

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

Lei Xu, Mohammed Asad Karim, Saket Dingliwal, Aparna Elangovan
AWS AI Labs
{ leixx, asadkr, skdin, aeg } @amazon.com

Motivation

Incorporating salient information has proven to be an effective approach to enhance text summarization performance.

Adapting this strategy to prompt-based summarization presents unique opportunities, particularly in guiding LLMs to generate summaries that fit different industrial use cases.

Our analysis provides valuable insights for applying these methods in real-world summarization applications, aiming to improve relevance and completeness in generated summaries.

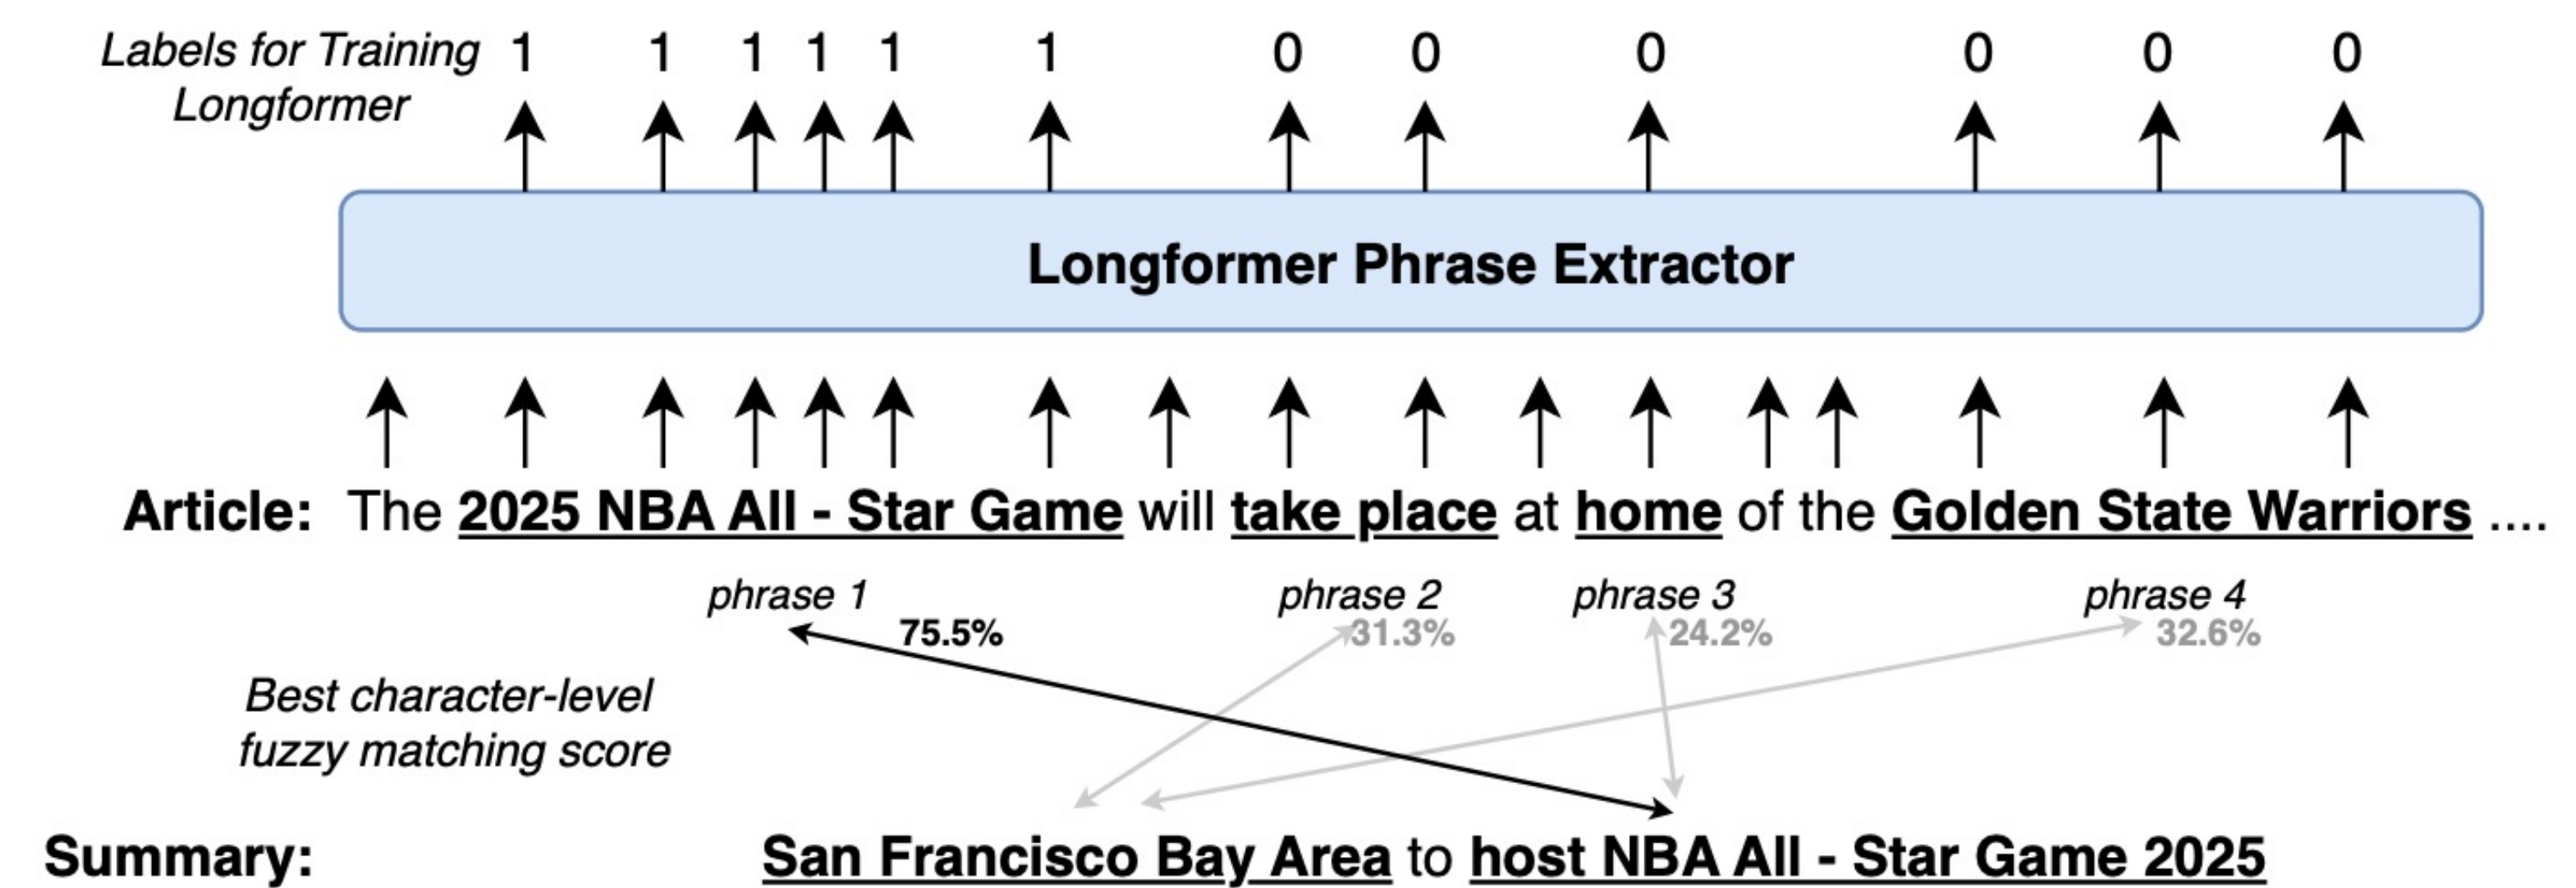
SigExt Method

- We use a fine-tuned Longformer to extract keyphrases from source documents.
- The training labels are obtained by aligning the source document with reference summary.
- During inference, top-K deduplicated keyphrases are added in the prompt.

GP-SigExt:

- We pre-trained the Longformer on 7 datasets, and found the pre-trained keyphrase extractor has good generalization capability.

Figure 1: SigExt Keyphrase Extractor



Prompt Example

Here is a news article: `<text>`
Please write a summary for the article in 2-3 sentences.
Consider include the following information: `<keyphrases>`

Experiment Settings

- We evaluate SigExt on 4 summarization datasets spanning across formal/conversational long/short documents.
- We evaluate 3 LLMs -- Claude, Mistral and Falcon.
- We use ROUGE and AlignScore metrics, and ACU-base human evaluation.

Table 1: Human evaluation using atomic content units (ACUs).

	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 2: Adding keyphrases extracted using SigExt into prompts can improve the alignment between generation and reference.

Method	SAMSum			CNN/DailyMail			ArXiv			MeetingBank			Avg. $\Delta R1-f$
	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	R1-f	RL-f	R1-r	
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

Figure 2: Precision-recall trade off by changing the number of keyphrases.




Table 3: Phrases are the optimal granularity for salient information compared to words or sentences.

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: Using salient information in prompt may increase hallucination (evaluated by AlignScore).

	SamSum	CNN	ArXiv	M.Bank	Claude-Ins.	R1-f	Align.	R1-f	Align.
	Claude Ins.	85.8	83.8	53.7	73.1				
+SigExt	88.0	82.3	60.0	74.7					
Mistral-7B									
+SigExt	88.9	88.8	56.9	79.1					
+Oracle	84.7	87.0	49.5	77.1					
Falcon-40B									
+SigExt	81.6	67.7							
+Oracle	81.6	75.0							

Salient information may introduce hallucination

Document:

Jake: *we don't have english today*
Nadia: *whoooooah!*
Vanessa: *cool! you sure though?*
Jake: *Smith called in sick, they couldn't find a replacement, oh yeah*

Keyphrases:

Vanessa; Smith called; Jake; replacement; english today

Summary (Mistral vanilla):

Jake announced that they don't have English class today due to Smith being sick and unable to find a replacement.

Summary (Mistral + SigExt):

Vanessa is excited that Jake is teaching English today instead of Smith, who called in sick and couldn't be replaced

Salient information improves completeness

Keyphrases:

toothpaste; suncream; patients; **404,500 prescriptions**; NHS; **indigestion pills costing £29million**; NHS handed; **£13million**; spending millions; prescribed

Summary (Mistral vanilla):

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.

Summary (Mistral + SigExt):

The NHS is spending millions of pounds on prescribing **toothpaste** and suncream to patients, according to a damning analysis. In 2014, the NHS handed out **404,500 prescriptions** for suncream at a cost of **£13million** and 4.7million prescriptions for **indigestion pills costing £29million**.



Contributions

- We provide a comprehensive analysis on the impact of adding salient information in prompts for summarization
 - Summary length:** salient information can be used to control summary length.
 - Reference alignment:** salient information can improve reference alignment.
 - Completeness:** salient information can improve completeness.
 - Hallucination:** salient information may increase hallucination depending on the underlying LLM.
- We present SigExt, a simple yet effective keyphrase extraction model using a finetuned Longformer. Once trained, it is LLM-agnostic, enabling performance boost for different LLMs by adding extracted keyphrases in prompts without requiring LLM finetuning.
- We demonstrate that SigExt has cross-domain generalization capability through a general-purpose version (GP-SigExt) pretrained on 7 datasets.