```
<s>[INST]Here is a news article:  
<text>
```

```
Please write a short summary for the article in 1-2 sentences. Consider include the  
following information: <keywords>[/INST]
```

ArXiv

```
<s>[INST]Here is a research paper:  
<text>
```

```
Please write a short abstract in about 3 sentences. [/INST]
```

ArXiv with SigExt

```
<s>[INST]Here is a research paper:  
<text>
```

```
Please write a short abstract in about 3 sentences. Consider include the following  
information: <keywords>[/INST]
```

MeetingBank

```
<s>[INST]Here is a conversation:  
<text>
```

```
Please write a 2-3 sentence summary. [/INST]
```

MeetingBank with SigExt

```
<s>[INST]Here is a conversation:  
<text>
```

```
Please write a 2-3 sentence summary. Consider include the following information: <  
keywords>[/INST]
```

A.3 Zero-shot Falcon and Flan-T5 Prompts

SAMSum

```
Here is a conversation:  
<text>
```

```
Please write a short 1 sentence summary. Summary:
```

SAMSum with SigExt

Here is a conversation:
<text>

Please write a short 1 sentence summary. Consider include the following information:
<keywords>

Summary:

CNN/DailyMail

Here is a news article:
<text>

Please write a short summary for the article in 1-2 sentences.

Make sure the summary is no more than 2 sentences. Summary:

CNN/DailyMail with SigExt

Here is a news article:
<text>

Please write a short summary for the article in 1-2 sentences. Consider include the following information: <keywords>.

Make sure the summary is no more than 2 sentences. Summary:

ArXiv

Here is a research paper:
<text>

Please write a short abstract in about 3 sentences.

Abstract:

ArXiv with SigExt

Here is a research paper:
<text>

Please write a short abstract in about 3 sentences. Consider include the following information: <keywords>.

Abstract:

MeetingBank

Here is a conversation:
<text>

Please write a 2-3 sentence summary.

Summary:

MeetingBank with SigExt

Here is a conversation:
<text>

Please write a 2-3 sentence summary. Consider include the following information: <keywords>

Summary:

B Loss Function

The training objective for SigExt is

$$\mathcal{L} = - \sum_{i \in 1 \dots T} \sum_{k \in l_i \dots r_i} [\text{label}(p_i) \log f(x_k) + \lambda(1 - \text{label}(p_i)) \log(1 - f(x_k))],$$

where $f(x_k)$ denotes the binary classification probability on token x_k by the classification head, and λ is the class balancing weight.

Document	Jake: <i>we don't have english today</i> Nadia: whooooah! Vanessa: cool! you sure though?
Keyphrases	Jake: Smith called in sick, they couldn't find a replacement, oh yeah
Mistral	Vanessa; Smith called; <i>Jake; replacement;</i> english today
+SigExt	Jake announced that they don't have English class today due to Smith being sick and unable to find a replacement. Vanessa is excited that <i>Jake is teaching English today</i> instead of Smith, who called in sick and couldn't be replaced
Document	Sam: Good afternoon, is this available? Lia: Hi, how can I help you? Sam: I'm interested in this product. Are the earphones available? Lia: Yes, they are. Sam: where can I collect them? Lia: Well, you can have them delivered to you through courier - this costs £4 (pay on delivery) Lia: I'm afraid collecting them personally is a little difficult, as the product can only be collected in Garstang (where the company is located) Sam: Okay. Sam: Are they new? Do you have any other colors that white? Lia: Completely new - we also offer grey and pink. Sam: Then I would like the white ones please. Lia: Okay, so I just need your address and telephone Sam: The total costs £187 - can I lower the price to £180?? Lia: <i>fraid not - company policy.</i> Sam: Oh, okay. Sam: How long does it take for the product to arrive? Lia: It generally arrives the next day, in rare instances it's longer.
Keyphrases	Lia; earphones available; <i>Garstang; £180; collected;</i> Sam; product; white; delivered; company.
Mistral	Sam is interested in purchasing earphones and asks about their availability, collection options, and delivery time.
+SigExt	Lia confirmed that the earphones are available, can be delivered for £4, and are only available for collection in Garstang, but offered Sam the <i>option to purchase them for £180</i>
Document	Kim: I'm going to Seoul! Agatha: wow finally! Mark: When? Bring us some kimchi! Kim: <i>People on a plane will kill me if I do</i> :D Kim: I'm going in April <unk>3 Agatha: Are you going with Jane? Kim: yes, of course - she'll be our tour guide Mark: so jealous... I'd love to come as well Kim: Really? I asked you like a hundred times
Keyphrases	Seoul; Kim; Jane; Agatha; Mark; going; <i>kimchi;</i> tour guide; Kim.
Mistral	Kim is going to Seoul in April with Jane as their tour guide and Agatha and Mark express their excitement and jealousy.
+SigExt	Kim is going to Seoul in April with Jane as their tour guide, and <i>they will bring some kimchi</i> back for Mark who is jealous and wishes he could come too

Table 8: Examples where adding keyphrases in prompt hurts the faithfulness of the summary. The failure pattern is if a keyphrase is negated in the document, the LLM would ignore the negation.

Gold Summary	NHS gave out 404,500 prescriptions for suncream at a cost of £13m in 2014. Also handed out 4.7million prescriptions for indigestion pills costing £29m. Other items routinely prescribed include vitamins, Vaseline and toothpaste. Critics branded prescriptions 'ludicrous' at time of financial crisis for NHS.
Key Phrases	toothpaste; suncream; patients; <i>404,500 prescriptions</i> ; NHS; <i>indigestion pills costing£29million</i> ; NHS handed; <i>£13million</i> ; spending millions; prescribed
Mistral-7B	The NHS is spending millions on prescriptions for suncream and other household essentials, according to a damning analysis. The figures come as the UK ranks below many European countries in health and wellness, and as the NHS faces a huge financial crisis.
+SigExt	The NHS is spending millions of pounds on prescribing <i>toothpaste</i> and suncream to patients, according to a damning analysis. In 2014, the NHS handed out <i>404,500 prescriptions</i> for suncream at a cost of <i>£13million</i> and 4.7million prescriptions for <i>indigestion pills costing£29million</i> .
Gold Summary	Cloned camel was born in 2009 and called Injaz, which means achievement. She was cloned from ovarian cells and born by surrogate mother. Injaz is now six years old and is said to have conceived naturally. She's expected to give birth late this year, proving cloned animals' fertility.
Key Phrases	Injaz; pregnant; slaughtered camel; cloned; ovarian cells; surrogate mother; first cloned camel; <i>2009</i> ; <i>conceived naturally</i> ; whose name
Mistral-7B	The world's first cloned camel, Injaz, is pregnant and expected to have a calf. This marks the first time that a cloned camel has been able to reproduce naturally, and scientists are excited to see if cloned camels are as fertile as naturally produced ones.
+SigExt	Injaz, the world's first cloned camel born from the ovarian cells of a slaughtered camel in <i>2009</i> , is now pregnant and expected to give birth to a calf later this year. The camel's <i>conception was natural</i> , proving that cloned camels are fertile and capable of reproducing.
Gold Summary	Lewis Hamilton won Sunday's Bahrain Grand Prix ahead of Kimi Raikkonen. Hamilton is out of contract at the end of the year and is yet to sign new deal. It has been suggested that Hamilton could replace Raikkonen at Ferrari. But team principal Maurizio Arrivabene says he is happy with driver line-up.
Key Phrases	Kimi Raikkonen; Ferrari; <i>next season</i> ; <i>replace Raikkonen</i> ; Mercedes; 30; Lewis Hamilton yet; Bahrain; Ferrari team principal Maurizio Arrivabene; Vettel
Mistral-7B	Ferrari team principal Maurizio Arrivabene has revealed that he uses a carrot-and-stick method with Kimi Raikkonen to keep his desire to remain with the Maranello marque high. The Finn has arguably never been happier during his Formula One career, thanks to the atmosphere generated within the team by the arrival of Arrivabene.
+SigExt	Ferrari team principal Maurizio Arrivabene has revealed that he uses a carrot-and-stick method to keep Kimi Raikkonen's desire to remain with the Maranello marque high. With Lewis Hamilton yet to sign a new contract, it has been suggested that he will <i>replace Raikkonen</i> at Ferrari for <i>next season</i> . In the Bahrain Grand Prix, Raikkonen finished on the podium for the first time in 26 races as runner-up to Hamilton.

Table 9: Examples of using SigExt with Mistral-7B model on CNN dataset.

	<p><i>Jenkin</i> : hey what is your <i>spirit animal</i> ? <i>Sophie</i> : what ? <i>Jenkin</i> : go on ? <i>Sophie</i> : I dont know a <i>fox</i> lol <i>Jenkin</i> : are you wiley ? <i>Sophie</i> : sometimes <i>Jenkin</i> : I am a <i>Sophie</i> : I think you are a bit mad like the mad</p>
Document	<p><i>Jenkin</i> : I have been <i>reading</i> about <i>animal spirits</i> its quite good <i>Sophie</i> : you will have to tell me about the <i>fox</i> .. do you decide what your <i>animal</i> is or does someone tell you ? <i>Jenkin</i> : There is a <i>pack</i> of <i>cards</i> and you <i>choose</i> the one that you are <i>drawn</i> to <i>Sophie</i> : oh right I <i>would choose</i> the <i>Fox</i> <i>Jenkin</i> : well I did n't know but I was <i>drawn</i> to the <i>dolphin</i> <i>Sophie</i> : oh <i>Jenkin</i> : I will <i>bring</i> them over <i>tomorrow</i> <i>Sophie</i> : oh yes please that will be great</p>
Reference	<p>Jenkin has been reading about spirit animals and he was drawn to a dolphin. Sophie would choose a fox. Jenkin will bring pack of cards with spirit animals to Sophie tomorrow.</p>
	<p><i>Jacky</i> : I think you were <i>right</i> yesterday . <i>David</i> : What about ? I 'm <i>right</i> about most <i>things</i> : P <i>Jacky</i> : Yeah , whole you ;) <i>Jacky</i> : About <i>taking</i> the <i>blame</i> etc . <i>David</i> : Okey , I remeber . We 'll <i>talk</i> later ? <i>Jacky</i> : With pleasure . I 'll call you when I get <i>home</i> .</p>
Document	
Reference	<p>According to Jacky, David did the right thing taking the blame. They will talk when Jack comes back home.</p>
	<p><i>Jill</i> : So <i>bored</i> ! <i>Nate</i> : Well ... ca n't help you there <i>Nate</i> : Still at <i>work</i> <i>Jill</i> : ugh I need to find a job <i>Jill</i> : I 've <i>watched</i> everything on <i>youtube</i> already <i>Nate</i> : Doubt it : P I 'll <i>call</i> you when I get off <i>work</i></p>
Document	
Reference	<p>Jill is bored and has watched YouTube. Nate is at work and will call Jill when he finishes it.</p>

Table 10: Visualization of overlapping words between the document and reference summary on the SAMSum dataset. The words are dispersed across the document, making it difficult to extract sentence-level salient information.

Dataset	Description	Input/Output
CNN	News article headline generation	773/58
SAMSum	Messenger-like conversations summarization	127/23
ArXiv	Research paper abstract generation	6446/166
MeetingBank	Meeting transcript summarization	3095/66

Table 11: Dataset description and input/output length.