

# Exploring Iterative Controllable Summarization with Large Language Models

Anonymous ACL submission

## Abstract

Large language models (LLMs) have demonstrated remarkable performance in abstractive summarization tasks. However, their ability to precisely control summary attributes (e.g., *length* or *topic*) remains underexplored, limiting their adaptability to specific user preferences. In this paper, we systematically explore the controllability of LLMs. To this end, we revisit summary attribute measurements and introduce iterative evaluation metrics, *failure rate* and *average iteration count* to precisely evaluate controllability of LLMs, rather than merely assessing errors. Our findings show that LLMs struggle more with numerical attributes than with linguistic attributes. To address this challenge, we propose a guide-to-explain framework (GTE) for controllable summarization. Our GTE framework enables the model to identify misaligned attributes in the initial draft and guides it in self-explaining errors in the previous output. By allowing the model to reflect on its misalignment, GTE generates well-adjusted summaries that satisfy the desired attributes with robust effectiveness, requiring surprisingly fewer iterations than other iterative approaches.

## 1 Introduction

Large language models (LLMs) have demonstrated superior performance in text summarization, outperforming encoder-decoder models by generating more contextually appropriate and natural summaries (Goyal et al., 2023; Zhang et al., 2024; Pu et al., 2023; Ryu et al., 2024b). However, given individuals’ diverse preferences for summary styles, it is essential to generate summaries tailored to individual needs (Zhang et al., 2023b). For instance, some users may prefer specific topic-focused summaries or retain exact phrases. While LLMs excel in generating fluent summaries, their ability to precisely control attributes remains underexplored (Liu et al., 2024), limiting their adaptability to diverse user preferences.

Controllable summarization has recently garnered attention (Zhong et al., 2021; Ahuja et al., 2022; Maddela et al., 2022; Mehra et al., 2023; Xu et al., 2023; Zhang et al., 2023b), with prior studies primarily utilizing encoder-decoder architectures (Mao et al., 2022; Zhang et al., 2022; Vig et al., 2022; He et al., 2022; Pagnoni et al., 2023; Wang et al., 2023; Urlana et al., 2024) that incorporate attribute-specific training signals to constrain summary generation. However, these systems lack scalability due to their dependence on attribute-specific training. In contrast, LLMs offer flexibility yet rely on in-context learning and prompting without explicit training-based constraints, hindering fine-grained attribute control (Yuan et al., 2024; Tang et al., 2023), necessitating further investigation.

To systematically explore the controllability of LLMs, we first revisit the measurements for four key attributes: *extractiveness*, *length*, *topic*, and *speaker*, and refine them for more precise measurement. Specifically, instead of previous attribute measurement strategies that rely solely on word presence for *topic*- or *speaker*-focused summaries, we adopt embedding similarity to incorporate semantic information into the measurements. Building on this refined attribute measurement, we move beyond the previous controllability evaluation, which primarily assesses the disparity between the reference and the generated summary (He et al., 2022; Zhang et al., 2023b). To fully explore the LLMs’ controllability, we shift focus from merely measuring attribute errors to evaluating whether LLMs can accurately control specified attributes via iterative refinement. Thus, we introduce two evaluation metrics: the *failure rate*—the proportion of control failures within the maximum iterations—and the *average iteration count* required for successful control. Our findings reveal that while LLMs excel at controlling linguistic attributes such as *topic* and *speaker*, they severely struggle with numerical attributes such as *extrac-*

*tiveness* and *length*. We assume that unlike linguistic attributes, which rely on semantic coherence, numerical ones demand strict quantitative constraints, making fine-grained control challenging.

To address this challenge, we propose a guide-to-explain (GTE) framework, which enables precise attribute control solely through LLMs without relying on additional attribute-specific training. We first design a step-by-step attribute identification phase to instruct the model on calculating misaligned attributes of its previously generated summary, then guide it to explain the rationales behind its errors. Through self-reflection, the model corrects its previous mistakes and generates a well-aligned summary in the regeneration phase. By integrating a self-refinement strategy—proven effective in complex reasoning tasks (Weng et al., 2023; Madaan et al., 2023; Dhuliawala et al., 2024; Gou et al., 2024)—into controllable summarization, we enhance LLMs’ attribute controllability while ensuring summary quality.

We evaluate GTE on mixed-attribute summarization datasets, MACSum<sub>Doc</sub> and MACSum<sub>Dial</sub> (Zhang et al., 2023b). GTE successfully controls each attribute with minimal iterations, significantly outperforming other iteration methods, while demonstrating robustness by consistently adjusting attributes across data samples. Further, we also demonstrate the high quality of the controlled summaries across multiple generic summarization evaluation metrics, including UniEval (Zhong et al., 2022) and QuestEval (Scialom et al., 2021). Finally, we analyze whether LLMs can control multiple attributes simultaneously, revealing their struggles in jointly managing correlated numerical attributes. Our contributions are as follows:

- We systematically explore LLM’s controllability in text summarization.
- We refine the measurement of summarization attributes and introduce *iterative evaluation* metrics for evaluating LLMs’ controllability.
- We propose a guide-to-explain (GTE), which guides the model to explain its misalignments and effectively adjusts misaligned attributes within just a few iterations.

## 2 Related work

**Controllable summarization** Controllable summarization has recently gained attention due to its

practical applications (Zhong et al., 2021; Ahuja et al., 2022; Maddela et al., 2022; Mehra et al., 2023; Xu et al., 2023; Zhang et al., 2023b; Ribeiro et al., 2023). Previous research employed encoder-decoder models to control attributes (Fan et al., 2018; Liu and Chen, 2021; Dou et al., 2021; He et al., 2022; Mao et al., 2022; Zhang et al., 2022; Goyal et al., 2022; Vig et al., 2022; Bahrainian et al., 2022; Liu et al., 2022; Pagnoni et al., 2023; Wang et al., 2023; Urlana et al., 2024). For example, CTRLSum (He et al., 2022) trains models by prepending a keyword sequence to the source document. Similarly, MACSum (Zhang et al., 2023b) adopts prompt learning by prepending each attribute’s value to the input source using a combination of hard prompts and soft prefixes. HYDRASUM (Goyal et al., 2022) leverages a single encoder and multiple decoder framework, using a mixture-of-experts where the decoders share probabilities to effectively control the attributes.

Most controllable summarization research relied on encoder-decoder frameworks. In addition, they required attribute-specific training or custom datasets to control each attribute, limiting the flexibility of attribute manipulation. Therefore, we propose a generalizable approach to enable flexible attribute control without the need for tailored training for each attribute, leveraging underexplored LLMs for controllable summarization (Tang et al., 2023; Yuan et al., 2024; Liu et al., 2024).

**Self-correction** Recently, the self-correction approach has been used to refine their initial attempts to solve complex problems (Weng et al., 2023; Shinn et al., 2023; Madaan et al., 2023; Dhuliawala et al., 2024; Gou et al., 2024), mirroring humans. In summarization tasks, self-correction has been employed to enhance the overall quality of summaries (Zhang et al., 2023a; Sun et al., 2024). Zhang et al. (2023a) utilizes iterative feedback from an evaluator to instruct ChatGPT to produce higher-quality summaries. Unlike prior works, we focus on generating summaries tailored to user preferences, which involve a multitude of factors to consider.

## 3 Problem formulation for controllable summarization

We conduct a systematic exploration of LLMs’ controllability evaluation. While generic evaluation metrics (e.g. *consistency*, *fluency*) favor higher scores, controllability requires aligning summaries with user-specified attribute values. Previous re-

Attribute	Metrics	Paper
Extractiveness	ROUGE, word overlap	Goyal et al. (2022); Zhang et al. (2023b)
Length	Absolute length, length ratio	Goyal et al. (2022); He et al. (2022); Maddela et al. (2022); Zhang et al. (2023b)
Topic, Query	ROUGE, LDA, topic word count, classifier	Zhong et al. (2021); He et al. (2022); Zhang et al. (2023b); Xu et al. (2023)
Speaker, Entity	ROUGE, speaker utterance word overlap	Maddela et al. (2022); Zhang et al. (2023b)

Table 1: Previous methods for measuring attributes.

search mainly assessed deviations from target values rather than ensuring precise control (Zhang et al., 2023b). To address this, we introduce iterative evaluation metrics, assessing whether LLMs can precisely adjust attributes, even through multiple iterations. Before evaluating a controllability, we first refine attribute measurement.

### 3.1 Controllable attribute measurement

We revisit attribute measurement to quantify key attributes for controllable summarization: *extractiveness*, *length*, *topic*, and *speaker*. Table 1 summarizes how previous controllable summarization studies have measured each attribute. However, the measurements for certain attributes have not yet been clearly defined. Thus, we outline our newly defined approach for attribute measurements below.

**Extractiveness** quantifies the degree of lexical overlap between a summary and its source document. Highly extractive summary is required when users need to retain the original context, such as in legal documents, whereas paraphrasing is preferred in general cases. To measure the *extractiveness*, MACSum (Zhang et al., 2023b) employs the average precision scores of ROUGE-2 and ROUGE-3. HYDRASUM (Goyal et al., 2022) calculates the proportion of words in the summary that appear in the source document as well as the average length of copied contiguous spans in the summary. Following the definition of *extractiveness*, we measure this attribute as the proportion of words in the summary directly reused from the source text.

**Length** refers to the number of words or sentences in the summary or the ratio of the summary’s length to the original text. By controlling the length, the amount of information in the summary can be adjusted according to user preferences. Prompts suggested by earlier works specify a fixed number of sentences, e.g., "3 sentences," but this approach fails to account for variations in sentence length and does not accurately reflect the summary’s actual length (Goyal et al., 2023; Liu et al., 2024; Yuan et al., 2024). As the summary length may

vary depending on the source document, we use the length ratio instead of absolute length.

**Topic** refers to generating a summary centered around one or multiple themes. Query-focused summarization (QFS), which generates summaries based on a specific query, and entity-based summarization, which focuses on a particular individual, are variations of topic-focused summarization. Zhang et al. (2023b) measured topic word frequency in summaries. Similarly, most QFS methods have relied solely on ROUGE scores, evaluating generated summaries by comparing them to human-annotated references (Zhong et al., 2021). However, even if topic words do not explicitly appear, the summary can still reflect the core context of the topic, especially for LLM-generated summaries, which tend to paraphrase the content. Therefore, rather than simply counting the frequency of word occurrences, we evaluate the semantic similarity between the summary and each topic-related word. We compute the embedding similarity  $\mathcal{B}$  between the topic word and each word in the summary  $s$ :  $\frac{1}{n} \sum_{i \in s} \mathcal{B}(\text{topic}, \text{word}_i)$ , where  $n$  is the number of words in the summary. If multiple topics  $k$  are present, we use the average embedding similarity across all topics:  $\frac{1}{k} \sum_{j \in k} \frac{1}{n} \sum_{i \in s} \mathcal{B}(\text{topic}_k, \text{word}_i)$ .

**Speaker** refers to generating a summary that focuses on the utterances of a specific speaker from a long document or dialogue. Zhang et al. (2023b) calculate the frequency of the speaker’s spoken words appearing in the summary. Similar to *topics*, measuring the proportion of words from a specific speaker’s dialogue included in the summary does not fully capture semantic elements. Therefore, we extract speaker utterances and construct speaker set  $\mathcal{U}$  and leverage BERTScore F1 (Zhang et al., 2020) to compute the embedding similarity between the summary  $s$  and  $\mathcal{U}$ :  $\text{BERTScore}(s, \mathcal{U})$ .

### 3.2 Iterative controllability evaluation

Previous evaluations of the model’s controllability were conducted using metrics such as mean absolute deviation (MAD) between predicted and

Attribute	Label	MACSum <sub>Doc</sub>			MACSum <sub>Dial</sub>		
		Distributions	Relabel	# of summaries	Distributions	Relabel	# of summaries
Extractiveness	normal	35.7 - 100.0% (85.2%)	85.0%	3731	53.2 - 100.0% (86.4%)	85.0%	1661
	high	55.0 - 100.0% (90.0%)	90.0%	287	63.0 - 100.0% (88.9%)	90.0%	340
	fully	84.6 - 100.0% (99.7%)	100.0%	260	75.9 - 100.0% (98.4%)	100.0%	337
Length	short	0.7 - 15.0% (4.8%)	7.5%	1059	0.2 - 20.8% (2.0%)	7.5%	300
	normal	0.5 - 48.6% (6.9%)	20.0%	2194	0.3 - 41.9% (3.7%)	20.0%	1693
	long	1.5 - 39.8% (13.9%)	32.5%	1025	0.7 - 32.4% (6.0%)	32.5%	345
Topic	-	74.8 - 88.8	74.0	2013	73.6 - 87.0	74.0	2317
Speaker	-	-	-	-	75.6 - 92.0	75.0	1796

Table 2: Data distributions of MACSum<sub>Doc</sub> and MACSum<sub>Dial</sub>.

reference summary attributes (He et al., 2022), or label-based control error rate (CER) and control correlation (Zhang et al., 2023b). While these metrics capture the differences between the reference and the model-generated summaries, they do not explicitly determine whether LLMs can ultimately exert effective control over the summaries. Since precisely controlling attributes in a single generation is challenging, we evaluate LLMs’ ability to iteratively refine and adjust attributes over multiple revisions. Specifically, we introduce two metrics: (1) the *failure rate*, which quantifies the proportion of cases where the model reaches predefined maximum iterations without achieving the desired modifications, and (2) the *average iteration count* required for successful attribute control. A failure or reaching the maximum number of iterations is denoted as  $\ominus$ . We set the maximum number of iterations to 20 due to cost constraints.

**Iteration threshold** We set attribute-specific thresholds and iteratively regenerate summaries until they are met. Each attribute is measured using the criteria outlined in Section 3.1 to determine its threshold. For *extractiveness* and *length*, we consider control successful if the attribute values fall within  $\pm 5$  of the target value. For *topic* and *speaker*, we use the minimum embedding similarity values of the reference summaries in the training dataset (Table 2) as thresholds to determine whether a summary is *topic*-focused or *speaker*-focused. The threshold can be adjusted based on the strictness of the evaluation criteria.

**Label reinterpretation** We use the two publicly available MACSum datasets (Zhang et al., 2023b) for controllable summarization. However, existing labels are ambiguous since the criteria are not numerically defined (e.g., how short must a summary be to qualify as *short*?). We believe that these ambiguous criteria may confuse LLMs, so we assign clear numerical values to each label. To provide de-

tailed criteria, we reinterpret the labels based on the attribute distributions in each training set. For *extractiveness*, we set labels as {normal: 85%, high: 90%, fully: 100%}, following the mean value. Unlike Zhang et al. (2023b), we define the summary length as a ratio of the original text rather than a fixed value. The existing labels do not distinctly differentiate these ratios, as the average length ratios between labels show minimal differences. Specifically, *short* is 2.0, while *normal* is 3.7, indicating a relatively small gap. To allow for greater variability, we expand the range. Importantly, our method generates summaries based on the specified numerical values, regardless of predefined labels. The broader range allows for a more adaptive and effective evaluation of LLM controllability. We set the lengths to {*short*: 7.5%, *normal*: 15%, *long*: 32.5%}, providing clearer distinctions with evaluating a broader range of controllability.

#### 4 Controllability of LLMs in text summarization

	Extractiveness	Length	Topic
Phi-3-medium	100.00% / $\ominus$	100.00% / $\ominus$	38.08% / 0.22
Llama3-8B	100.00% / $\ominus$	100.00% / $\ominus$	57.14% / 0.12
Llama3-70B	49.91% / 8.05	49.36% / 8.24	0.00% / 0.24
GPT-3.5	49.73% / 9.80	76.42% / 0.00	0.00% / 0.00
GPT-4o	39.31% / 6.63	69.84% / 0.00	0.38% / 0.02

Table 3: The left number represents the averaged control *failure rate*, and the right side denotes the *average iteration count* for successful control.

As research on leveraging LLMs for controllable summarization remain limited, we evaluate the controllability of various LLMs using the iterative evaluation method from Section 3.2. We first provide an initial control prompt and generate a summary. If the generated summary fails to meet the specified attribute threshold, the result is fed back into the LLM’s input, prompting it to regenerate until the attribute is correctly controlled. As shown in Table 3,

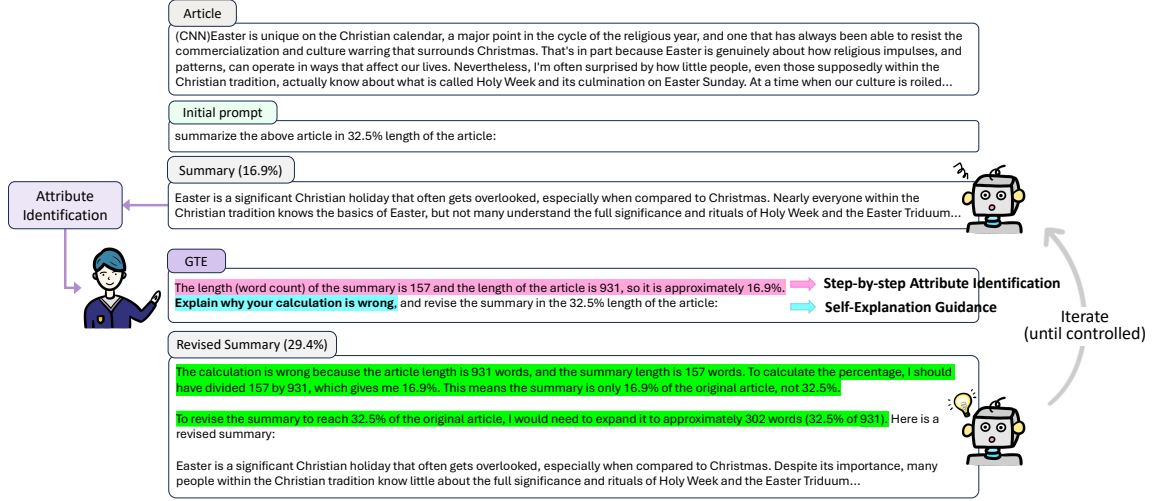


Figure 1: Overview of guide-to-explain system (GTE). The pink parts (■) represent the step-by-step attribute-identification, and the blue parts (■) correspond to the self-explanation guidance.

smaller-scale LLMs, such as Phi-3-medium (Abdin et al., 2024) and Llama3-8B (Dubey et al., 2024), partially control *topic*, but fail to control *extractiveness* and *length*. Similarly, large-scale LLMs such as Llama3-70B, GPT-3.5 (Brown, 2020), and GPT-4o (Achiam et al., 2023) effectively control *topic*, demonstrating low failure rates. However, they also struggle to control *extractiveness* and *length*, showing failure rates of around 50%. Notably, when initial attempts fail, GPT-3.5 and GPT-4o could not adjust despite multiple iterations, ultimately reaching the maximum iteration limit, resulting in an iteration count of zero. As a result, we find that generating summaries while controlling attributes remains challenging for LLMs, even with iterative attempts, especially for numerical attributes such as *extractiveness* and *length*.

## 5 Method: Guide-to-Explain (GTE)

Therefore, we introduce a guide-to-explain (GTE) framework to adjust various attributes via LLMs. As shown in Figure 1, our GTE framework consists of two key phase: the step-by-step attribute identification and the self-explanation guidance. We provide an step-by-step attribute identification to teach precise adjustment for incorrectly generated responses then guide LLMs to reflect by explaining the rationales behind their mistakes. Our approach allows the model to make appropriate adjustments in subsequent iterations.

### 5.1 Step-by-step attribute identification

We first instruct the LLM to generate an initial summary  $s'$  that reflects the specified attribute. If the

LLM fails to control the attributes accurately, we provide step-by-step attribute identification (SAI) to guide the model on how to adjust the attributes. LLM may struggle to measure numerical attributes such as *extractiveness* or the *length* ratio. Therefore, we instruct the model on how to measure the attribute step-by-step so that it can revise its previously generated summary precisely.

### 5.2 Self-explanation guidance

After the identification phase, we provide self-explanation guidance (SEG) to the model, guiding the model to explain why it initially failed to control the attributes. This mirrors how humans solve complex problems by reviewing their mistakes to improve future responses. Building on this, in the next iteration, the document ( $d$ ), initial instruction ( $i$ ), and previously generated summary ( $s'$ ) are provided as inputs, along with SAI and SEG. Despite LLMs being known to struggle with number-related tasks (Akhtar et al., 2023; Imani et al., 2023), our guidance helps the model effectively control numerical attributes by self-explaining its miscalculations before generating summaries, especially when combined with the step-by-step attribute identification phase. We introduce GTE as a framework that integrates step-by-step attribute identification and self-explanation guidance.

### 5.3 Overall process

By receiving  $[d; i, s'; \text{SAI}; \text{SEG}]$  as input, the model first reflects on the reasons for the initial error before generating a revised summary. If the revised summary still fails to satisfy the attributes, we re-

Model	Extractiveness ( $\downarrow / \downarrow$ )				Length ( $\downarrow / \downarrow$ )				Topic( $\downarrow$ )	Speaker( $\downarrow$ )
	normal	high	full	avg	short	normal	long	avg		
Phi-3-medium-Iter	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	38.08% / 0.22	-
Phi-3-medium-GTE	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	37.97% / 0.04	-
Llama3-8B-Iter	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	100.00% / $\circ$	57.14% / 0.12	-
Llama3-8B-GTE	12.63% / 3.52	11.63% / 2.53	0.00% / 1.46	11.70% / 3.26	26.40% / 3.08	10.92% / 2.26	13.18% / 3.85	14.99% / 2.80	25.56% / 0.91	-
Llama3-70B-Iter	54.82% / 8.44	37.21% / 7.47	2.70% / 3.78	49.91% / 8.05	18.40% / 6.58	54.61% / 10.42	67.44% / 12.00	49.36% / 8.24	0.00% / 0.24	-
Llama3-70B-SAI	26.55% / 6.57	18.60% / 7.81	0.00% / 1.86	24.14% / 6.52	4.80% / 5.42	2.73% / 3.81	10.85% / 4.84	5.12% / 4.39	0.00% / 0.10	-
Llama3-70B-GTE	<b>0.21%</b> / 3.28	<b>0.00%</b> / 2.83	<b>0.00%</b> / 1.50	<b>0.18%</b> / 3.22	<b>0.00%</b> / 1.10	<b>0.00%</b> / 1.61	<b>2.32%</b> / 3.14	<b>0.55%</b> / 1.90	0.00% / 0.01	-
GPT-3.5-Iter	45.18% / 9.80	60.47% / 0.00	94.59% / 0.00	49.73% / 9.80	53.60% / 0.00	80.89% / 0.00	88.37% / 0.00	76.42% / 0.00	0.00% / 0.00	-
GPT-3.5-GTE	17.56% / 3.86	51.16% / 5.00	67.57% / 4.00	23.58% / 3.90	5.60% / 4.63	44.03% / 6.62	78.29% / 7.00	43.33% / 5.95	0.00% / 0.00	-
GPT-4o-Iter	34.69% / 6.77	55.81% / 0.00	78.38% / 3.00	39.31% / 6.63	72.00% / 0.00	64.85% / 0.00	79.07% / 0.00	69.84% / 0.00	0.38% / 0.02	-
GPT-4o-SAI	35.12% / 5.50	48.84% / 15.50	62.16% / 6.00	38.03% / 6.13	60.00% / 8.79	61.09% / 9.40	78.29% / 2.00	64.90% / 8.60	0.00% / 0.04	-
GPT-4o-GTE	<b>0.00%</b> / 2.76	<b>0.00%</b> / 4.70	<b>0.00%</b> / 2.03	<b>0.00%</b> / 2.87	<b>0.00%</b> / 1.20	<b>0.00%</b> / 1.21	<b>0.00%</b> / 1.96	<b>0.00%</b> / 1.42	0.00% / 0.02	-

Table 4: The results of controllability measured on the MACSum<sub>Doc</sub> dataset. The bold denotes the best performance.

Model	Extractiveness ( $\downarrow / \downarrow$ )				Length ( $\downarrow / \downarrow$ )				Topic ( $\downarrow$ )	Speaker ( $\downarrow$ )
	normal	high	fully	avg	short	normal	long	avg		
Llama3-70B-Iter	31.78% / 8.13	43.59% / 8.40	8.16% / 5.39	29.63% / 7.59	12.00% / $\circ$	93.75% / 6.00	98.00% / $\circ$	81.79% / 6.00	0.00% / 0.01	0.00% / 0.00
Llama3-70B-SAI	14.41% / 5.91	23.08% / 5.31	<b>0.00%</b> / 3.72	13.27% / 5.50	<b>0.00%</b> / 1.25	62.05% / 5.70	92.00% / 9.33	57.10% / 5.62	0.00% / 0.02	0.00% / 0.00
Llama3-70B-GTE	<b>0.00%</b> / 2.31	<b>0.00%</b> / 2.56	4.08% / 3.64	<b>0.61%</b> / 2.49	<b>0.00%</b> / 1.00	<b>36.61%</b> / 4.73	<b>80.00%</b> / 5.70	<b>37.65%</b> / 4.53	0.00% / 0.01	0.00% / 0.00
GPT-4o-Iter	79.24% / 4.36	82.05% / 3.67	59.18% / 1.00	76.54% / 4.00	6.00% / $\circ$	98.21% / $\circ$	100.00% / $\circ$	84.26% / $\circ$	0.31% / 0.01	0.00% / 0.00
GPT-4o-SAI	84.75% / 4.00	87.18% / 1.50	53.06% / 5.10	80.25% / 4.32	2.00% / 4.50	96.43% / $\circ$	100.00% / $\circ$	82.41% / 4.50	0.00% / 0.01	0.00% / 0.00
GPT-4o-GTE	<b>17.80%</b> / 7.94	<b>25.64%</b> / 7.92	<b>8.16%</b> / 4.58	<b>17.28%</b> / 7.53	<b>0.00%</b> / 1.40	<b>9.82%</b> / 2.75	<b>44.00%</b> / 4.21	<b>13.58%</b> / 2.90	0.00% / 0.02	0.00% / 0.00

Table 5: The results of controllability measured on the MACSum<sub>Dial</sub> dataset.

peat the GTE process until the model generates an attribute-aligned summary. In Figure 1, we illustrate in detail how the GTE framework actually operates. If the initial draft fails to properly adjust an attribute, we first provide a step-by-step guide on how to measure the attribute. Then, we guide the model to explain the miscalculation and regenerate the summary. As intended, the model mimics step-by-step attribute identification, reflects on the summary based on the identification feedback, and generates an improved summary.

## 6 Experimental setup

We evaluate the controllability of various LLMs, including Phi-3-medium (Abdin et al., 2024), Llama3 series (Dubey et al., 2024), and GPT series (Brown, 2020; Achiam et al., 2023). To analyze model performance by size, we utilize both the 8B<sup>1</sup> and quantized 70B versions<sup>2</sup> of Llama3, as well as GPT-3.5 and GPT-4o. We use BERTScore (Zhang et al., 2020)<sup>3</sup> to measure embedding similarity.

We used two datasets for our experiments: MACSum<sub>Doc</sub> and the MACSum<sub>Dial</sub> datasets (Zhang et al., 2023b), which comprise committee meeting transcripts and news contents, respectively. Both datasets are designed for mixed-attribute summarization that controls multiple attributes simultaneously. Notably, only MACSum<sub>Dial</sub> has *speaker*

attribute. Since we evaluate LLM performance on individual attributes, we use attributes separately.

## 7 Results and Discussions

**Main results** We denote the naive iteration approach, which repeatedly adjusts attributes, as Iter. The strategy that only provides step-by-step attribute identification is defined as SAI, while our full guiding framework is referred to as GTE. As shown in Table 4, our GTE demonstrate remarkably lower failure rates and require fewer iterations when adjusting summaries across all attributes, including challenging numerical attributes in MACSum<sub>Doc</sub>. Surprisingly, our method reduced the failure rate to nearly 0% when applied to Llama3-70B and GPT-4o, successfully controlling both *extractiveness* and *length* within just 1–3 iterations. For smaller models such as Phi-3-medium and Llama3-8B, which initially exhibited high failure rates, our approach significantly reduced their failure rates, demonstrating its effectiveness across different model scales. In particular, for long length, the most challenging attribute, our method achieved a remarkably low failure rate of just 2.32% within an average of 3.14 iterations.

LLMs encounter more difficulties with the MACSum<sub>Dial</sub> dataset (Table 5). The dataset, which is derived from QMSum (Zhong et al., 2021), consists of lengthy and diverse content parliamentary and committee meetings, making it more challenging compared to the CNN-news-based MACSum<sub>Doc</sub>. Notably, length control proved to be

<sup>1</sup>meta-llama/meta-llama-3-8b-instruct

<sup>2</sup>casperhansen/llama-3-70b-instruct-awq

<sup>3</sup>[https://github.com/Tiiiger/bert\\_score](https://github.com/Tiiiger/bert_score)

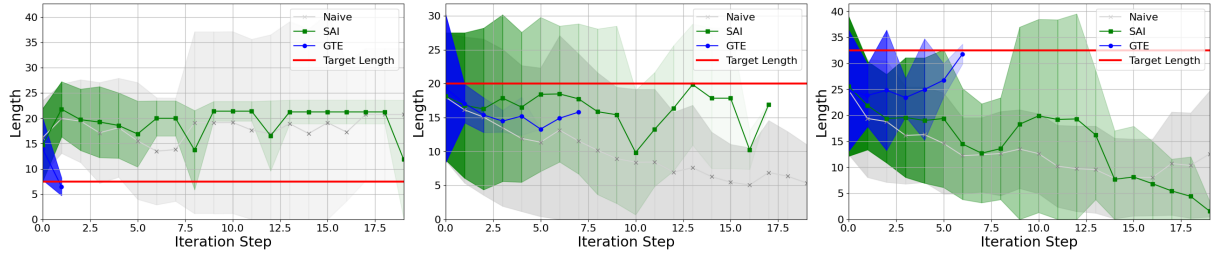


Figure 2: The graphs show how the length changes for each iteration. The intensity of the distribution color is proportional to the number of data points, and the markers represent the average values. The red line indicates the target length, with values of 7.5, 20, and 32.5 from left to right.

the most challenging attribute in the MACSum<sub>Dial</sub> dataset. This difficulty is likely due to the original dataset being a lengthy parliamentary corpus, making it inherently challenging to generate summaries of a specific length. While the model handled short-length summaries well, the difficulty increased significantly as the requested summary length grew. In fact, for long length, both GPT-4o-Iter and GPT-4o-SAI showed a 100% failure rate. However, our framework meaningfully improved length controllability. With GPT-4o, the average failure rate dropped below 50%. Especially, for normal-length summaries, the failure rate further reduced from over 90% to 9.82%. Regarding *extractiveness*, the Iter and SAI of GPT-4o exhibit relatively low iteration count since the models mostly exceed the maximum iteration. While they fail nearly 80%, our GTE demonstrates a markedly lower failure rate at 17.28% with low iterations, demonstrating the effectiveness of our framework.

**Gradual change across iteration steps** To analyze how attribute changes each step, we track *length* adjustments per iteration (Figure 2). While all methods start with a similar distribution at the initial points, GTE consistently converges within approximately three iterations, maintaining a stable length adjustment pattern across samples. In contrast, Iter and SAI show inconsistent changes across samples, leading higher variance in length changes. This demonstrates that our method enables robust attribute control with fewer iterations, regardless of the data sample. For the experiment, we utilize Llama3-70B and randomly selected 110 samples from the MACSum<sub>Doc</sub> test set.

**Attribute types** We observe that LLMs control linguistic attributes (*topic* and *speaker*) better than numerical attributes (*extractiveness* and *length*). This aligns with previous research in mathematical reasoning, where LLM struggle with numer-

ical features (Akhtar et al., 2023), highlighting a broader challenge in precisely handling numerical constraints. From the perspective of the summarization task, *extractiveness* and *length* control the structure of the summary, whereas *topic* and *speaker* control its contents. Our findings suggest that LLMs are proficient at adjusting content to align with user preferences, but they struggle to generate summaries in a desired structural format.

**Quality of controlled summary** We evaluate the quality of summaries generated by GTE. We mainly use UniEval (Zhong et al., 2022) and QuestEval (Scialom et al., 2021), as they highly correlate with human judgement and assess the overall quality of the summary itself. UniEval is a multi-dimensional evaluator that assesses *coherence*, *consistency*, *fluency*, and *relevance* of summaries. QuestEval measures precision and recall by leveraging question answering framework to compare the content between the source document and the generated summary without relying on the reference summary. Table 6 shows that our method’s summaries outperform across all UniEval dimensions and QuestEval, demonstrating effective attribute control while maintaining overall summary quality. *Relevance* assesses how well a summary retains key information compared to the reference. While Iter and SAI generate misaligned summaries with lower *relevance* scores, GTE effectively aligns them, leading to a substantial gain.

Although previous studies have shown that ROUGE scores (Lin, 2004) are insufficient for assessing summary quality (Scialom et al., 2021; Zhong et al., 2022; Ryu et al., 2024a), and given that our goal is to control the summary rather than make it similar to the reference, we still include ROUGE and BERTScore (Zhang et al., 2020) in our evaluation to provide clearer assessment. Our framework exhibits higher scores than other ap-

Model	UniEval					QuestEval	BERTScore	ROUGE-1
	Coherence	Consistency	Fluency	Relevance	Overall			
Iter (Ext)	0.820	0.800	0.859	0.696	0.794	0.523	0.826	0.194
SAI (Ext)	0.884	0.843	0.905	0.785	0.864	0.554	0.848	0.229
Iter (Len)	0.836	0.803	0.836	0.759	0.808	0.484	0.829	0.235
SAI (Len)	0.934	0.834	0.942	0.887	0.899	0.548	0.867	0.270
GTE (Ext)	<b>0.941</b>	<b>0.873</b>	0.937	0.880	<b>0.908</b>	<b>0.590</b>	0.861	0.236
GTE (Len)	0.937	0.840	<b>0.944</b>	<b>0.901</b>	0.905	0.553	<b>0.868</b>	<b>0.272</b>

Table 6: Quality of the controlled summaries.

proaches, demonstrating through various evaluation metrics that GTE not only enhances controllability, but also improves overall summary quality.

## 8 Mixed attributes

We extend our evaluation to assess whether LLMs could precisely handle mixed-attribute control. While models manage to control linguistic attributes but struggle with numerical attributes. Simultaneous control over all attributes remains challenging for all iterative methods, including GTE. Our GTE framework guides LLMs to identify the causes of their errors and regenerate summaries by incorporating this feedback. However, in a mixed attribute setting, the model needs to handle multiple SAI and GS for each attribute simultaneously, increasing the cognitive load and making precise control of all attributes more challenging. Therefore, unlike single-attribute evaluation, which assesses whether attributes are precisely controlled, we evaluate mixed-attribute control by measuring errors using root mean squared error (rMSE). We compare the error between the attributes of the generated summary and the requested values, providing a more flexible evaluation of attribute control.

**Sequential-planning** Discovering the challenges in precisely controlling all attributes simultaneously, we introduce a sequential planning strategy, *min-planning*, which gradually adjusts attributes from the ill-controlled with the initial draft using GTE. Figure 3 shows the results comparing single attribute control with iterations to mixed-attribute control using *min-planning* on the MACSum<sub>Doc</sub> dataset. We refer to the initial summary in the mixed-attribute control as the *mixed-draft*. The *min-planning* method shows a modest improvement in controlling both attributes compared to the *mixed-draft*. Attributes are still not fully controlled compared to single-attribute models, highlighting the challenge of balancing multiple attributes. We anticipate that modifying one attribute often dis-

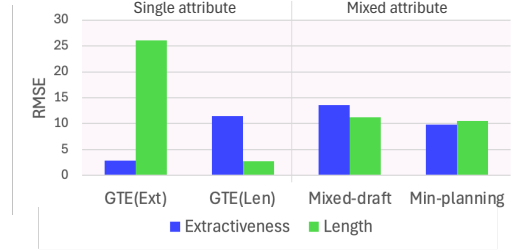


Figure 3: Performance in mixed-attribute.

rupts previously adjusted attributes due to correlations. For example, even if *length* is adjusted first, it may still change when *extractiveness* is controlled. Additionally, *min-planning* only adjusts each attribute once without iterations, which may account for its failure to fully control the attributes. A single refinement is often insufficient to control all attributes, whereas GTE iteratively regenerates the summary until the attribute is successfully adjusted in single-attribute control. Exploring ways for LLMs to control multiple attributes simultaneously would be promising future work.

## 9 Conclusion

In this work, we systematically explore the controllability of LLMs. To this end, we revisit the measurement of various attributes. We evaluate the controllability of LLMs via iterative assessment, finding that LLMs struggle more with numeric attribute control than linguistic attributes. To address this, we propose a guide-to-explain (GTE) approach, where the model is guided to explain its misalignments through attribute identification and then grounds this explanation to generate better-controlled summaries in subsequent iterations. GTE enables LLMs to control challenging numerical attributes with lower failure and fewer iterations. Furthermore, we demonstrate the high quality of controlled summaries using various generic summarization evaluation metrics.

## Limitation

We explore the controllability of various attributes in LLMs and introduced a novel guide-to-explain (GTE) framework to address challenges in numerical attributes. While GTE enhanced successfully control over challenging numerical attributes, it still struggled with highly correlated mixed numerical attributes. Additionally, *min-planning*, which adjusts attributes in order of least alignment, also faced difficulties achieving precise control. Even after properly adjusting one attribute, modifying the correlated numerical attribute caused the previously adjusted attribute to change. We believe further research could explore more effective methods for addressing these challenges.

## Ethics

We used publicly available MACSum datasets for our research, conducting experiments with Phi-3, Llama3<sup>4</sup>, GPT-3.5, and GPT-4o from April to October 2024.

## References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ojas Ahuja, Jiacheng Xu, Akshay Gupta, Kevin Horecka, and Greg Durrett. 2022. **ASPECTNEWS: Aspect-oriented summarization of news documents**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6494–6506, Dublin, Ireland. Association for Computational Linguistics.
- Mubashara Akhtar, Abhilash Shankarampeta, Vivek Gupta, Arpit Patil, Oana Cocarascu, and Elena Simperl. 2023. **Exploring the numerical reasoning capabilities of language models: A comprehensive analysis on tabular data**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15391–15405, Singapore. Association for Computational Linguistics.

<sup>4</sup>Meta Llama3 Community License, Copyright © Meta Platforms, Inc. All Rights Reserved. More details can be found at: [Llama3 License](#)

- Seyed Ali Bahrainian, Sheridan Feucht, and Carsten Eickhoff. 2022. **NEWS: A corpus for news topic-focused summarization**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 493–503, Dublin, Ireland. Association for Computational Linguistics.
- Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2024. **Chain-of-verification reduces hallucination in large language models**. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3563–3578, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2021. **GSum: A general framework for guided neural abstractive summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4830–4842, Online. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Angela Fan, David Grangier, and Michael Auli. 2018. **Controllable abstractive summarization**. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujia Yang, Nan Duan, and Weizhu Chen. 2024. **CRITIC: Large language models can self-correct with tool-interactive critiquing**. In *The Twelfth International Conference on Learning Representations*.
- Tanya Goyal, Nazneen Rajani, Wenhao Liu, and Wojciech Kryscinski. 2022. **HydraSum: Disentangling style features in text summarization with multi-decoder models**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 464–479, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tanya Goyal et al. 2023. **News summarization and evaluation in the era of gpt-3**. *Preprint*, arXiv:2209.12356.
- Junxian He, Wojciech Kryscinski, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2022. **CTRLsum: Towards generic controllable text summarization**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5879–5915, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

708	Shima Imani, Liang Du, and Harsh Shrivastava. 2023.	767
709	<a href="#">MathPrompter: Mathematical reasoning using large</a>	768
710	<a href="#">language models</a> . In <i>Proceedings of the 61st Annual</i>	769
711	<i>Meeting of the Association for Computational</i>	770
712	<i>Linguistics (Volume 5: Industry Track)</i> , pages 37–	771
713	42, Toronto, Canada. Association for Computational	772
714	Linguistics.	773
715	Chin-Yew Lin. 2004. <a href="#">ROUGE: A package for auto-</a>	774
716	<a href="#">matic evaluation of summaries</a> . In <i>Text Summariza-</i>	
717	<i>tion Branches Out</i> , pages 74–81, Barcelona, Spain.	
718	Association for Computational Linguistics.	
719	Yixin Liu, Alexander Fabbri, Jiawen Chen, Yilun Zhao,	
720	Simeng Han, Shafiq Joty, Pengfei Liu, Dragomir	
721	Radev, Chien-Sheng Wu, and Arman Cohan. 2024.	
722	<a href="#">Benchmarking generation and evaluation capabili-</a>	
723	<a href="#">ties of large language models for instruction control-</a>	
724	<a href="#">lable summarization</a> . In <i>Findings of the Association</i>	
725	<i>for Computational Linguistics: NAACL 2024</i> , pages	
726	4481–4501, Mexico City, Mexico. Association for	
727	Computational Linguistics.	
728	Yizhu Liu, Qi Jia, and Kenny Zhu. 2022. <a href="#">Length control</a>	
729	<a href="#">in abstractive summarization by pretraining informa-</a>	
730	<a href="#">tion selection</a> . In <i>Proceedings of the 60th Annual</i>	
731	<i>Meeting of the Association for Computational Lin-</i>	
732	<i>guistics (Volume 1: Long Papers)</i> , pages 6885–6895,	
733	Dublin, Ireland. Association for Computational Lin-	
734	guistics.	
735	Zhengyuan Liu and Nancy Chen. 2021. <a href="#">Controllable</a>	
736	<a href="#">neural dialogue summarization with personal named</a>	
737	<a href="#">entity planning</a> . In <i>Proceedings of the 2021 Con-</i>	
738	<i>ference on Empirical Methods in Natural Language</i>	
739	<i>Processing</i> , pages 92–106, Online and Punta Cana,	
740	Dominican Republic. Association for Computational	
741	Linguistics.	
742	Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler	
743	Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon,	
744	Nouha Dziri, Shrimai Prabhumoye, Yiming Yang,	
745	Shashank Gupta, Bodhisattwa Prasad Majumder,	
746	Katherine Hermann, Sean Welleck, Amir Yazdan-	
747	bakhsh, and Peter Clark. 2023. <a href="#">Self-refine: Itera-</a>	
748	<a href="#">tive refinement with self-feedback</a> . In <i>Advances in</i>	
749	<i>Neural Information Processing Systems</i> , volume 36,	
750	pages 46534–46594. Curran Associates, Inc.	
751	Mounica Maddela, Mayank Kulkarni, and Daniel	
752	Preotiuc-Pietro. 2022. <a href="#">EntSUM: A data set for entity-</a>	
753	<a href="#">centric extractive summarization</a> . In <i>Proceedings</i>	
754	<i>of the 60th Annual Meeting of the Association for</i>	
755	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	
756	pages 3355–3366, Dublin, Ireland. Association for	
757	Computational Linguistics.	
758	Ziming Mao, Chen Henry Wu, Ansong Ni, Yusen Zhang,	
759	Rui Zhang, Tao Yu, Budhaditya Deb, Chenguang	
760	Zhu, Ahmed Awadallah, and Dragomir Radev. 2022.	
761	<a href="#">DYLE: Dynamic latent extraction for abstractive</a>	
762	<a href="#">long-input summarization</a> . In <i>Proceedings of the</i>	
763	<i>60th Annual Meeting of the Association for Compu-</i>	
764	<i>tational Linguistics (Volume 1: Long Papers)</i> , pages	
765	1687–1698, Dublin, Ireland. Association for Compu-	
766	tational Linguistics.	
	Dhruv Mehra, Lingjue Xie, Ella Hofmann-Coyle,	
	Mayank Kulkarni, and Daniel Preotiuc-Pietro. 2023.	
	<a href="#">EntSUMv2: Dataset, models and evaluation for more</a>	
	<a href="#">abstractive entity-centric summarization</a> . In <i>Proceed-</i>	
	<i>ings of the 2023 Conference on Empirical Methods</i>	
	<i>in Natural Language Processing</i> , pages 5538–5547,	
	Singapore. Association for Computational Linguis-	
	tics.	
	Artidoro Pagnoni, Alex Fabbri, Wojciech Kryscinski,	
	and Chien-Sheng Wu. 2023. <a href="#">Socratic pretraining:</a>	
	<a href="#">Question-driven pretraining for controllable summa-</a>	
	<a href="#">rization</a> . In <i>Proceedings of the 61st Annual Meeting</i>	
	<i>of the Association for Computational Linguistics (Vol-</i>	
	<i>ume 1: Long Papers)</i> , pages 12737–12755, Toronto,	
	Canada. Association for Computational Linguistics.	
	Xiao Pu, Mingqi Gao, and Xiaojun Wan. 2023.	
	Summarization is (almost) dead. <i>arXiv preprint</i>	
	<i>arXiv:2309.09558</i> .	
	Leonardo F. R. Ribeiro, Mohit Bansal, and Markus	
	Dreyer. 2023. <a href="#">Generating summaries with control-</a>	
	<a href="#">lable readability levels</a> . In <i>Proceedings of the 2023</i>	
	<i>Conference on Empirical Methods in Natural Lan-</i>	
	<i>guage Processing</i> , pages 11669–11687, Singapore.	
	Association for Computational Linguistics.	
	Sangwon Ryu, Heejin Do, Yunsu Kim, Gary Lee, and	
	Jungseul Ok. 2024a. <a href="#">Multi-dimensional optimization</a>	
	<a href="#">for text summarization via reinforcement learning</a> .	
	In <i>Proceedings of the 62nd Annual Meeting of the</i>	
	<i>Association for Computational Linguistics (Volume 1:</i>	
	<i>Long Papers)</i> , pages 5858–5871, Bangkok, Thailand.	
	Association for Computational Linguistics.	
	Sangwon Ryu, Heejin Do, Yunsu Kim, Gary Geunbae	
	Lee, and Jungseul Ok. 2024b. <a href="#">Key-element-informed</a>	
	<a href="#">slm tuning for document summarization</a> . In <i>Inter-</i>	
	<i>speech 2024</i> , pages 1940–1944.	
	Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier,	
	Benjamin Piwowarski, Jacopo Staiano, Alex Wang,	
	and Patrick Gallinari. 2021. <a href="#">QuestEval: Summariza-</a>	
	<a href="#">tion asks for fact-based evaluation</a> . In <i>Proceedings of</i>	
	<i>the 2021 Conference on Empirical Methods in Natu-</i>	
	<i>ral Language Processing</i> , pages 6594–6604, Online	
	and Punta Cana, Dominican Republic. Association	
	for Computational Linguistics.	
	Noah Shinn, Federico Cassano, Edward Berman, Ash-	
	win Gopinath, Karthik Narasimhan, and Shunyu Yao.	
	2023. <a href="#">Reflexion: Language agents with verbal rein-</a>	
	<a href="#">forcement learning</a> . <i>Preprint</i> , arXiv:2303.11366.	
	Shichao Sun, Ruifeng Yuan, Ziqiang Cao, Wenjie Li,	
	and Pengfei Liu. 2024. <a href="#">Prompt chaining or stepwise</a>	
	<a href="#">prompt? refinement in text summarization</a> . In <i>Find-</i>	
	<i>ings of the Association for Computational Linguistics</i>	
	<i>ACL 2024</i> , pages 7551–7558, Bangkok, Thailand	
	and virtual meeting. Association for Computational	
	Linguistics.	
	Yuting Tang, Ratish Puduppully, Zhengyuan Liu, and	
	Nancy Chen. 2023. <a href="#">In-context learning of large lan-</a>	
	<a href="#">guage models for controlled dialogue summarization:</a>	



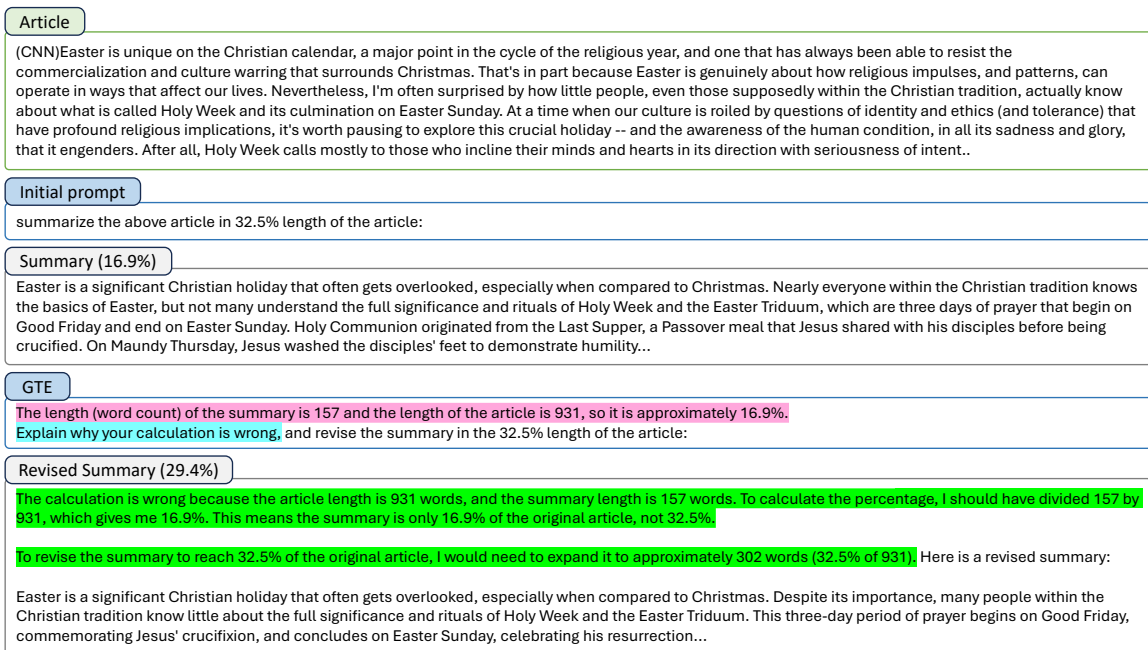


Figure 4: Length guide-to-explain (GTE).

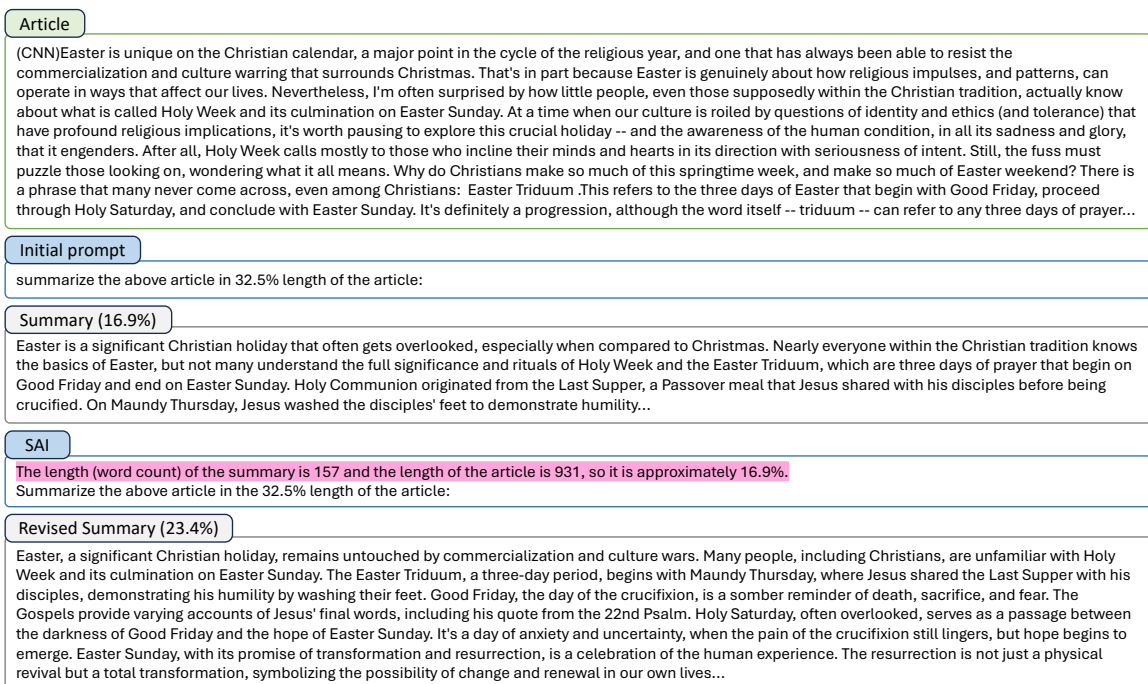


Figure 5: Length attribute-identification (SAI).

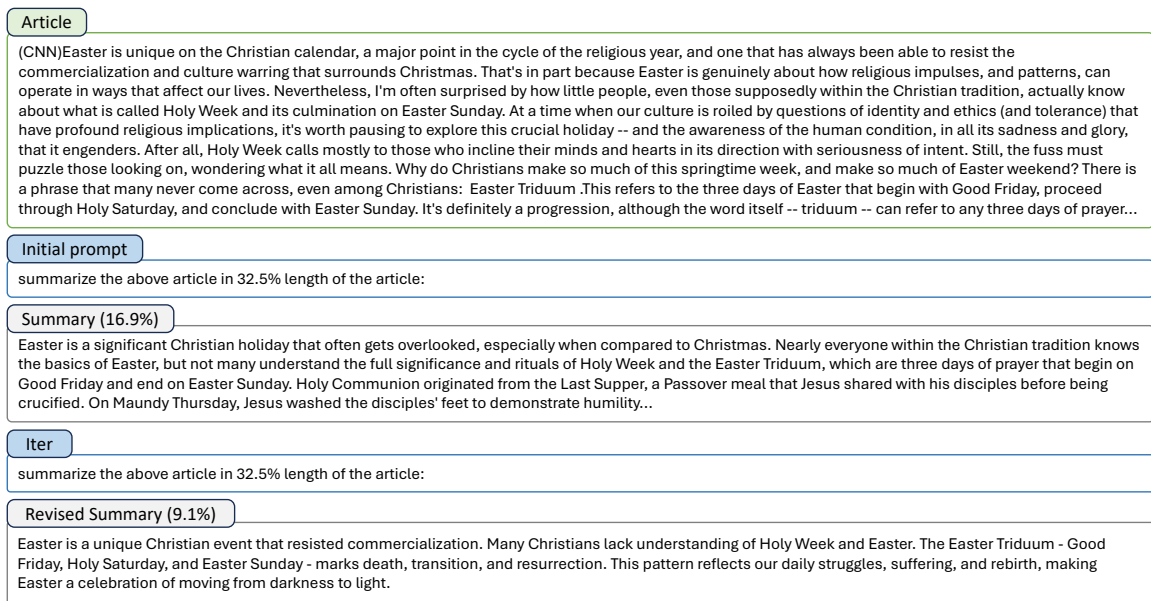


Figure 6: Length iteration (Iter).

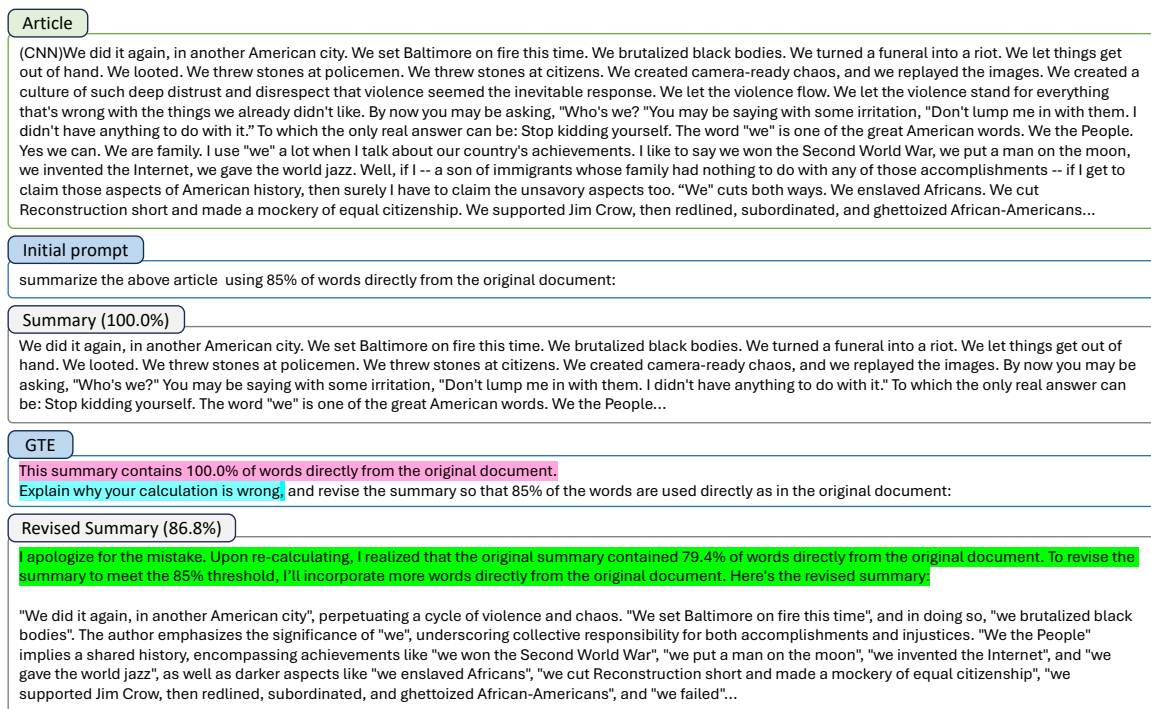


Figure 7: Extractiveness guide-to-explain (GTE).

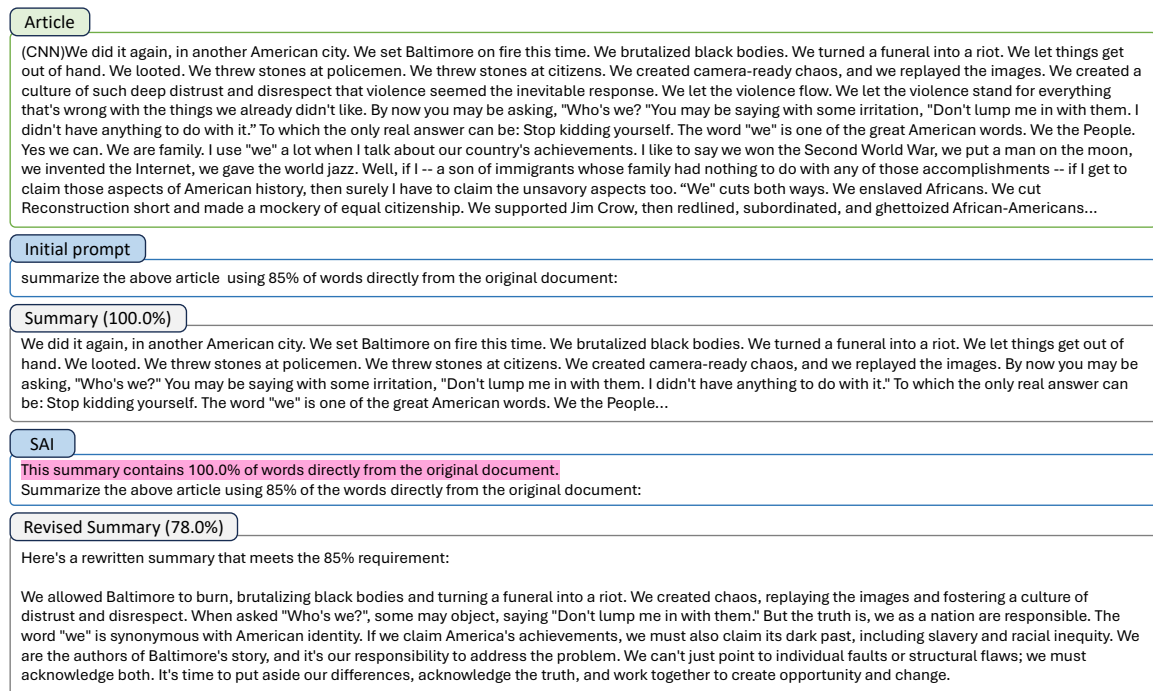


Figure 8: Extractiveness attribute-identification (SAI).

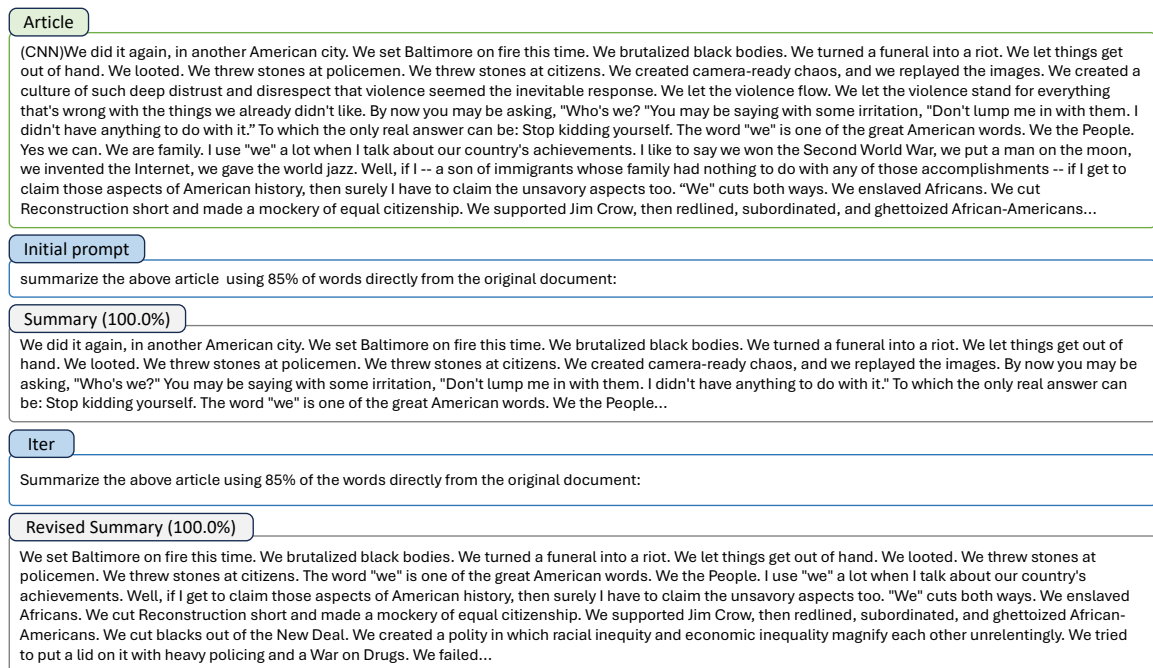


Figure 9: Extractiveness iteration (Iter).

Article

(CNN)Prison life won't be pretty for Aaron Hernandez, the former NFL player and convicted murderer sentenced to life without parole. After correction officers evaluate him, he will be shipped to Massachusetts' flagship maximum-security prison, one of the most high-tech jails in the United States with no history of breakouts: the Souza-Baranowski Correctional Center in Shirley, about 40 miles outside downtown Boston. It's called Souza, for short, and it's the state's newest prison, opened in 1998, with a matrix of 366 cameras recording live 24 hours a day and a microwave detection perimeter without wire. "I don't know the date, but he'll be going there. That's the maximum-security facility," Department of Corrections spokesman Darren Duarte said. Legal advocates for inmates describe Souza as sterile and violent at once. Its diverse demographic includes the young and the old, many of whom are also doing life. One stubborn problem is that opiates are smuggled to inmates, the legal advocates said. "It's very shiny and clean looking and very sterile," said Leslie Walker, executive director of Prisoners' Legal Services of Massachusetts, who has been visiting the Souza prison about every six weeks for the past 15 years and serves indigent prisoners there. But, she added: "It is a very dangerous prison that is right now experiencing a veritable flood of opiates."...

Initial prompt

summarize the above article focusing on hernandez:

Summary (72.6)

We did it again, in another American city. We set Baltimore on fire this time. We brutalized black bodies. We turned a funeral into a riot. We let things get out of hand. We looted. We threw stones at policemen. We threw stones at citizens. We created camera-ready chaos, and we replayed the images. By now you may be asking, "Who's we?" You may be saying with some irritation, "Don't lump me in with them. I didn't have anything to do with it." To which the only real answer can be: Stop kidding yourself. The word "we" is one of the great American words. We the People...

GTE

The similarity between this summary and the topic hernandez is 72.6, so this summary is not centered around the topic hernandez. Explain why you were unable to summarize around the topic, and revise the summary to center around the topic hernandez:

Revised Summary (76.7%)

I was unable to summarize around the topic hernandez because the original article provided a lot of context about the prison system and the conditions that hernandez will face, but did not provide much information that is specific to hernandez himself. The article does not delve into hernandez's thoughts, feelings, or personal experiences, and instead focuses on the general details of his situation.

Here is a revised summary that centers around the topic hernandez:

Aaron Hernandez, the former NFL player and convicted murderer, is facing a harsh reality as he begins his life sentence without parole. After being convicted, Hernandez was immediately transferred to the Massachusetts Correctional Institution-Cedar Junction in Walpole, where he will be processed before being sent to the Souza-Baranowski Correctional Center in Shirley. As a high-profile inmate, Hernandez's safety will be a concern for prison officials, who may need to take measures to protect him from other inmates or gang members. Despite his celebrity status, Hernandez will face harsh conditions, including being locked in his cell for 19 hours a day.

Figure 10: Topic guide-to-explain (GTE).