

On Context Utilization in Summarization with Large Language Models

Anonymous ACL submission

Abstract

Large language models (LLMs) excel in abstractive summarization tasks, delivering fluent and pertinent summaries. Recent advancements have extended their capabilities to handle long-input contexts, exceeding 100k tokens. However, in question answering, language models exhibit uneven utilization of their input context. They tend to favor the initial and final segments, resulting in a U-shaped performance pattern concerning where the answer is located within the input. This bias raises concerns, particularly in summarization where crucial content may be dispersed throughout the source document(s). Besides, in summarization, mapping facts from the source to the summary is not trivial as salient content is usually re-phrased. In this paper, we conduct the first comprehensive study on context utilization and position bias in summarization. Our analysis encompasses 5 LLMs, 10 datasets, and 5 evaluation metrics. We introduce a new evaluation benchmark called MiddleSum on the which we benchmark two alternative inference methods to alleviate position bias: hierarchical summarization and incremental summarization¹.

1 Introduction

Large language models (LLMs) have drastically transformed the landscape of NLP recently (Brown et al., 2020). With instruction tuning (Ouyang et al., 2022; Chung et al., 2022), LLMs made a major leap forward in conditional (prompted) content generation, and can generate satisfying outputs without the need to finetune on a task specifically. In abstractive summarization specifically, this approach has arguably opened a new paradigm; summaries generated by LLMs are highly fluent, grammatical and relevant (Goyal et al., 2022). Despite noticeably lower scores on automatic metrics such as ROUGE (Lin, 2004) or BERTScore (Zhang et al., 2019), summaries generated by LLMs are

largely preferred by humans over summaries from state-of-the-art fine-tuned models like BRIO (Liu et al., 2022b, 2023c). In fact, on XSum, GPT-3.5 summaries are even on par with re-annotated human-written summaries, and much better than the dataset ground-truth, according to human evaluators (Zhang et al., 2023b). LLMs also show promising capability in evaluating summaries generated by other systems, including other LLMs (Fu et al., 2023; Luo et al., 2023; Shen et al., 2023a).

Despite this success, a few major technological bottlenecks remain with LLMs, including the maximum length of their context window. The standard context window length for open-source LLMs is 2k tokens (Brown et al., 2020; Scao et al., 2022; Penedo et al., 2023; Touvron et al., 2023a), which drastically limits their usefulness for long-input summarization (Shaham et al., 2022). Several techniques were proposed to extend the context window, including ALiBi (Press et al., 2021), LeX (Sun et al., 2022), position interpolation (Chen et al., 2023) and YaRN (Peng et al., 2023). While some of them claim up to 128k tokens processing capacity (Peng et al., 2023), it remains unclear how much such methods help on long-context summarization.

Scaling up context length would only succeed if a key question gets addressed first: *do LLMs make proper use of their entire context?* Recent work (Liu et al., 2023a) suggested that, surprisingly, such a simple assumption may not hold: through experiments on multi-document question answering and key-value retrieval, the authors find that LLMs mostly focus on the *beginning* and *end* of the (long) context window. Plotting performance with regards to the position of the important information exhibits a *U-shape*, with performance high at first (beginning of the source), then dropping, and rising again at the end. Worryingly, in the middle of the context window, LLMs’ performance can drop to even *below random chance*, calling for greater examination of LLMs’ behaviors with regard to the

¹We will share all code and data upon publication.

position of information within the source.

In this work, we investigate in depth how LLMs use their context window in abstractive summarization. Unlike in question-answering, mapping facts in the output to a specific snippet in the source is not straightforward in abstractive summarization, due to the high-level of re-phrasing and compression. We conduct a large-scale study with 5 open-source LLMs, 10 datasets covering many aspects of summarization, and 5 highly diverse automatic metrics. Our contributions are threefold:

- We conduct the first large-scale analysis on context utilization in abstractive summarization, and the impact of the position of salient information on performance. We show that the U-shape or *middle-curse* exhibited by (Liu et al., 2023a) also holds in abstractive summarization.
- We craft an evaluation dataset (**MiddleSum**) where important information is concentrated in the middle of the context, enabling us to automatically quantify how much LLMs are affected by the *middle-curse*.
- We benchmark two alternative methods for inference on MiddleSum: *hierarchical summarization* and *incremental summarization*, showing their promise at alleviating the *middle curse* (especially in the scientific paper domain).

2 Experimental Setup

Datasets We cover a broad set of diverse abstractive summarization tasks, varying length and domain. We include 5 datasets of standard length (source is below 2k tokens, which always fits in the context window): (i) **CNN/DailyMail** (Hermann et al., 2015), (ii) **XSum** (Narayan et al., 2018), (iii) **Reddit-TIFU** (Kim et al., 2019), (iv) **SAMSum** (Gliwa et al., 2019), and (v) **Multi-XScience** (Lu et al., 2020). We also include another 5 long-input summarization datasets: (i) **Arxiv** and (ii) **PubMed** (Cohan et al., 2018), (iii) **GovReport** (Huang et al., 2021), (iv) **SummScreenFD** (Chen et al., 2022), and (v) **Multi-News** (Fabbri et al., 2019). A high-level view of each dataset is shown in Table 1, and detailed statistics are presented in Appendix A. For all datasets, we run experiments on the test set, subsampling 1,000 data points if its size is greater than 1,000, or using the entire test set otherwise.

Models We experiment with several popular high-performing *open-source* instruction-tuned

Dataset	Input length	# Documents		In Flan?	
	Standard Long	Single	Multi	Yes	No
CNN/DM (Hermann et al., 2015)	✓	✓		✓	
XSum (Narayan et al., 2018)	✓	✓		✓	
Reddit-TIFU (Kim et al., 2019)	✓	✓			✓
SAMSum (Gliwa et al., 2019)	✓	✓		✓	
Multi-XScience (Lu et al., 2020)	✓		✓		✓
Arxiv (Cohan et al., 2018)		✓		✓	✓
PubMed (Cohan et al., 2018)		✓		✓	✓
GovReport (Huang et al., 2021)		✓		✓	✓
SummScreenFD (Chen et al., 2022)		✓		✓	✓
Multi-News (Fabbri et al., 2019)		✓	✓	✓	✓

Table 1: Summarization datasets under study. In standard length datasets, the context and summary fit within a 2k tokens LLM context window.

LLMs. Instruction-tuning datasets such as Flan (Wei et al., 2021) include some of the datasets we study: CNN/DM, XSum, SAMSum and Multi-News. We analyze the following models through HuggingFace *transformers* (Wolf et al., 2020):

- **Flan-UL2** is a 20B parameters encoder-decoder model pre-trained on 1T tokens. It is based on the UL2 20B model (Tay et al., 2022), with the addition of Flan-T5 (Chung et al., 2022) instruction fine-tuning. The context window is 2k.
- **Llama-2** (Touvron et al., 2023b) is a recently introduced powerful decoder-only model pre-trained on 2T tokens, ranging from 7B to 70B parameters, and with a 4k context window. We use the **7B** and **13B** models.
- **Xgen-7B** (Nijkamp et al., 2023) is a 7B decoder-only model pre-trained on up to 1.5T tokens. It supports 8k tokens context window.
- **Mistral-7B** (Jiang et al., 2023a) is also an 8k-context 7B decoder-only model, with performance slightly better than Llama-2-13B.

We use the instruction-tuned (or *chat*) checkpoints for Llama-2, Xgen-7B and Mistral-7B. To run inference, we use the following prompt: Read the following text and summarize it: [text]. Summarize the above text in [n] sentences. Summary: where n is set to an average number of target sentences per dataset (see Appendix A). We infer all models in *bfloat16* and sample summaries with top-k sampling (Fan et al., 2018) using $k = 50$.

Evaluation Measures Summarization evaluation is especially challenging in the LLM era, as most automatic metrics poorly correlate with human preferences (Goyal et al., 2022; Liu et al., 2022a). To get a broad picture of performance, we evaluate with metrics as diverse as possible. First,

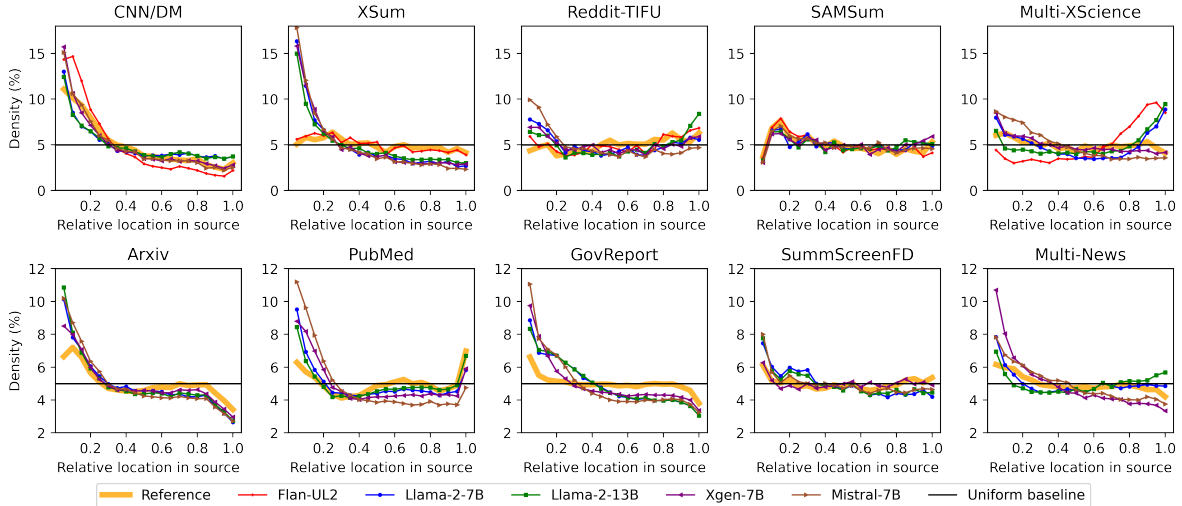


Figure 1: Distribution of the relative location of summary bigrams within the source. We split each source document into 20 bins of the same number of words, and plot the distribution of summary bigrams over source bins.

we consider **reference-based** metrics: **ROUGE-2** (Lin, 2004), which measures bigram overlap, **BERTScore** (Zhang et al., 2019), which measures semantic similarity with BERT (Devlin et al., 2018) embeddings, and **A3CU** (Liu et al., 2023c), which extracts facts in the form of Atomic Content Units (ACUs) (Liu et al., 2022a), and checks the presence of ACUs between prediction and reference. As **reference-free** metrics, we include **SummaC** (Laban et al., 2022), a leading factual consistency evaluation metric relying on entailment scores between pairs of source and summary sentences. We also leverage **GPT-3.5** (gpt-3.5-turbo-0125) as a summarization evaluator, which is proven to be a strong natural language generation evaluator (Wang et al., 2023a; Shen et al., 2023a; Jain et al., 2023). We prompt the model with the source and generated summary (which fits in GPT-3.5’s 16k context window) and ask to output a score on a likert scale from 1 to 5. We refer to Appendix B for the full prompt template. For GPT-3.5 evaluation, we subsample 300 data points per dataset in order to reduce costs.

We report the performance of LLMs on all 10 datasets, alongside a comparison to SOTA, in Appendix C. FLan-UL2 dominates on standard-length datasets, but Llama-2-13B has the upper hand on the long-input ones. Performance itself is not our focus in this paper, but rather which position-related factors influence it. We discard FLan-UL2 on long-input datasets due to poor performance.

3 Experiments

In this section, we describe a series of experiments aimed at understanding how LLMs treat information in their input depending on the position.

3.1 RQ1: Where in the source do LLMs take their information from?

We investigate summaries generated by LLMs, and map them to specific parts of the input. Unlike in question-answering or *extractive* summarization, mapping salient information from a summary to the source is not trivial in *abstractive* summarization.

We follow the approach used in (Kim et al., 2019; Zhao et al., 2022a) and compute the relative position of bigrams from generated summaries within the source documents, as a proxy for the position of salient information. We only use unique bigrams from summaries, and for each bigram, find all its occurrences within the source, if there are any. We then split the source into 20 bins of the same number of words, and compute the fraction of matched bigrams found in each bin. On top of the LLMs described above, we include the position of bigrams from reference summaries, and a uniform baseline.

As seen in Fig. 1, all summarization datasets except XSum, and Reddit-TIFU show some *lead bias*: salient bigrams from the reference (orange curves) are more likely to be found at the beginning of the source. However, LLMs show a significantly stronger lead bias on all datasets: bigrams from LLMs summaries are much more likely to be found in the first 20% words of the source. It is especially striking on XSum (except for FLan-UL2), Reddit-TIFU, Arxiv, PubMed and GovReport. On XSum, FLan-UL2 closely matches the reference distribution, which we attribute to its better instruction tuning. Results in Appendix D confirm that bigram distribution for LLMs and references are statistically different (p-value of Kolmogorov-




Figure 2: Distribution of relative location of input context sentences aligned with sentences from summaries. X-axis corresponds to the source sentence bin, y-axis to the fraction of aligned sentences in each bin.

Smirnov test (Massey Jr, 1951) inferior to 0.001) in all cases except Flan-UL2 on XSum and SAM-Sum and Llama-2-7B on XSum. We conclude that **LLMs focus on contents at the beginning of the source document(s)**.

3.2 RQ2: Where do LLMs look at within their context window?

In the previous experiment’s design, LLMs may not see the entire source in long-input summarization datasets, due to their limited context window, which is shorter than the source on the long-input datasets. We now focus on input information accessible to LLMs, and only consider salient information if it falls within the context window. Besides, since the same bigram may occur multiple times throughout the source, we adjust the methodology for saliency estimation. We align sentences in generated summaries to sentences in the context, following the procedure described in (Zhou et al., 2018) and also used in (Adams et al., 2023). Specifically, we greedily select source sentences until the ROUGE-1 F1 score between the set of selected source sentences and the summary stops increasing. The resulting set of source sentences forms a proxy of the visible salient input information being rephrased by the model when summarizing. We split each truncated source document into 10 bins of the same number of sentences, and map each aligned source sentence to its bin. Note that bins are not directly comparable across models, as context length varies across models.

As we can see in Fig. 2, sentences from the first 10% or last 10% of the input context are much more represented than others. A clear U-shape emerges on PubMed and SummScreenFD for all LLMs. This is intriguing knowing that Llama-2 LLMs and the other two LLMs have different context window lengths, and the last 10% of each context window may contain content of varying saliency. In other words, **LLMs seem to be mostly**

re-phrasing information from the beginning or the end of their context window.

3.3 RQ3: Does LLMs performance depend on the position of salient information?

Results from the last experiment raise the question of whether LLMs’ summarization performance changes depending on *where salient information is located within the input*. As an approximation for salient information, we consider the alignment between summary sentences and source sentences like in Fig. 2, but this time using the *reference* summaries. Each reference summary is mapped to the source sentences it maximizes ROUGE-1 F1 against, which may be scattered across the whole source. We convert each source sentence to its cumulative word count from the beginning of the source, and take the average as an approximation of the mean position of salient information within the source. We keep data points with mean salient position fitting within the LLM context window.

We examine performance changes with regard to this salient position. To do so, we compute the Spearman correlation coefficient between salient position and each evaluation metric in Table 2. A high absolute Spearman value means that summary quality (as measured by this metric) can change (and deteriorate) with the position of important information within the context. There are several takeaway findings from this Table. First, we notice that on standard-length datasets, reference-based evaluation metrics are negatively correlated to position of salient information. The correlation is only moderate, yet remarkably consistent across datasets (except Reddit-TIFU) and models. This is surprising, since such datasets fit entirely in context and are not affected by truncation. In contrast, reference-free metrics show either no significant or *positive* correlation to information position. For long-input datasets, the negative trend for reference-based metrics is confirmed. On these

Metric	Model	CNN/DM	XSum	Reddit	SAMSum	Multi-X	AVG	Arxiv	PubMed	GovReport	SummScreenFD	Multi-N	AVG
ROUGE-2	Flan-UL2	-0.296	-0.124	0.048	-0.069	-0.201	-0.128	-	-	-	-	-	-
	Llama-2-7B	-0.160	-0.023	0.063	-0.059	-0.100	-0.056	0.022	-0.113	-0.109	-0.079	-0.210	-0.098
	Llama-2-13B	-0.166	-0.086	0.031	-0.078	-0.039	-0.068	-0.017	-0.081	-0.166	-0.139	-0.213	-0.123
	Xgen-7B	-0.228	-0.042	0.066	-0.039	-0.041	-0.056	0.028	-0.091	-0.405	0.063	-0.283	-0.138
	Mistral-7B	-0.289	-0.031	0.006	-0.024	-0.052	-0.078	-0.270	-0.279	-0.585	-0.132	-0.324	-0.318
BERTScore	Flan-UL2	-0.331	-0.185	0.062	-0.144	-0.399	-0.187	-	-	-	-	-	-
	Llama-2-7B	-0.173	-0.012	0.062	-0.130	-0.385	-0.128	-0.031	-0.203	-0.104	-0.067	-0.256	-0.132
	Llama-2-13B	-0.193	-0.102	0.038	-0.089	-0.352	-0.140	-0.082	-0.209	-0.063	-0.152	-0.279	-0.157
	Xgen-7B	-0.252	-0.106	0.046	-0.075	-0.343	-0.146	-0.017	-0.125	-0.353	-0.093	-0.345	-0.187
	Mistral-7B	-0.278	-0.052	0.014	-0.108	-0.416	-0.168	-0.348	-0.367	-0.567	-0.356	-0.403	-0.408
A3CU	Flan-UL2	-0.258	-0.090	0.050	-0.123	-0.069	-0.098	-	-	-	-	-	-
	Llama-2-7B	-0.182	-0.076	0.028	-0.121	-0.090	-0.088	-0.038	-0.209	-0.154	-0.129	-0.217	-0.149
	Llama-2-13B	-0.190	-0.098	0.009	-0.166	0.016	-0.086	-0.104	-0.232	-0.111	-0.198	-0.228	-0.175
	Xgen-7B	-0.212	-0.130	0.023	-0.126	-0.025	-0.094	-0.036	-0.255	-0.211	-0.076	-0.287	-0.173
	Mistral-7B	-0.291	-0.110	0.004	-0.105	-0.119	-0.126	-0.160	-0.283	-0.305	-0.010	-0.283	-0.208
SummaC	Flan-UL2	-0.012	0.548	0.270	0.186	-0.035	0.191	-	-	-	-	-	-
	Llama-2-7B	0.088	0.552	0.375	0.227	0.224	0.293	0.090	0.108	0.126	-0.020	0.205	0.102
	Llama-2-13B	0.162	0.556	0.394	0.173	0.096	0.276	0.090	0.265	0.192	-0.144	0.232	0.127
	Xgen-7B	0.001	0.161	0.220	0.117	0.004	0.101	-0.208	-0.087	-0.313	-0.141	0.046	-0.141
	Mistral-7B	-0.045	0.515	0.149	0.069	0.154	0.128	-0.250	-0.103	-0.387	0.124	-0.010	-0.125
GPT-3.5	Flan-UL2	-0.020	0.196	0.027	-0.009	-0.193	0.000	-	-	-	-	-	-
	Llama-2-7B	0.036	0.036	-0.152	0.077	-0.153	-0.031	0.008	-0.116	-0.013	-0.068	-0.120	-0.062
	Llama-2-13B	0.039	0.072	-0.038	0.066	-0.052	0.017	-0.010	-0.084	-0.084	-0.060	-0.051	-0.058
	Xgen-7B	0.056	-0.007	-0.101	0.006	-0.174	-0.044	-0.058	-0.096	-0.317	-0.063	-0.055	-0.118
	Mistral-7B	-0.115	0.124	-0.133	0.036	-0.204	-0.058	-0.446	-0.322	-0.580	-0.188	-0.163	-0.342

Table 2: Spearman correlation coefficient between each LLM’s metric, and the mean position of salient information within the context window. Flan-UL2 is not applied to long-context summarization datasets due to its too short context window. **Multi-X** is short for Multi-XScience, **Multi-N** is Multi-News dataset, **AVG** columns represent the average over standard-length and long-input datasets, respectively. Numbers in gray correspond to non-significant Spearman scores (p-value greater than 0.05).

lengthy datasets, SummaC and GPT-3.5 tend to switch from positive to negative correlation, especially for Xgen-7B and Mistral-7B. We highlight that since GPT-3.5 itself is affected by the *middle-curse* from Liu et al. (2023a), it may not accurately evaluate summarization when salient content lays in the middle of the context. In light of these results, we conclude that *LLMs’ summarization performance is sensitive to the position of salient information in the context window*.

4 Analysis

4.1 How is information in the *middle* treated?

Previous experiments show that LLMs place more emphasis on the beginning and the end of their context. We now narrow down on how LLMs treat the *middle*. To remove the effect of spread of salient information, we perform two controlled experiments in multi-document summarization. This setup enables us to shuffle the order of the input, which is not realistic for the single-document setup as it would break coherence. We only consider data points with the same number k of documents: $k = 7$ documents for Multi-XScience ($n = 329$), and $k = 5$ documents for Multi-News ($n = 219$).

In the first experiment, we vary the position of salient information throughout the input. We keep a single document (the abstract of the query paper on Multi-XScience, and the document with the highest BERTScore with the reference on Multi-News), and place it at position j for $j \in \{1, \dots, k\}$, using

$k - 1$ documents from a random data point for the other slots. The single relevant document is accompanied by a [RELEVANT] header, while the other documents have an [IRRELEVANT] header, and we prompt the LLM to only summarize the relevant document. For reference-free evaluation metrics, we use the single relevant document as source. In this experiment and the following, we also include a *random* baseline of shuffled inputs and model predictions. In Fig. 3, we see a noticeable drop in performance for all metrics when the salient document is not in the first or final position. Flan-UL2 seems to focus on the *end* of the context, Xgen-7B and Mistral-7B on the *beginning*, and Llama-2 models on *both*. Performance can fall quite below random range, especially for SummaC, confirming the worrying trend from Liu et al. (2023a).

A more fine-grained analysis with GPT-3.5 in Fig. 4 evaluating specific attributes (following the method in Adams et al. (2023)) reveals more details. *Coherence* and *quality* remain high and stable. In other words, the text outputs are of good quality. But for *informativeness* and *attributability*, the U-shape appears again: it shows that the LLMs are struggling to generate content sticking to the document inserted in the middle.

In the second experiment, we take the opposite approach, and put salient information at the beginning and the end, while the middle of the prompt is filled with noise. We keep the first and last documents, and fill the $k - 2$ middle ones with random

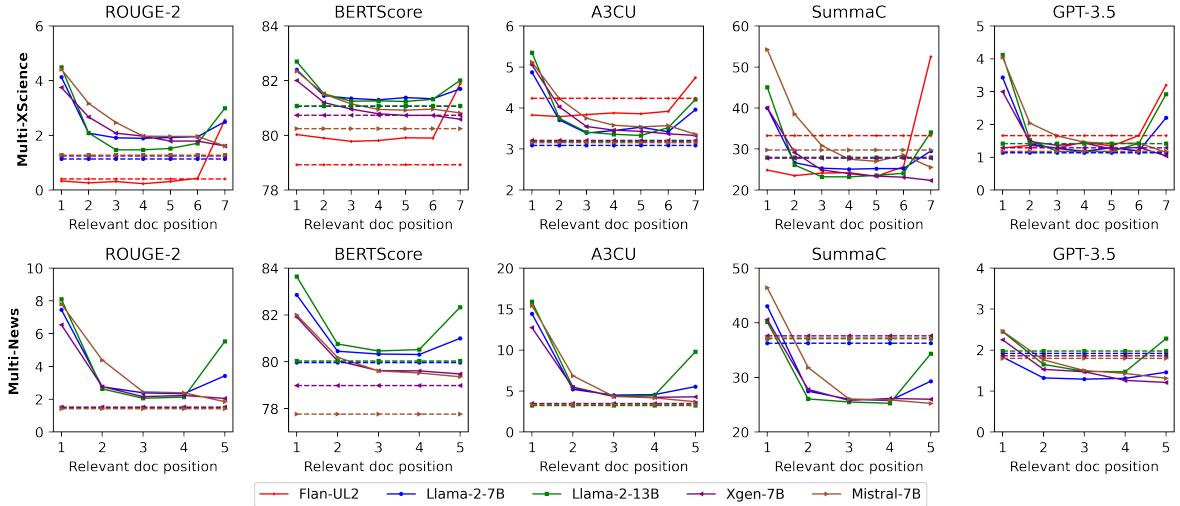


Figure 3: Multi-document summarization performance on Multi-XScience (top row) and Multi-News (bottom row) when a unique relevant document is used, and its position is varied (x-axis). Dashed horizontal lines correspond to the random baseline.

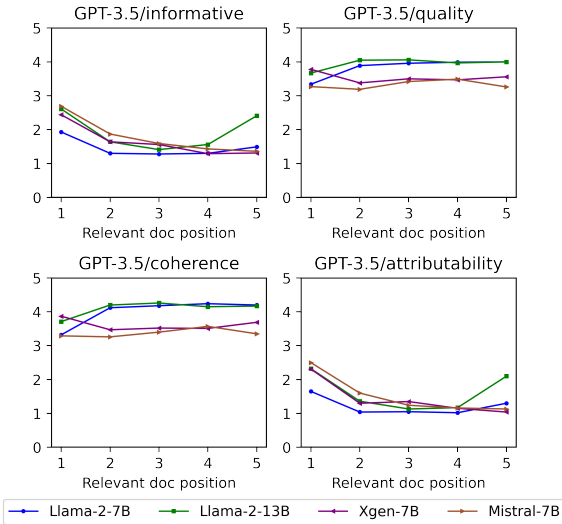


Figure 4: Fine-grained evaluation of multi-document summarization on Multi-News with GPT-3.5 when varying the position of a unique relevant input document.

Dataset	Model	Input documents	R-2	BS	A3CU	SummaC	GPT-3.5
Multi-X	Llama-2-7B	All 7	4.64	82.83	5.88	54.56	4.17
		First + last	4.62	82.82	5.78	47.00	4.50
		First + 5 random + last	4.43	82.64	5.37	43.25	4.01
		Random	1.22	81.14	3.08	30.30	1.14
		All 7	4.78	83.00	6.64	42.62	4.35
	Llama-2-13B	First + last	4.73	82.86	5.76	43.74	4.53
		First + 5 random + last	4.61	82.80	5.72	46.42	4.31
		Random	1.16	81.16	3.22	28.38	1.42
	Xgen-7B	All 7	5.37	82.68	6.59	44.34	4.19
		First + last	5.01	82.73	5.86	49.08	4.45
		First + 5 random + last	3.89	82.16	5.03	55.29	3.01
	Mistral-7B	Random	1.25	80.74	3.21	27.77	1.29
All 7		5.40	82.60	6.35	63.78	4.26	
First + last		5.12	82.67	6.15	60.91	4.76	
	First + 5 random + last	4.45	82.40	5.35	58.19	4.06	
	Random	1.25	80.64	3.13	25.97	1.18	
	Multi-N	Llama-2-7B	All 5	10.76	85.04	19.06	60.09
First + last			9.50	84.43	15.88	54.52	3.80
First + 3 random + last			7.57	83.36	12.39	50.35	2.94
Random			1.52	80.22	3.11	33.34	1.92
Llama-2-13B		All 5	10.42	84.60	18.15	57.26	3.83
		First + last	9.55	84.58	16.99	49.84	3.73
		First + 3 random + last	8.27	83.79	14.84	50.09	3.18
Xgen-7B		Random	1.55	80.05	3.17	32.80	1.98
		All 5	9.04	83.18	17.05	60.55	3.32
		First + last	7.82	83.27	14.18	51.59	3.60
Mistral-7B		First + 3 random + last	6.30	81.85	11.66	49.02	2.66
		Random	1.52	78.99	3.48	37.62	1.86
	All 5	9.52	83.55	17.03	63.02	3.15	
	First + last	9.11	83.69	14.99	67.14	3.51	
	First + 3 random + last	6.59	81.66	12.50	52.14	2.45	
	Random	1.57	79.21	3.30	37.18	1.80	

Table 3: Performance in multi-document summarization on Multi-XScience (7 documents) and Multi-News (5 documents) when infilling the middle of the context window with random documents. **R-2** is ROUGE-2, **BS** refers to BERTScore.

documents. We also run a baseline just using the first and last documents as input, expected to be close to the result with random documents in between. As displayed in Table 3, filling with random noise between the first and last document (which amounts to a prompt mostly irrelevant to the reference) leads to a moderate drop in performance. For instance, on Multi-XScience, with 5 random documents between the first and last, Llama-2-13B maintains 4.31 as GPT-3.5 score compared to 4.35 when using all 7 documents.

We conclude from these two experiments that *LLMs can focus on the beginning and/or the end of their input, but largely ignore the middle. The U-shape or middle curse from Liu et al. (2023a) also applies to abstractive summarization.*

4.2 Can we alleviate the middle curse?

To evaluate loss of performance due to the *middle curse* in a natural setup, we subsample data points from each of the 5 long-input summarization datasets. We obtain sentences from the (un-truncated) source aligned with the reference summary, following the procedure from §3.2. Only data points where the start index of the earliest aligned source sentence is at least 1,200 words, are kept, ensuring no salient information at the start. We randomly sample 50 data points from each of Arxiv, PubMed, GovReport and Multi-News, and 25 from SummScreenFD, forming an evaluation dataset of




Figure 5: Reference-based evaluation on the MiddleSum dataset. We also report (gray bars) performance achieved by uniformly sampling subsets of the same size as MiddleSum from the original datasets, alongside bootstrapping variance (black lines).

225 samples which we name **MiddleSum**.²

We evaluate LLMs on MiddleSum, keeping only reference-based evaluation as the dataset is built using saliency with regard to the reference. As expected, in Fig. 5 we see that LLMs perform noticeably worse on MiddleSum (green bars) as compared to the full set (gray bars), confirming that MiddleSum is a more challenging task.

We benchmark alternative inference methods on MiddleSum: *hierarchical* summarization and *incremental* summarization, both of which are explored in the concurrent work of Chang et al. (2023). Namely, let us divide an input x of length n into k consecutive blocks of size at most m (yielding $k = \lceil \frac{n}{m} \rceil$): $x = (x_1, \dots, x_k)$.

Hierarchical summarization consists in summarizing each block and then summarizing the concatenation of summaries:

$$y_i = \text{LLM}(x_i) \quad \forall i \in \{1, \dots, k\} \quad (1)$$

$$y = \text{LLM}(y_1, \dots, y_k) \quad (2)$$

Incremental summarization consists in updating a summary of the text so far with content from the current text block (we have $y_0 = \emptyset$):

$$\forall i \in \{1, \dots, k\}, y_i = \text{LLM}(y_{i-1}, x_i) \quad (3)$$

In both methods, the final output is y . Noting l the output length, standard inference has complexity in $\mathcal{O}(l.n^2)$, while both alternative methods have complexity in $\mathcal{O}(k.l.m^2) = \mathcal{O}(l.n.m)$, which is lower. For both methods and all models, we use a block size m of 1,500 words (roughly 2,000 tokens), and preserve coherence by ending blocks at the earliest end of sentence reaching the length.

Results are shown in Fig. 5 (blue and purple bars), with statistics, bigrams analysis and results split over datasets in Appendix E. Both methods show promising results, notably on Mistral-7B for which they improve performance significantly.

²We will release it alongside our code

Across domains, *hierarchical* and *incremental* inference are very successful on scientific publications, which we hypothesize is due to the natural division in structured sections of such inputs. Yet, they seem to harm summaries on the other domains.

4.3 Is scaling context length really useful?

Experiments from §4.1 confirm that LLMs struggle to summarize information contained in the middle of their context window. This poses issues for long-input summarization: after the initial part with high saliency, important information becomes sparser, and at the same time LLMs processing capability weakens. To investigate this issue, we infer long-document summarization with length truncated at $m * 2k$ tokens, varying m from 1 to 6. We use our longest context LLMs Xgen-7B and Mistral-7B, as well as two LLMs extending Llama-2-7B context window with **position interpolation** (Chen et al., 2023), a method gaining traction as an efficient way to scale LLMs' context window. We use two LLMs applying position interpolation: Vicuna-7B-1.5-16k³, and Llama-2-7B-32k⁴, with context of 16k tokens and 32k tokens, respectively.

Results for GovReport in Fig. 6 (and Appendix F for Arxiv and SummScreen) confirm our intuition: all metrics plateau or even *decrease* (see Mistral-7B) from 4k context window upwards. Two conflicting forces are at play when increasing length: giving more information to the model helps it retrieving key elements to make a summary, while reasoning over a longer context is more challenging. Yet, such a drop for Xgen-7B and Mistral-7B at 8k inference length is concerning. Both position interpolated models show more robustness ; thus these patterns may vary a lot across LLMs. Our results suggest that *in the current LLMs inference and evaluation framework, there is no need to exceed 4k tokens in the context window.*

³In HuggingFace: `lmsys/vicuna-7b-v1.5-16k`

⁴In HuggingFace: `togethercomputer/LLaMA-2-7B-32K`




Figure 6: Long-input summarization performance on GovReport with 4 models and all 5 metrics. X-axis represents the truncated maximum source length. Xgen-7b and Mistral-7b cannot infer beyond 8k tokens.

5 Related Work

Summarization with LLMs It is widely acknowledged that LLMs have propelled forward abstractive summarization research (Pu et al., 2023), with their summaries being highly rated by human annotators (Goyal et al., 2022; Zhang et al., 2023b). Liu et al. (2023b) proposes to train smaller models like BART (Lewis et al., 2020) or BRIO (Liu et al., 2022b) with contrastive learning using LLMs like ChatGPT as evaluator providing signal on which generated summary candidate is better. Summary chain-of-thought designs a custom chain-of-thought method which first prompts the LLM to list important facts, then integrates these facts into a coherent summary (Wang et al., 2023c). SummIt utilizes ChatGPT to iteratively write then refine summaries given feedback from an evaluator LLM (Zhang et al., 2023a). Chain-of-density gradually makes GPT-4 generated summaries contain more and more entities while keeping length budget constant, creating more informative albeit a bit less readable summaries (Adams et al., 2023). Ravaut et al. (2022) noticed that data points with higher compression are generally harder to summarize with pre-trained models.

Position bias in LLMs Sun et al. (2021) showed that for Transformer-based models, most recent tokens play a greater role compared to older tokens for next-token prediction. It was later found that for in-context learning, the order of examples within the prompt impacts GPT-3’s performance (Liu et al., 2021; Lu et al., 2021). Reliance on positional information affects LLMs capabilities in arithmetic (Shen et al., 2023b), in multiple-choice question-answering (Zheng et al., 2023; Pezeshkpour and Hruschka, 2023), and as text generation evaluators (Wang et al., 2023b); making it hard to rank LLMs (Alzahrani et al., 2024). Liu et al. (2023a) was the first to show that LLMs’ performance weakens in the middle of the prompt, yet, how LLMs make use of their full context window

remains poorly understood. The *passkey retrieval* evaluation, which consists in prompting the LLM to recall a complex string or long number inserted in its prompt, is becoming popular recently as a way of verifying LLM’s processing capability at each position (Liu et al., 2023a; Jiang et al., 2024). However, this task does not measure position bias on complex, abstract reasoning tasks like summarization. A line of work attempts to solve the *middle curse* through compressing the prompt (Jiang et al., 2023b,c), with very promising results albeit at the cost of prompt fluency. Another approach marginalizes results over different permutations of the input to suppress dependency on input order (Tang et al., 2023). Concurrent work to ours also finds that in zero-shot summarization, LLMs tend to prefer lead content (Chhabra et al., 2024).

6 Conclusion

Behind the recent hype around LLMs and their amazing instruction following and content generation capacities, our study showcases a major weakness in abstractive summarization: LLMs suffer from the *middle curse* and struggle to use information in the middle of their context window. LLMs do not make a consistent use of their context window as they mostly look at the beginning and (to a lesser extent) the end, which at first glance may be hidden by the prevalent *lead bias* in summarization datasets. Extending context window beyond 4k tokens, which has been an intense area of focus lately, is not justified in the current inference and evaluation setup in abstractive summarization. We benchmarked two alternative inference methods (*hierarchical* inference and *incremental* inference) on MiddleSum, an evaluation subset designed to showcase the *middle curse*. Despite promising results, especially on scientific paper datasets, these methods are far from a silver bullet to the *middle curse*. We call for a better evaluation of LLMs, which accounts for the salient spans of the source which are effectively being processed.

562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583

584

585
586
587
588

589
590
591
592
593
594
595

596
597
598
599
600
601

602
603
604
605

606
607
608
609
610
611
612

Limitations

Our work only considers open-source LLMs for summary generation and ignores closed-source LLMs such as the popular OpenAI’s GPT-3.5 and GPT-4, or Anthropic’s Claude. We made this decision to advocate for openness in LLM research ; yet we acknowledge that it would be interesting to also investigate properties of summaries generated by these paying API LLMs.

Another limitation lays in the saliency estimation. We approximate salient content in the source through maximizing ROUGE-1 overlap with summary sentences. Other methods are also well-suited for this task, albeit at greater computational cost ; for instance semantic similarity through BERTScore or BARTScore ; or saliency estimation through a LLM in zero-shot.

Lastly, we can only evaluate a finite number of LLMs, and we settled for the evaluation of 5 recent and popular open-source LLMs. Findings may change as LLMs undergo changes and improvements in their training process.

References

Griffin Adams, Alexander Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: Gpt-4 summarization with chain of density prompting. *arXiv preprint arXiv:2309.04269*.

Norah Alzahrani, Hisham Abdullah Alyahya, Yazeed Alnumay, Sultan Alrashed, Shaykhah Alsubaie, Yusef Almushaykeh, Faisal Mirza, Nouf Alotaibi, Nora Altwaresh, Areeb Alowisheq, et al. 2024. When benchmarks are targets: Revealing the sensitivity of large language model leaderboards. *arXiv preprint arXiv:2402.01781*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Yapei Chang, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2023. Boookscore: A systematic exploration of book-length summarization in the era of llms. *arXiv preprint arXiv:2310.00785*.

Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2022. SummScreen: A dataset for abstractive screenplay summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8602–8615, Dublin, Ireland. Association for Computational Linguistics.

Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.

Anshuman Chhabra, Hadi Askari, and Prasant Mohapatra. 2024. Revisiting zero-shot abstractive summarization in the era of large language models from the perspective of position bias. *arXiv preprint arXiv:2401.01989*.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric 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.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

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.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.

Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.

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.

Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.

667	Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. <i>Advances in neural information processing systems</i> , 28.	723
668		724
669		725
670		726
671		727
672	Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji, and Lu Wang. 2021. Efficient attentions for long document summarization. In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 1419–1436, Online. Association for Computational Linguistics.	728
673		729
674		730
675		731
676		732
677		733
678		734
679	Sameer Jain, Vaishakh Keshava, Swarnashree Mysore Sathyendra, Patrick Fernandes, Pengfei Liu, Graham Neubig, and Chunting Zhou. 2023. Multi-dimensional evaluation of text summarization with in-context learning. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 8487–8495, Toronto, Canada. Association for Computational Linguistics.	735
680		736
681		737
682		738
683		739
684		740
685		741
686		742
687	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. 2023a. Mistral 7b. <i>arXiv preprint arXiv:2310.06825</i> .	743
688		744
689		745
690		746
691		747
692	Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. <i>arXiv preprint arXiv:2401.04088</i> .	748
693		749
694		750
695		751
696		752
697	Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023b. LlmLingua: Compressing prompts for accelerated inference of large language models. <i>arXiv preprint arXiv:2310.05736</i> .	753
698		754
699		755
700		756
701	Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023c. Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression. <i>arXiv preprint arXiv:2310.06839</i> .	757
702		758
703		759
704		760
705		761
706	Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. Abstractive summarization of Reddit posts with multi-level memory networks. In <i>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)</i> , pages 2519–2531, Minneapolis, Minnesota. Association for Computational Linguistics.	762
707		763
708		764
709		765
710		766
711		767
712		768
713		769
714	Philippe Laban, Tobias Schnabel, Paul N Bennett, and Marti A Hearst. 2022. Summac: Re-visiting nli-based models for inconsistency detection in summarization. <i>Transactions of the Association for Computational Linguistics</i> , 10:163–177.	770
715		771
716		772
717		773
718		774
719	Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training	775
720		776
721		777
722		778
	for natural language generation, translation, and comprehension. In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 7871–7880, Online. Association for Computational Linguistics.	779
	Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In <i>Text Summarization Branches Out</i> , pages 74–81, Barcelona, Spain. Association for Computational Linguistics.	780
	Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? <i>arXiv preprint arXiv:2101.06804</i> .	781
	Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023a. Lost in the middle: How language models use long contexts. <i>arXiv preprint arXiv:2307.03172</i> .	782
	Yixin Liu, Alexander R Fabbri, Pengfei Liu, Dragomir Radev, and Arman Cohan. 2023b. On learning to summarize with large language models as references. <i>arXiv preprint arXiv:2305.14239</i> .	783
	Yixin Liu, Alexander R Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, et al. 2022a. Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation. <i>arXiv preprint arXiv:2212.07981</i> .	784
	Yixin Liu, Alexander R Fabbri, Yilun Zhao, Pengfei Liu, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2023c. Towards interpretable and efficient automatic reference-based summarization evaluation. <i>arXiv preprint arXiv:2303.03608</i> .	785
	Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022b. BRIO: Bringing order to abstractive summarization. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.	786
	Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. <i>arXiv preprint arXiv:2104.08786</i> .	787
	Yao Lu, Yue Dong, and Laurent Charlin. 2020. MultiXScience: A large-scale dataset for extreme multi-document summarization of scientific articles. In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 8068–8074, Online. Association for Computational Linguistics.	788
	Zheheng Luo, Qianqian Xie, and Sophia Ananiadou. 2023. Chatgpt as a factual inconsistency evaluator for text summarization.	789

778	Frank J Massey Jr. 1951. The kolmogorov-smirnov	831
779	test for goodness of fit. <i>Journal of the American</i>	832
780	<i>statistical Association</i> , 46(253):68–78.	833
781	Shashi Narayan, Shay B. Cohen, and Mirella Lapata.	834
782	2018. Don’t give me the details, just the summary!	835
783	topic-aware convolutional neural networks for ex-	836
784	treme summarization. In <i>Proceedings of the 2018</i>	
785	<i>Conference on Empirical Methods in Natural Lan-</i>	
786	<i>guage Processing</i> , pages 1797–1807, Brussels, Bel-	
787	gium. Association for Computational Linguistics.	
788	Erik Nijkamp, Tian Xie, Hiroaki Hayashi, Bo Pang,	
789	Congying Xia, Chen Xing, Jesse Vig, Semih Yavuz,	
790	Philippe Laban, Ben Krause, et al. 2023. Xgen-7b	
791	technical report. <i>arXiv preprint arXiv:2309.03450</i> .	
792	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	
793	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	
794	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	
795	2022. Training language models to follow instruc-	
796	tions with human feedback. <i>Advances in Neural</i>	
797	<i>Information Processing Systems</i> , 35:27730–27744.	
798	Bo Pang, Erik Nijkamp, Wojciech Kryściński, Sil-	
799	vio Savarese, Yingbo Zhou, and Caiming Xiong.	
800	2022. Long document summarization with top-	
801	down and bottom-up inference. <i>arXiv preprint</i>	
802	<i>arXiv:2203.07586</i> .	
803	Guilherme Penedo, Quentin Malartic, Daniel Hesslow,	
804	Ruxandra Cojocaru, Alessandro Cappelli, Hamza	
805	Alobeidli, Baptiste Pannier, Ebtesam Almazrouei,	
806	and Julien Launay. 2023. The refinedweb dataset	
807	for falcon llm: outperforming curated corpora with	
808	web data, and web data only. <i>arXiv preprint</i>	
809	<i>arXiv:2306.01116</i> .	
810	Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and En-	
811	rico Shippole. 2023. Yarn: Efficient context window	
812	extension of large language models. <i>arXiv preprint</i>	
813	<i>arXiv:2309.00071</i> .	
814	Pouya Pezeshkpour and Estevam Hruschka. 2023.	
815	Large language models sensitivity to the order of	
816	options in multiple-choice questions. <i>arXiv preprint</i>	
817	<i>arXiv:2308.11483</i> .	
818	Ofir Press, Noah A Smith, and Mike Lewis. 2021.	
819	Train short, test long: Attention with linear biases	
820	enables input length extrapolation. <i>arXiv preprint</i>	
821	<i>arXiv:2108.12409</i> .	
822	Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023.	
823	Summarization is (almost) dead. <i>arXiv preprint</i>	
824	<i>arXiv:2309.09558</i> .	
825	Mathieu Ravaut, Shafiq Joty, and Nancy Chen. 2022.	
826	Towards summary candidates fusion. In <i>Proceed-</i>	
827	<i>ings of the 2022 Conference on Empirical Methods</i>	
828	<i>in Natural Language Processing</i> , pages 8488–8504,	
829	Abu Dhabi, United Arab Emirates. Association for	
830	Computational Linguistics.	
	Teven Le Scao, Angela Fan, Christopher Akiki, El-	837
	lie Pavlick, Suzana Ilić, Daniel Hesslow, Roman	838
	Castagné, Alexandra Sasha Luccioni, François Yvon,	839
	Matthias Gallé, et al. 2022. Bloom: A 176b-	840
	parameter open-access multilingual language model.	841
	<i>arXiv preprint arXiv:2211.05100</i> .	842
	Abigail See, Peter J. Liu, and Christopher D. Manning.	843
	2017. Get to the point: Summarization with pointer-	
	generator networks. In <i>Proceedings of the 55th An-</i>	
	<i>nuual Meeting of the Association for Computational</i>	
	<i>Linguistics (Volume 1: Long Papers)</i> , pages 1073–	
	1083, Vancouver, Canada. Association for Computa-	
	tional Linguistics.	844
	Uri Shaham, Maor Ivgi, Avia Efrat, Jonathan Ber-	845
	ant, and Omer Levy. 2023. Zeroscrolls: A zero-	846
	shot benchmark for long text understanding. <i>arXiv</i>	847
	<i>preprint arXiv:2305.14196</i> .	
	Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori	848
	Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong,	849
	Mor Geva, Jonathan Berant, and Omer Levy. 2022.	850
	SCROLLS: Standardized CompaRison over long lan-	851
	guage sequences. In <i>Proceedings of the 2022 Con-</i>	852
	<i>ference on Empirical Methods in Natural Language</i>	853
	<i>Processing</i> , pages 12007–12021, Abu Dhabi, United	854
	Arab Emirates. Association for Computational Lin-	855
	guistics.	856
	Chenhui Shen, Liying Cheng, Yang You, and Lidong	857
	Bing. 2023a. Are large language models good evalu-	858
	ators for abstractive summarization? <i>arXiv preprint</i>	859
	<i>arXiv:2305.13091</i> .	860
	Ruoqi Shen, Sébastien Bubeck, Ronen Eldan, Yin Tat	861
	Lee, Yuanzhi Li, and Yi Zhang. 2023b. Positional de-	862
	scription matters for transformers arithmetic. <i>arXiv</i>	863
	<i>preprint arXiv:2311.14737</i> .	864
	Simeng Sun, Kalpesh Krishna, Andrew Mattarella-	865
	Micke, and Mohit Iyyer. 2021. Do long-range lan-	866
	guage models actually use long-range context? In	867
	<i>Proceedings of the 2021 Conference on Empirical</i>	868
	<i>Methods in Natural Language Processing</i> , pages 807–	869
	822, Online and Punta Cana, Dominican Republic.	870
	Association for Computational Linguistics.	871
	Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shao-	872
	han Huang, Alon Benhaim, Vishrav Chaudhary, Xia	873
	Song, and Furu Wei. 2022. A length-extrapolatable	874
	transformer. <i>arXiv preprint arXiv:2212.10554</i> .	875
	Raphael Tang, Xinyu Zhang, Xueguang Ma, Jimmy	876
	Lin, and Ferhan Ture. 2023. Found in the mid-	877
	dle: Permutation self-consistency improves listwise	878
	ranking in large language models. <i>arXiv preprint</i>	879
	<i>arXiv:2310.07712</i> .	880
	Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Gar-	881
	cia, Jason Wei, Xuezhi Wang, Hyung Won Chung,	882
	Dara Bahri, Tal Schuster, Steven Zheng, et al. 2022.	883
	U12: Unifying language learning paradigms. In <i>The</i>	884
	<i>Eleventh International Conference on Learning Rep-</i>	885
	<i>resentations</i> .	886

887	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier	9732–9747, Toronto, Canada. Association for Com-	944
888	Martinet, Marie-Anne Lachaux, Timothée Lacroix,	putational Linguistics.	945
889	Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal		
890	Azhar, et al. 2023a. Llama: Open and effi-	Wenhan Xiong, Anchit Gupta, Shubham Toshniwal,	946
891	cient foundation language models. <i>arXiv preprint</i>	Yashar Mehdad, and Wen-tau Yih. 2022. Adapting	947
892	<i>arXiv:2302.13971</i> .	pretrained text-to-text models for long text sequences.	948
		<i>arXiv preprint arXiv:2209.10052</i> .	949
893	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-		
894	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023a.	950
895	Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti	Summit: Iterative text summarization via chatgpt.	951
896	Bhosale, et al. 2023b. Llama 2: Open founda-	<i>arXiv preprint arXiv:2305.14835</i> .	952
897	tion and fine-tuned chat models. <i>arXiv preprint</i>		
898	<i>arXiv:2307.09288</i> .	Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q	953
		Weinberger, and Yoav Artzi. 2019. Bertscore: Eval-	954
899	Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang	uating text generation with bert. <i>arXiv preprint</i>	955
900	Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou.	<i>arXiv:1904.09675</i> .	956
901	2023a. Is chatgpt a good nlg evaluator? a preliminary		
902	study. <i>arXiv preprint arXiv:2303.04048</i> .	Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy	957
		Liang, Kathleen McKeown, and Tatsunori B	958
903	Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai	Hashimoto. 2023b. Benchmarking large language	959
904	Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui.	models for news summarization. <i>arXiv preprint</i>	960
905	2023b. Large language models are not fair evaluators.	<i>arXiv:2301.13848</i> .	961
906	<i>arXiv preprint arXiv:2305.17926</i> .		
		Chao Zhao, Faeze Brahman, Kaiqiang Song, Wenlin	962
907	Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023c.	Yao, Dian Yu, and Snigdha Chaturvedi. 2022a. Nar-	963
908	Element-aware summarization with large language	raSum: A large-scale dataset for abstractive narra-	964
909	models: Expert-aligned evaluation and chain-of-	tive summarization. In <i>Findings of the Association</i>	965
910	thought method. In <i>Proceedings of the 61st Annual</i>	<i>for Computational Linguistics: EMNLP 2022</i> , pages	966
911	<i>Meeting of the Association for Computational Lin-</i>	182–197, Abu Dhabi, United Arab Emirates. Associ-	967
912	<i>guistics (Volume 1: Long Papers)</i> , pages 8640–8665,	ation for Computational Linguistics.	968
913	Toronto, Canada. Association for Computational Lin-		
914	guistics.	Yao Zhao, Misha Khalman, Rishabh Joshi, Shashi	969
		Narayan, Mohammad Saleh, and Peter J Liu.	970
915	Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin	2022b. Calibrating sequence likelihood improves	971
916	Guu, Adams Wei Yu, Brian Lester, Nan Du, An-	conditional language generation. <i>arXiv preprint</i>	972
917	drew M Dai, and Quoc V Le. 2021. Finetuned lan-	<i>arXiv:2210.00045</i> .	973
918	guage models are zero-shot learners. <i>arXiv preprint</i>		
919	<i>arXiv:2109.01652</i> .	Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou,	974
		and Minlie Huang. 2023. Large language models are	975
920	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien	not robust multiple choice selectors. <i>arXiv e-prints</i> ,	976
921	Chaumond, Clement Delangue, Anthony Moi, Pier-	pages arXiv–2309.	977
922	ric Cistac, Tim Rault, Remi Louf, Morgan Funtow-		
923	icz, Joe Davison, Sam Shleifer, Patrick von Platen,	Qingyu Zhou, Nan Yang, Furu Wei, Shaohan Huang,	978
924	Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu,	Ming Zhou, and Tiejun Zhao. 2018. Neural docu-	979
925	Teven Le Scao, Sylvain Gugger, Mariama Drame,	ment summarization by jointly learning to score and	980
926	Quentin Lhoest, and Alexander Rush. 2020. Trans-	select sentences. In <i>Proceedings of the 56th Annual</i>	981
927	formers: State-of-the-art natural language processing.	<i>Meeting of the Association for Computational Lin-</i>	982
928	In <i>Proceedings of the 2020 Conference on Empirical</i>	<i>guistics (Volume 1: Long Papers)</i> , pages 654–663,	983
929	<i>Methods in Natural Language Processing: System</i>	Melbourne, Australia. Association for Computational	984
930	<i>Demonstrations</i> , pages 38–45, Online. Association	Linguistics.	985
931	for Computational Linguistics.		
		A Statistics	986
932	Wen Xiao, Iz Beltagy, Giuseppe Carenini, and Arman		
933	Cohan. 2022. PRIMERA: Pyramid-based masked	In Table 4, we include statistics on each of the	987
934	sentence pre-training for multi-document summariza-	abstractive summarization datasets under consid-	988
935	tion. In <i>Proceedings of the 60th Annual Meeting of</i>	eration. We use the non-anonymized version for	989
936	<i>the Association for Computational Linguistics (Vol-</i>	CNN/DM (See et al., 2017). For Reddit-TIFU,	990
937	<i>ume 1: Long Papers)</i> , pages 5245–5263, Dublin,	we use the Long subset, and SummScreenFD,	991
938	Ireland. Association for Computational Linguistics.	is the ForeverDreaming (FD) subset of Summ-	992
		Screen. GovReport and SummScreenFD are part of	993
939	Jiawen Xie, Qi Su, Shaoting Zhang, and Xiaofan Zhang.	the long-input benchmarks Scrolls (Shaham et al.,	994
940	2023. Alleviating exposure bias via multi-level con-	2022) and ZeroScrolls (Shaham et al., 2023).	995
941	trastive learning and deviation simulation in abstrac-		
942	tive summarization. In <i>Findings of the Association</i>		
943	<i>for Computational Linguistics: ACL 2023</i> , pages		

Dataset	Domain	# Docs	# Data points			# Sentences			# Words		# Tokens		
			Train	Val	Test	Doc.	Summ.	Inf.	Doc.	Summ.	Doc.	Summ.	Max gen.
CNN/DM (Hermann et al., 2015)	News	1.00	287,113	13,334	<u>11,490</u>	33.37	3.79	3	773.23	57.75	994.56	84.47	192
XSum (Narayan et al., 2018)	News	1.00	204,045	11,332	<u>11,334</u>	19.11	1.00	1	433.05	23.19	566.79	31.63	64
Reddit-TIFU (Long (Kim et al., 2019)	Social Media	1.00	33,701	4,214	<u>4,221</u>	22.21	1.45	2	444.20	23.37	532.18	29.82	128
SAMSum (Kim et al., 2019)	Dialogue	1.00	14,732	818	819	8.96	2.01	2	126.93	23.12	175.54	29.69	128
Multi-XScience (Lu et al., 2020)	Science	5.06	30,369	5,066	<u>5,093</u>	30.55	4.86	5	773.36	120.65	965.99	157.77	384
Arxiv (Cohan et al., 2018)	Science	1.00	203,037	6,436	<u>6,440</u>	250.37	6.23	6	6,446.11	166.72	8,940.00	225.58	512
PubMed (Cohan et al., 2018)	Science (medical)	1.00	119,924	6,633	<u>6,658</u>	101.61	7.59	7	3,142.92	208.03	4,602.62	324.97	512
GovReport (Huang et al., 2021)	Legal	1.00	17,517	973	973	282.86	23.14	22	8,363.22	649.01	11,025.02	879.10	768
SummScreenFD (Chen et al., 2022)	TV Transcripts	1.00	3,673	338	337	727.06	5.26	5	7,618.20	123.34	10,067.39	157.44	512
Multi-News (Fabbri et al., 2019)	News	2.73	44,972	5,622	<u>5,622</u>	79.02	9.88	10	2,101.49	256.55	2,998.52	324.29	512

Table 4: Statistics on the 10 datasets used for experiments. **Doc.** is the source document, **Summ.** the ground-truth summary, **Inf.** refers to the number of desired sentences to be in the summary prompted to each LLM during inference. Statistics are computed on the entire test set. **# Tokens** is calculated with Llama-2’s tokenizer. **Max gen.** is the maximum tokens size that we set when decoding summaries. Underlined test sizes correspond to datasets where we subsample randomly 1,000 test data points for evaluation.

In Fig. 7, we illustrate how much of the source document(s) is visible with a 4k context window (Llama-2).

B GPT-3.5 Evaluation

To evaluate LLM-generated summaries with GPT-3.5, we use the following prompt template:

Score the following summary generated by another system given the source on a scale from 1 to 5 with regards to overall general summary quality. 1-point indicates a low quality summary, and 5 points a very high quality summary. A high quality summary is grammatical, fluent, informative, relevant, coherent and factually consistent with the source. Let’s think step-by-step and just output the score.

Source:
[source]

Instruction:
Summarize the above text in [n] sentences.

Summary:
[summary generated by the LLM to evaluate]

Your score:

When evaluating for the specific aspect of **informativeness**, the first paragraph becomes:

Score the following summary generated by another system given the source on a scale from 1 to 5 with regards to how informative the summary is. 1 point indicates a not informative summary, and 5 points a very informative summary. An informative summary captures the important information in

the article and presents it accurately and concisely. Let’s think step-by-step and just output the score.

When evaluating for the specific aspect of **overall quality**, the first paragraph becomes:

Score the following summary generated by another system given the source on a scale from 1 to 5 with regards to its quality. 1 point indicates a low quality summary, and 5 points a very high quality summary. A high quality summary is comprehensible and understandable. Let’s think step-by-step and just output the score.

When evaluating for the specific aspect of **coherence**, the first paragraph becomes:

Score the following summary generated by another system given the source on a scale from 1 to 5 with regards to its coherence. 1 point indicates an incoherent summary, and 5 points a very coherent summary. A coherent summary is well-structured and well-organized. Let’s think step-by-step and just output the score.

When evaluating for the specific aspect of **attributability**, the first paragraph becomes:

Score the following summary generated by another system given the source on a scale from 1 to 5 with regards to how attributable it is. 1 point indicates a not very attributable summary, with many hallucinations, and 5 points a summary very attributable to the source, consistent with the source. In a very attributable summary, all the information is fully attributable to the source. Let’s think step-by-step and just output the score.

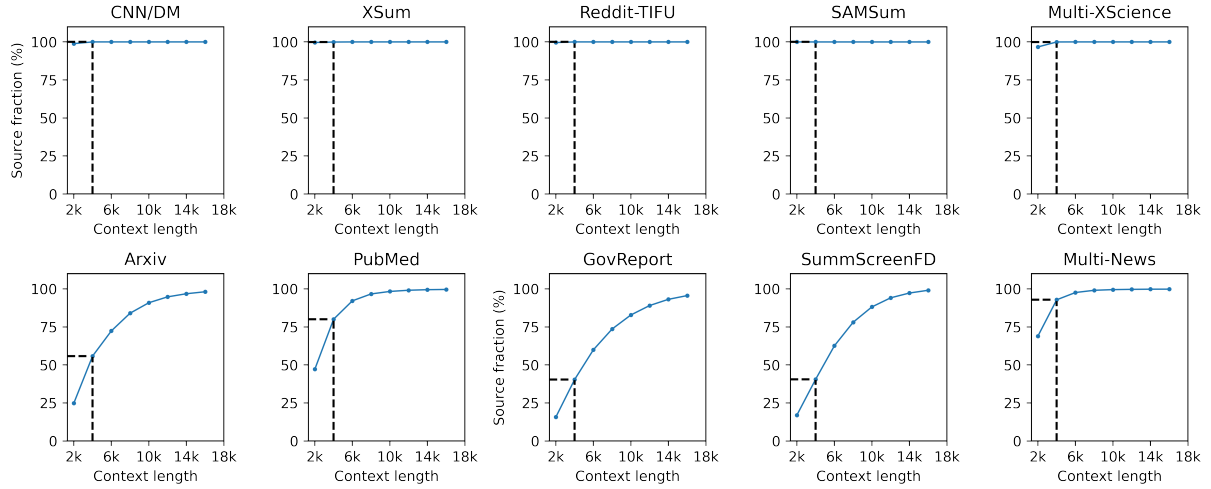


Figure 7: Fraction of the source which fits into the context window, for several context window lengths with Llama-2 tokenization. The black dashed lines correspond to Llama-2 context window length of 4k tokens. On standard-length datasets, 4k is enough to access 100% of all source documents ; but on the long-input datasets such as GovReport or SummScreenFD, such a context window may not even fit 50% of the source.

Model	Metric	CNN/DM	XSum	Reddit-TIFU	SAMSum	Multi-X	AVG	Arxiv	PubMed	GovReport	SummScreenFD	Multi-N	AVG
SOTA	ROUGE-2	24.17	27.09	11.13	29.88	4.60	19.37	21.93	23.26	30.90	10.70	13.60	20.07
	# sents	2.89	1.00	1.34	2.08	2.28	1.19	1.78	1.27	1.80	5.49	2.79	2.63
Flan-UL2	ROUGE-2	20.28	22.74	8.61	28.21	3.04	16.58	9.37	7.42	4.76	4.33	7.79	6.73
	BERTScore	88.05	91.94	87.42	92.29	81.87	88.31	83.82	83.07	83.53	84.85	84.97	84.05
	A3CU	32.69	32.11	16.89	49.48	5.98	27.43	14.79	13.83	12.00	8.37	16.99	13.20
	SummaC	69.96	24.27	35.76	30.19	57.98	43.63	67.56	60.96	73.80	56.00	76.80	67.02
	GPT-3.5	3.16	3.52	1.61	2.92	3.23	2.89	2.66	2.85	3.13	2.35	3.32	2.86
Llama-2-7B	# sents	3.00	1.27	2.00	1.83	7.77	3.17	5.80	6.61	12.88	5.77	18.69	9.95
	ROUGE-2	14.16	7.27	4.17	15.53	4.87	9.20	13.84	12.89	16.22	5.36	12.37	12.14
	BERTScore	87.32	87.47	85.87	89.95	83.32	86.79	83.84	82.82	85.28	85.41	85.63	84.60
	A3CU	29.04	14.18	12.15	35.64	6.39	19.48	16.78	16.60	17.23	9.66	22.23	16.50
	SummaC	35.58	25.24	26.38	24.56	49.39	32.23	53.22	51.82	70.47	39.01	57.49	54.40
Llama-2-13B	GPT-3.5	4.10	4.24	2.83	3.61	4.42	3.84	4.21	4.19	3.43	2.71	3.91	3.69
	# sents	3.01	1.16	2.00	1.98	5.22	2.67	5.92	7.22	27.75	5.16	12.79	11.77
	ROUGE-2	14.10	8.61	4.22	14.16	5.29	9.28	13.52	15.24	17.28	5.62	12.58	12.85
	BERTScore	87.40	87.94	85.85	89.45	83.58	86.84	83.88	84.24	85.29	85.42	85.84	84.93
	A3CU	29.57	16.30	12.94	34.19	7.29	20.06	16.44	19.21	17.01	10.30	23.09	17.21
Xgen-7B	SummaC	33.83	24.07	25.76	24.81	41.70	30.03	55.09	56.00	76.44	38.74	53.12	55.88
	GPT-3.5	4.12	4.34	2.91	3.45	4.45	3.85	4.06	4.13	3.66	2.69	3.89	3.69
	# sents	3.93	2.24	2.62	2.34	5.94	3.41	8.07	13.50	22.46	10.60	9.05	12.74
	ROUGE-2	14.55	6.00	3.98	14.51	5.54	8.92	12.31	13.99	14.68	4.24	11.07	11.26
	BERTScore	87.07	87.12	85.84	89.53	83.42	86.60	83.07	82.87	83.94	83.91	84.95	83.75
Mistral-7B	A3CU	27.88	12.75	12.56	33.18	7.45	18.76	15.28	18.79	15.65	8.44	21.33	18.90
	SummaC	52.25	37.95	28.40	25.63	44.36	37.72	53.28	60.56	65.22	42.03	56.45	55.51
	GPT-3.5	3.82	3.97	2.78	3.52	4.37	3.69	3.96	3.99	2.78	2.42	3.58	3.35
	# sents	3.10	1.12	2.64	2.25	7.73	3.37	12.00	11.88	25.83	27.13	16.43	18.65
	ROUGE-2	16.37	7.05	4.34	14.66	5.57	9.60	9.77	14.32	11.36	3.11	12.61	10.23
Mistral-7B	BERTScore	87.50	87.45	85.71	89.78	83.16	86.72	81.44	82.85	82.43	81.46	85.07	82.65
	A3CU	30.60	13.14	13.08	32.21	6.96	19.20	12.66	16.31	14.92	8.40	22.22	14.90
	SummaC	53.67	24.79	30.51	26.82	62.49	39.66	57.81	69.03	67.07	35.76	68.50	59.63
	GPT-3.5	3.92	4.30	2.73	3.63	4.47	3.81	2.83	3.63	2.01	1.88	3.60	2.79

Table 5: Performance achieved by the 5 LLMs on each dataset for 5 automatic metrics. # sents is the average number of sentences in generated summaries. **Multi-X** is short for Multi-XScience, **Multi-N** is Multi-News dataset, **AVG** columns represent the average over standard-length and long-input datasets, respectively. SOTA numbers are taken from (Xie et al., 2023) on CNN/DM, from (Zhao et al., 2022b) XSum, Reddit-TIFU and SAMSum, from (Pang et al., 2022) for Arxiv and PubMed, from (Xiong et al., 2022) for GovReport and SummScreenFD, from (Xiao et al., 2022) for Multi-News and Multi-XScience. Due to a lack of reported results for other metrics, we only include ROUGE-2 scores for SOTA models. Best scores (outside of SOTA) are in bold.

Model	CNN/DM	XSum	Reddit-TIFU	SAMSum	Multi-XScience	Arxiv	PubMed	GovReport	SummScreenFD	Multi-News
Flan-UL2	0.000	0.012	0.000	0.017	0.000	–	–	–	–	–
Llama-2-7B	0.000	0.000	0.000	0.490	0.000	0.000	0.000	0.000	0.000	0.000
Llama-2-13B	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Xgen-7B	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mistral-7B	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 6: P-value of a 2-sample Kolmogorov-Smirnov test between the position distribution of bigrams in LLM-generated summaries and bigrams in reference summaries. We round numbers to 3 decimals. Numbers in gray correspond to non-significant differences (p-value above 0.05).

Dataset	Domain	# Docs	# Data points	# Sentences		# Words		# Tokens	
				Doc.	Summ.	Doc.	Summ.	Doc.	Summ.
Arxiv	Science	1.00	50	299.36	5.84	7,605.52	165.80	10,846.00	226.22
PubMed	Science (medical)	1.00	50	157.60	6.96	5,090.84	203.04	7,783.02	329.12
GovReport	Legal	1.00	50	445.64	22.80	13,308.70	656.10	17,462.90	883.10
SummScreen	TV transcripts	1.00	25	762.24	3.56	9,732.44	88.60	12,878.08	115.00
MultiNews	News	3.18	50	211.34	9.86	5,939.32	269.84	8,130.78	336.16
Overall	Mixed	1.48	225	332.24	10.50	8,180.13	297.57	11,258.16	407.15

Table 7: Statistics on the MiddleSum evaluation dataset, breaking down on each domain. **Doc.** is the source document, **Summ.** the ground-truth summary. **# Tokens** is calculated with Llama-2’s tokenizer.

C Baseline Performance

In Table 5, we report zero-shot performance with the 5 LLMs described in §2. We note that for standard-length datasets, Flan-UL2 is dominating, while for long-context ones, Llama-2-13B takes the lead.

D Statistical Significance on RQ1

In Table 6, we run a 2-sample Kolmogorov-Smirnov statistical significance test to compare the bigrams position distribution of LLMs with the reference summaries.

E More Results on MiddleSum

In Table 7, we show statistics on the MiddleSum evaluation dataset.

In Figs. 8a to 8d, we repeat the salient bigrams analysis from Fig. 1 for each of the four long-input LLMs on MiddleSum. We note that both *hierarchical* and *incremental* inference notably decrease reliance on lead bigrams compared to standard inference.

In Figs. 9 to 13, and in Table 8 (exact numbers), we show reference-based evaluation on each of the 5 subsets of MiddleSum.

F More Results on Long-Input Tasks

In Figs. 14 and 15, we conduct the analysis on increasingly longer input context from Fig. 6 on Arxiv and SummScreen, respectively.

G Software

We use the following Python libraries, all open source:

- *numpy*, version 1.24.3
- *torch*, version 2.0.1
- *scikit-learn*, version 1.0.2
- *sentencepiece*, version 0.1.97
- *nltk*, version 3.8.1
- *spacy*, version 3.6.0
- *scipy*, version 1.10.1
- *rouge-score*, version 0.1.2
- *bert-score*, version 0.3.13
- *summac*, version 0.0.03
- *tiktoken*, version 0.4.0
- *openai*, version 0.28.0
- *huggingface-hub*, version 0.17.2
- *datasets*, version 2.14.5
- *accelerate*, version 0.21.0
- *tokenizers*, version 0.14.1
- *transformers*, version 4.34.0

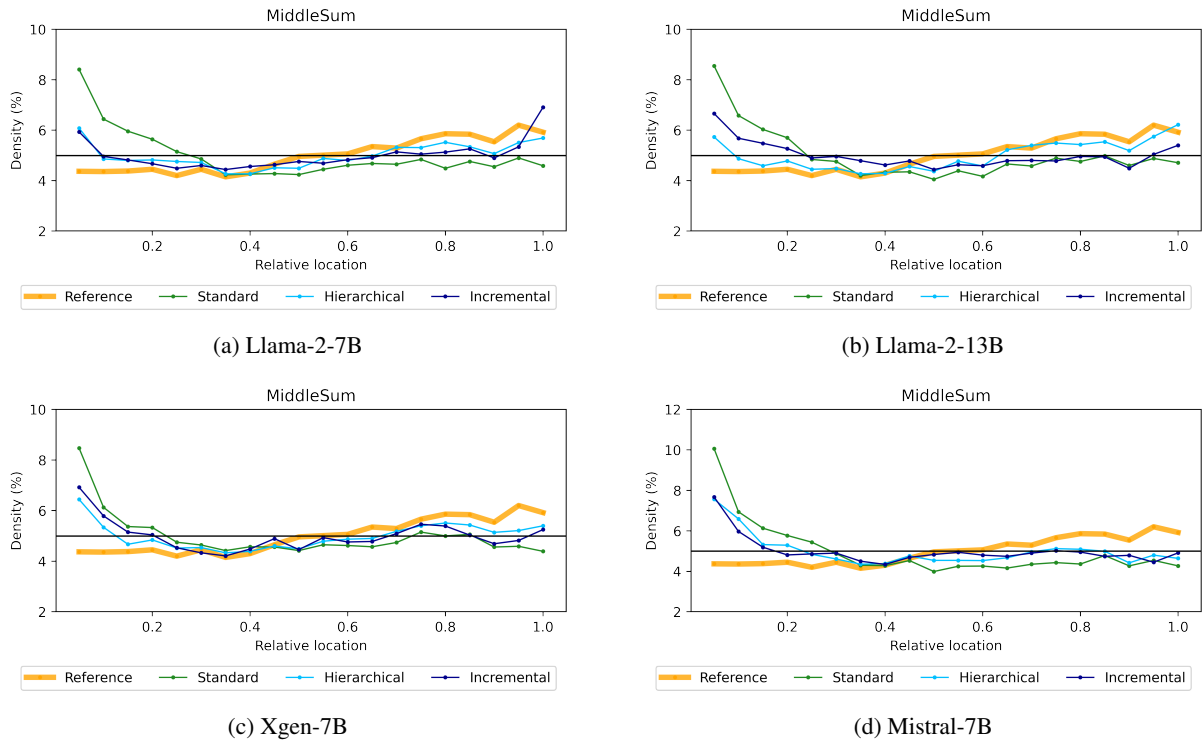


Figure 8: Relative location of summary bigrams within the source on MiddleSum for 4 LLMs.




Figure 9: Reference-based evaluation on the **Arxiv** subset of MiddleSum dataset.




Figure 10: Reference-based evaluation on the **PubMed** subset of MiddleSum dataset.




Figure 11: Reference-based evaluation on the **GovReport** subset of MiddleSum dataset.




Figure 12: Reference-based evaluation on the **SummScreenFD** subset of MiddleSum dataset.




Figure 13: Reference-based evaluation on the **Multi-News** subset of MiddleSum dataset.




Figure 14: Long-input summarization performance on **Arxiv** with 4 models and all 5 metrics. X-axis represents the truncated maximum source length. Xgen-7b and Mistral-7b cannot infer beyond 8k tokens.




Figure 15: Long-input summarization performance on **SummScreen** with 4 models and all 5 metrics. X-axis represents the truncated maximum source length. Xgen-7b and Mistral-7b cannot infer beyond 8k tokens.

Model	Metric	Inference	MiddleSum (MS)	MS/Arxiv	MS/PubMed	MS/GovReport	MS/SummScreen	MS/Multi-News
Llama-2-7B	ROUGE-2	Standard	10.96	12.62	10.97	13.26	4.07	10.43
		Hierarchical	11.33	14.63	13.36	13.51	4.85	7.06
		Incremental	12.56	14.43	13.54	17.26	5.19	8.68
	BERTScore	Standard	84.19	83.43	83.13	84.63	85.33	85.01
		Hierarchical	84.13	83.86	83.74	84.26	85.23	84.10
		Incremental	84.26	83.73	83.80	84.67	85.70	84.14
	A3CU	Standard	12.81	13.55	12.12	11.71	7.31	16.61
		Hierarchical	13.21	15.64	15.14	12.26	9.71	11.57
		Incremental	12.88	15.34	14.45	10.51	10.24	12.55
Llama-2-13B	ROUGE-2	Standard	11.07	11.68	11.63	13.56	5.09	10.38
		Hierarchical	10.24	13.45	13.26	10.21	4.93	6.71
		Incremental	11.90	12.53	13.84	17.34	5.06	7.33
	BERTScore	Standard	84.04	83.29	83.15	84.43	85.24	84.70
		Hierarchical	83.17	83.74	83.75	80.07	85.03	84.20
		Incremental	83.43	82.45	83.45	84.96	83.24	82.95
	A3CU	Standard	12.94	13.01	13.42	10.92	8.47	16.65
		Hierarchical	12.85	16.28	14.11	10.40	8.36	12.86
		Incremental	12.47	14.64	14.80	11.75	10.76	9.54
Xgen-7B	ROUGE-2	Standard	8.92	11.94	9.34	10.02	4.53	6.57
		Hierarchical	9.87	13.06	11.37	11.77	2.79	6.83
		Incremental	9.11	10.97	10.55	13.16	3.71	4.44
	BERTScore	Standard	82.51	82.45	81.96	82.65	83.28	82.60
		Hierarchical	82.86	82.73	82.56	82.67	83.79	83.03
		Incremental	82.49	81.94	81.82	83.11	83.39	80.83
	A3CU	Standard	12.50	15.14	12.85	11.99	8.48	12.05
		Hierarchical	11.74	14.55	12.61	10.76	7.33	11.25
		Incremental	11.51	12.73	13.15	9.13	8.97	9.69
Mistral-7B	ROUGE-2	Standard	7.16	9.28	9.56	5.86	2.19	6.41
		Hierarchical	10.78	11.67	13.71	13.36	4.76	7.39
		Incremental	10.15	11.36	12.66	12.89	3.46	7.02
	BERTScore	Standard	81.07	80.98	81.19	80.41	80.10	82.20
		Hierarchical	83.10	82.56	83.00	83.81	83.90	82.62
		Incremental	82.25	81.92	82.39	83.67	83.14	80.56
	A3CU	Standard	10.35	11.55	11.46	7.94	8.53	11.34
		Hierarchical	11.15	12.15	14.21	10.62	7.48	9.47
		Incremental	10.43	14.25	11.00	8.90	7.44	9.05

Table 8: Reference-based results for all models and inference methods on MiddleSum, breaking down by subset. The best number across inference methods is in bold.