

# Understanding LLM Reasoning for Abstractive Summarization

Haohan Yuan    Haopeng Zhang\*  
ALOHA Lab, University of Hawaii at Manoa  
{haohany, haopengz}@hawaii.edu

## Abstract

Reasoning has substantially improved Large Language Models (LLMs) on analytical tasks such as mathematics and code generation, but its value for abstractive summarization remains unclear. To address this gap, we adapt general reasoning strategies to the summarization setting and conduct a large-scale comparative study of 8 reasoning strategies and 3 Large Reasoning Models (LRMs) across 8 diverse datasets, evaluating both summary quality and factual faithfulness. Our results show that reasoning is not a universal solution and its effectiveness depends strongly on the strategy and the summarization setting. In particular, we find a trade-off between summary quality and factual faithfulness. Explicit reasoning strategies often improve reference-based quality, but may weaken factual grounding, whereas implicit reasoning in LRMs shows the opposite tendency. We further find that increasing an LRM’s internal reasoning budget does not reliably improve summarization and can even reduce factual consistency. These findings suggest that, for summarization, more reasoning is not always better. Effective reasoning should preserve faithful compression rather than induce over-elaboration. Our source code is publicly available.<sup>1</sup>

## 1 Introduction

LLMs have demonstrated remarkable capabilities across many natural language processing tasks (Zhang et al., 2024c; OpenAI et al., 2024b; Yuan et al., 2025; Zhang, 2025). A significant recent development is the paradigm of LLM reasoning, where models are prompted to articulate intermediate thoughts before responding. While techniques such as Chain-of-Thought (CoT) (Wei et al., 2022) and Self-Consistency (Wang et al., 2022)

have established state-of-the-art performance in tasks like mathematics and logical inference (Kojima et al., 2022; Huang and Chang, 2023; Chen et al., 2023; Tan et al., 2025), their utility for abstractive summarization remains largely assumed rather than proven.

Despite the surge of interest in reasoning, its application to summarization lacks systematic study. Prior work by Adams et al. (2023) found that step-wise reasoning can influence the balance between informativeness and readability in news summarization. Similarly, Chang et al. (2023) utilized hierarchical reasoning for book-length texts, while Zhang et al. (2023b) demonstrated the utility of reflective revision and refinement. However, summarization differs fundamentally from logic puzzles, as it requires information compression rather than derivation. Consequently, several critical questions remain unanswered: **RQ1:** Does reasoning actually improve summarization performance compared to standard prompting? **RQ2:** When do reasoning strategies excel across different datasets and context? **RQ3:** How do different reasoning methods impact summarization results?

To bridge this gap, we present the first large-scale, systematic evaluation of reasoning methods in LLM-based summarization. We establish a unified evaluation framework to assess eight representative strategies, categorized into three paradigms: *Augmentation-based* methods, including Chain-of-Thought (CoT), Cited Summarization (Cite), Extract-to-Abstract (E2A), and Question-Answer Guided (QAG); *Organization-based* methods, including Decomposition (Deco) and Plan-then-Write (Plan); and *Reflective* methods, including Iterative Refine (IR) and Self-Consistency (SC). We propose task-specific adaptations of Cite, QAG, Plan, and SC for abstractive summarization (via prompts and rubrics) and evaluate them within a unified framework. Furthermore, we benchmark these explicit reasoning methods against three lead-

\*corresponding author

<sup>1</sup>[https://github.com/alohalab-ai/Summarization\\_reasoning](https://github.com/alohalab-ai/Summarization_reasoning)

ing Large Reasoning Models (LRMs) (OpenAI et al., 2024a): o1, o3, and GPT-5.

To ensure robustness, our evaluation spans eight datasets covering diverse domains, lengths, and formats, including news, dialogue, social media, knowledge bases, scientific writing, narrative novels, and structured tables. We employ a comprehensive suite of metrics: reference-based quality (ROUGE (Lin, 2004), BERTScore (Zhang et al., 2020)), factual consistency (SummaC (Laban et al., 2022), AlignScore (Zha et al., 2023)), LLM-as-a-judge (G-Eval (Liu et al., 2023)), and human evaluation.

Our empirical analysis yields counter-intuitive findings that challenge prevailing assumptions. First, **reasoning is not a panacea**. A simple *Vanilla* prompt often matches or surpasses complex reasoning strategies, particularly in few-shot settings. Second, we observe a distinct **trade-off between quality and faithfulness**. Explicit reasoning strategies like SC and IR tend to improve linguistic fluency and reference-based scores, but often at the cost of hallucination. Conversely, implicit reasoning within LRMs excels at factual grounding but may sacrifice stylistic polish. We further note that LLM-as-a-judge metrics tend to overestimate faithfulness compared to human evaluation.

The contributions of this work are fourfold: (1) We adapt general reasoning paradigms (Cite, QAG, Plan, and SC) to the context of abstractive summarization. (2) We establish the first systematic evaluation for reasoning in summarization, encompassing 8 explicit methods, 3 LRMs, and 8 diverse datasets. (3) We identify a critical trade-off between summary quality and factual faithfulness, demonstrating that increased reasoning complexity does not guarantee better summarization. (4) We provide practical guidance: explicit reasoning strategies (SC, IR) are optimal for zero-shot summary quality, while GPT-5 performs best in scenarios demanding rigorous factual consistency.

## 2 Related Work

**Reasoning in Large Language Models.** Reasoning has been studied as a way to improve LLMs (Xu et al., 2025). Most approaches focus on explicit reasoning chains that help models think step by step (Wei et al., 2022; Yao et al., 2023; Khot et al., 2023). Prior research has mainly targeted structured tasks with clear intermediate steps, such as solving math problems, performing symbolic rea-

soning, or answering questions (Wang et al., 2022; Diao et al., 2024). These methods help models break down complex problems, reason through each step, and verify their own conclusions. However, reasoning for open-ended text generation, especially summarization, remains much less explored (Li et al., 2024a).

**Reasoning in Summarization.** Recent work has begun to shift reasoning methods toward abstractive summarization. Some techniques focus on entity- or structure-based reasoning, like Chain-of-Density (Adams et al., 2023) and Hierarchical Decomposition (Chang et al., 2023). Others use reflective reasoning, such as Iterative Refinement, which allows models to review and improve their outputs (Zhang et al., 2023b; Yun et al., 2025; Zhang et al., 2025a). However, most earlier studies test only a few reasoning methods and rely on a narrow range of datasets.

Building on this foundation, our work conducts the first systematic evaluation of a range of reasoning methods. We include both previously studied methods, CoT (Wei et al., 2022; Adams et al., 2023; Zhang et al., 2024a), IR (Zhang et al., 2023b), E2A (Zhang et al., 2023a) and Deco (Chang et al., 2023; Chen et al., 2025), as well as newly adapted ones such as Cite (Ou and Lapata, 2025), Plan (Sun et al., 2024), QAG (Sinha, 2025) and SC (Wang et al., 2022; Chen et al., 2023).

## 3 Method

Reasoning can be either *implicit and in-model* (e.g., in LRMs) or *explicit and prompt-induced*, which employs intermediate reasoning cues. As shown in Figure 1, we categorize explicit reasoning methods into three schemes: (1) *augmentation-based reasoning*; (2) *organization-based reasoning*; and (3) *reflective reasoning*. We describe each scheme in detail below.

### 3.1 Problem Formulation

Abstractive summarization is formulated as generating a summary  $s$  from an input document  $d$  using the language model  $\mathcal{M}$ . The decoding objective can be written as:

$$\hat{s} = \arg \max_s p_{\mathcal{M}}(s | d), \quad (1)$$

where  $\hat{s}$  denotes the generated summary. We represent the summarization process as:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d). \quad (2)$$

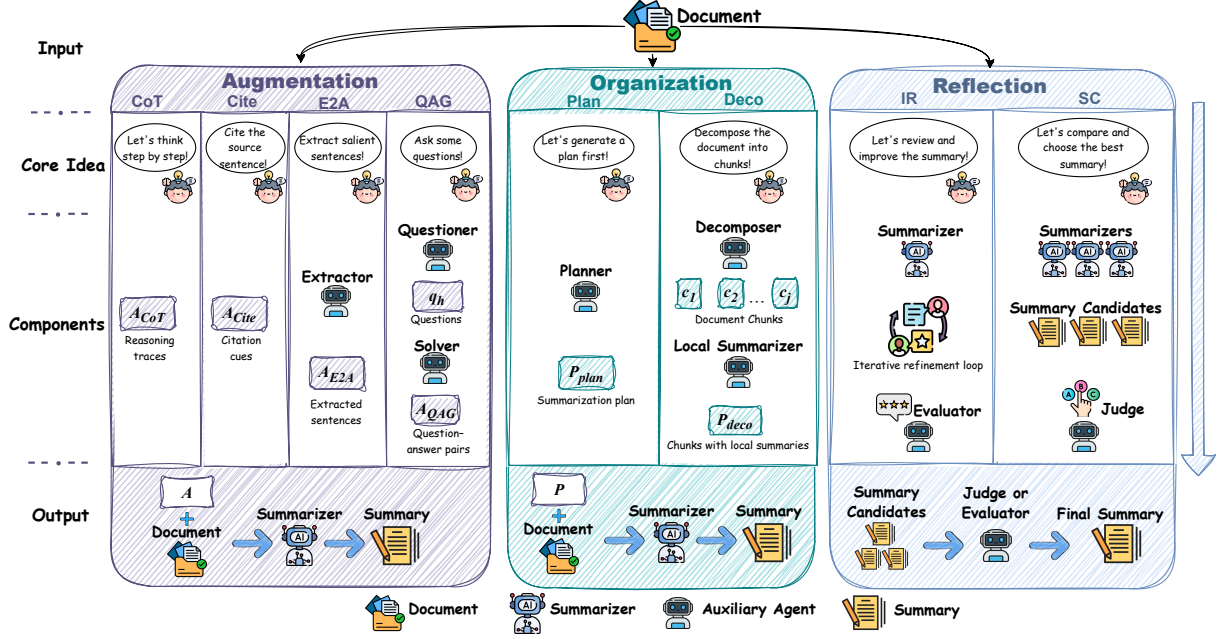


Figure 1: Overview of the 8 reasoning prompting strategies within our framework. We categorize them into three schemes: *Augmentation* expands the input context with auxiliary information; *Organization* structures the generation process via planning; and *Reflection* refines the output through self-evaluation and selection.

In the *Vanilla* setting, the model receives a brief instruction along with the document  $d$  and generates the summary in a single pass, i.e.,  $\hat{s} = \mathcal{M}_{\text{sum}}(d)$ . *LRMs* (e.g., O1, O3, GPT-5) use the same instruction, but rely on implicit, in-model reasoning.

### 3.2 Augmentation-based Reasoning (CoT, E2A, QAG, Cite).

In this scheme, the model is guided by the *augmentation* ( $A$ ) that expands the input with intermediate representations, such as stepwise reasoning cues, extracted evidence, or question–answer pairs. We define this generation process as:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d; A), \quad (3)$$

where  $A$  can be obtained by the following methods:

**CoT (Chain-of-Thought).** We adopt the CoT reasoning pattern to summarization (Wei et al., 2022; Adams et al., 2023; Zhang et al., 2024a). The augmentation  $A_{\text{CoT}}$  provides stepwise reasoning cues, directly embedded in the input prompt. The model generates the summary in a single pass:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d; A_{\text{CoT}}). \quad (4)$$

**E2A (Extract-to-Abstract).** We employ a two-stage extract–abstract pipeline (Dou et al., 2021; Zhang et al., 2023a). First, the model extracts  $K$

salient sentences from the document to form the augmentation  $A_{\text{E2A}}$ :

$$A_{\text{E2A}} = \{e_k\}_{k=1}^K = \mathcal{M}_{\text{extract}}(d), \quad (5)$$

where  $e_k$  denotes each extracted sentence. Next, the model generates the final summary with this additional context:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d; A_{\text{E2A}}). \quad (6)$$

**QAG (Question–Answer Guided).** Prior works (Durmus et al., 2020; Liu et al., 2024; Sinha, 2025) have shown that expert-designed questions can help models improve summarization faithfulness. We design a three-stage procedure that automatically generates questions. First, the model generates  $H$  document-specific questions:

$$\{q_h\}_{h=1}^H = \mathcal{M}_{\text{question}}(d). \quad (7)$$

Next, a new model answers each question:

$$a_h = \mathcal{M}_{\text{answer}}(d, q_h). \quad (8)$$

Finally, the summary is generated conditioned on the question–answer pairs, which form the augmentation  $A_{\text{QAG}}$  to generate the final summary:

$$A_{\text{QAG}} = \{(q_h, a_h)\}_{h=1}^H, \quad \hat{s} = \mathcal{M}_{\text{sum}}(d; A_{\text{QAG}}). \quad (9)$$

The number of questions  $H$  is decided by the question model according to the document length.

**Cite (Cited Summarization).** Inspired by cited summarization (Ou and Lapata, 2025), we propose a single-stage, retrieval-augmented setup. The augmentation  $A_{\text{Cite}}$  is a set of  $N$  citation units from the document:  $A_{\text{Cite}} = \{c_i\}_{i=1}^N$ . Each citation unit  $c_i$  serves as a referenceable sentence. The model is prompted to produce both the summary  $\hat{s}$  and a set of evidence alignments  $\mathcal{R}$ :

$$(\hat{s}, \mathcal{R}) = \mathcal{M}_{\text{cite}}(d; A_{\text{Cite}}). \quad (10)$$

Here,  $\mathcal{R}$  links each summary sentence to its supporting citations in  $A_{\text{Cite}}$ .

### 3.3 Organization-based Reasoning (Deco, Plan)

This reasoning scheme structures the summarization process through organization. We denote the organizational guidance prompt as  $P$ . The model generates the summary as:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d; P), \quad (11)$$

where  $P$  encodes the inferred structure or an explicit plan outlining the summary’s organization.

**Deco (Decomposition).** We implement a hierarchical decomposition strategy designed for long-document summarization (Chang et al., 2023). The process has three stages. First, the model segments the document  $d$  into  $J$  coherent chunks:

$$\{c_j\}_{j=1}^J = \mathcal{M}_{\text{seg}}(d). \quad (12)$$

Second, the model generates a local summary  $u_j$  for each chunk:

$$u_j = \mathcal{M}_{\text{sum}}(c_j). \quad (13)$$

Finally, the local summaries are combined into a structural plan  $P_{\text{Deco}}$ , which guides the final summarization:

$$P_{\text{Deco}} = \{u_j\}_{j=1}^J, \quad \hat{s} = \mathcal{M}_{\text{sum}}(d; P_{\text{Deco}}). \quad (14)$$

This method captures local information in each chunk and uses it to improve the global summary.

**Plan (Plan-then-Write).** Following previous work (Sun et al., 2024), we propose a plan then write strategy by introducing a two-stage structural framework. First, the model generates a compact plan  $P_{\text{Plan}}$ :

$$P_{\text{Plan}} = \mathcal{M}_{\text{plan}}(d), \quad (15)$$

where  $P_{\text{Plan}}$  includes elements such as  $\{\text{domain}, \text{goal}, \text{style}, \text{salient\_info}\}$ , which outlines the summary’s focus and structure. Second, the model generates the final summary conditioned on this plan:

$$\hat{s} = \mathcal{M}_{\text{sum}}(d; P_{\text{Plan}}). \quad (16)$$

This method mimics the human outlining process to organize key elements before abstraction.

### 3.4 Reflective Reasoning (IR, SC)

This reasoning scheme refines summaries through iterative reflection or self-evaluation. It allows the model to assess its own outputs and improve them based on the feedback. We denote reflective feedback or evaluation cues as  $R$ . The process typically involves three steps: generating one or more candidate summaries, evaluating them using  $R$ , and then revising or selecting the final summary  $\hat{s}$ .

**IR (Iterative Refine).** We adopt the evaluate then revise framework from prior work (Zhang et al., 2023b; Lin et al., 2025). The model iteratively improves its summary through self-evaluation and targeted revision, with early stopping to prevent over-refinement. The process starts with an initial draft:

$$s^{(0)} = \mathcal{M}_{\text{sum}}(d). \quad (17)$$

For each iteration  $t \in \{1, \dots, T\}$ , the model performs a two-step reflective loop. First, an evaluator module generates structured feedback  $R_{\text{IR}}^{(t)}$ :

$$R_{\text{IR}}^{(t)} = \mathcal{M}_{\text{evaluate}}(d, s^{(t-1)}). \quad (18)$$

Second, a reviser module incorporates this feedback to produce an improved summary:

$$s^{(t)} = \mathcal{M}_{\text{revise}}(d, s^{(t-1)}, R_{\text{IR}}^{(t)}). \quad (19)$$

The loop terminates when  $R_{\text{IR}}^{(t)}$  signals  $\langle \text{STOP} \rangle$ , or after  $T$  iterations. The final summary  $\hat{s} = s^{(T)}$  is taken from the last iteration.

**SC (Self-Consistency).** Building on self-consistency from majority voting (Wang et al., 2022; Chen et al., 2023; Li et al., 2024b), we propose a rubric-based selection framework for abstractive summarization. First, the model samples  $N$  candidate summaries from its conditional distribution, using a non-zero decoding temperature:

$$s^{(q)} \sim p_{\mathcal{M}}(s | d), \quad (20)$$

Dataset	Domain	Format	Statistics				
			Doc.Token	Sum.Token	Density (%)	Compression	Coverage (%)
CNN/DM (Hermann et al., 2015)	News	SDS	717	57	44.18	13.51	53.50
SAMSum (Gliwa et al., 2019)	Dialogue	SDS	101	20	16.68	5.18	31.62
Reddit (Kim et al., 2019)	Social Media	SDS	461	31	13.33	18.90	35.54
WikiHow (Koupae and Ganesan, 2018)	Knowledge Base	SDS	608	65	9.68	10.98	36.29
ArXiv (Cohan et al., 2018)	Scientific Paper	SDS	6,514	168	16.15	42.24	58.67
Multi-News (Fabbri et al., 2019)	News	MDS	2,098	272	0.23	7.69	45.67
BookSum (Kryściński et al., 2021)	Novel	LNS	3,524	290	1.67	12.49	31.12
SciGen (Moosavi et al., 2021)	Scientific Paper	TTS	217	123	77.83	2.32	14.85

Table 1: Statistics of the included summarization datasets. The ‘‘Format’’ column indicates the task type: SDS = Single-Document Summarization, MDS = Multi-Document Summarization, LNS = Long-form Narrative Summarization, and TTS = Table-to-Text Summarization.

where  $q \in \{1, \dots, N\}$ . Next, an LLM-based judge scores each candidate  $s^{(q)}$  against the document  $d$  using a rubric that evaluates criteria such as faithfulness, coverage, and fluency:

$$\text{score}_q = \mathcal{M}_{\text{judge}}(d, s^{(q)}). \quad (21)$$

Finally, the final summary is selected as:

$$q^* = \arg \max_q \text{score}_q. \quad (22)$$

where the candidate  $q^*$  has the highest score, and the final summary  $\hat{s} = s^{(q^*)}$  is obtained.

## 4 Experimental Setup

### 4.1 Datasets

As shown in Table 1, we use an evaluation suite of eight representative datasets covering multiple domains (Chu et al., 2025), which are categorized into three groups based on their input structure and task type. Full details on data pre-processing are provided in Appendix A.

*Short-Form Documents:* (1) **CNN/DM** (Hermann et al., 2015) contains concise news articles paired with human-written highlights. (2) **SAM-Sum** (Gliwa et al., 2019) includes messenger-style dialogues and short abstractive summaries. (3) **Reddit** (Kim et al., 2019) features informal, user-generated posts from the TIFU subreddit. (4) **WikiHow** (Koupae and Ganesan, 2018) provides instructional articles with step-wise summaries. These datasets show the single-document summarization (SDS) setting, with short inputs (100–700 tokens) and moderate compression ratios.

*Long-Form Documents:* (5) **ArXiv** (Cohan et al., 2018) consists of long scientific papers and their abstracts. (6) **Multi-News** (Fabbri et al., 2019) merges multiple related news articles into a single summary, requiring cross-document reasoning. (7) **BookSum** (Kryściński et al., 2021) includes chapters and narratives summarized by humans. These

datasets evaluate multi-document (MDS) and long-form narrative summarization (LNS) capabilities, with inputs up to around 6,500 tokens.

*Table-to-Text:* (8) **SciGen** (Moosavi et al., 2021) pairs scientific tables with descriptive textual summaries, testing structured-to-text generation (TTS).

### 4.2 Implementation Details

**Models.** Our experimental evaluation uses the GPT family via Azure AI API<sup>2</sup>. We utilize GPT-4.1 (1,000K context window) as our main model. We compare its performance against three LRMs: o1 (200K context), o3 (200K context), and GPT-5 (400K context).

**Decoding Parameters.** All configurations run once for each document with a fixed random seed for consistency. We use temperature=0 and keep top\_p, top\_k, and other sampling settings at their default values. For GPT-4.1, we set max\_output\_tokens=1,000. For the LRMs, we use think\_ability=medium and max\_completion\_tokens=10,000 because these models include internal ‘‘thinking’’ tokens in their completion budget.

### 4.3 Evaluation Metrics

We evaluate summaries using automatic metrics grouped by assessment focus: reference similarity, factual faithfulness, and LLM as a judge. **Reference similarity:** We report ROUGE (average of ROUGE-1/2/L) (Lin, 2004), which measures n-gram overlap and longest common subsequence between system and reference summaries. For semantic similarity, we use BERTScore F1 (Zhang et al., 2020), which computes contextual embedding similarity between the two summaries. **Factual faithfulness:** We adopt SummaC (Laban et al., 2022), which estimates sentence-level consistency via natural language inference, and AlignScore (Zha et al.,

<sup>2</sup><https://ai.azure.com/>

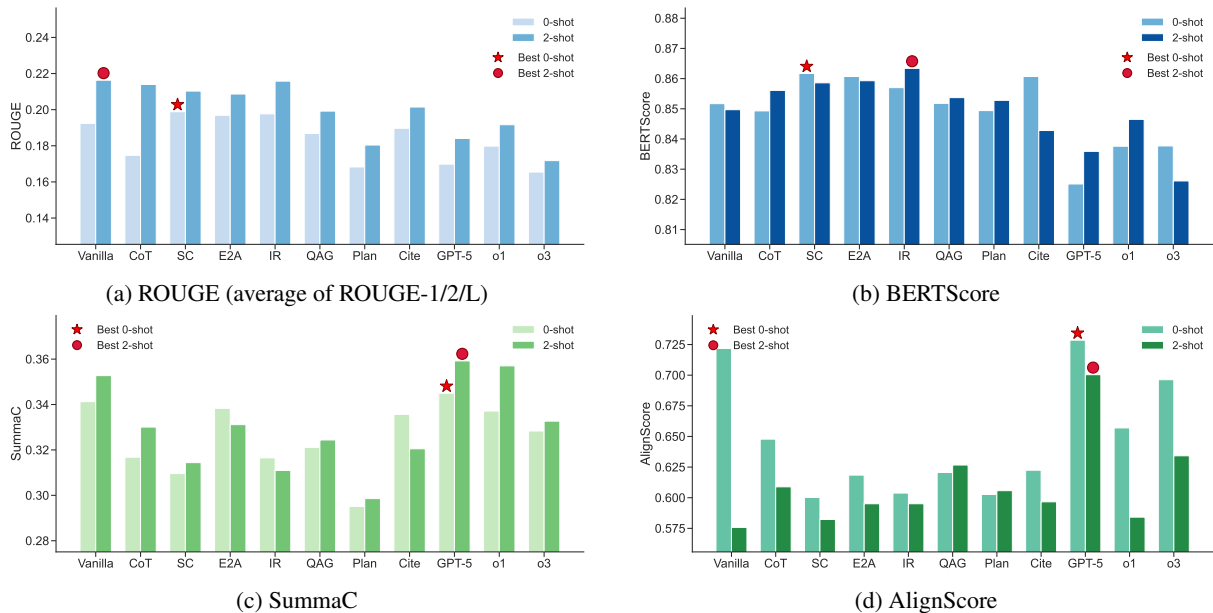


Figure 2: Summarization performance across 8 datasets. Metrics include reference-based quality (ROUGE, BERTScore) and factual faithfulness (SummaC, AlignScore). Red markers highlight the top methods.

2023), a finetuned metric measuring information alignment between paired texts. **LLM as a judge:** We use G-Eval (Liu et al., 2023), a GPT-based evaluation framework that provides human-aligned scores across three dimensions: Completeness, Conciseness, and Faithfulness. G-Eval offers a comprehensive view of overall summary quality beyond traditional overlap-based metrics. Additionally, we conduct a **human evaluation** using the same rubric as G-Eval, which allows a direct comparison between LLM-based judgments and human assessments.

## 5 Results and Analysis

### 5.1 RQ1: Does reasoning actually improve summarization performance compared to standard prompting?

**I. Reasoning helps with summarization only in specific configurations.** As shown in Figure 2, we find the impact of reasoning is highly context-dependent. For reference similarity (ROUGE and BERTScore), explicit reasoning methods such as SC and IR achieve the highest scores in the 0-shot setting. However, this advantage disappears in the 2-shot setting, where the Vanilla baseline outperforms other methods. For factual faithfulness (SummaC and AlignScore), we observe a different pattern: the LRM GPT-5 delivers the best results across both settings, suggesting benefits from implicit, model-internal reasoning. Meanwhile,

Vanilla remains a stable baseline. Overall, reasoning does not always make summarization better. Its benefits depend on the setup. The practical guide on selecting a strategy for each domain is provided in table 15.

**II. Reasoning-based methods show a quality-faithfulness trade-off.** Figure 3 illustrates the relationship between summary quality (BERTScore) and factual faithfulness (AlignScore). We observe a negative correlation, with a correlation coefficient of  $-0.685$  and a p-value of  $0.014$ , indicating that the trade-off is statistically significant. Most explicit reasoning methods, such as SC, IR, QAG, Plan, and CoT, group in the high-BERTScore range of about 85 to 86. However, they have a low-to-mid AlignScore of around 60 to 62. In contrast, the LRMs move towards a lower BERTScore (approximately 82 to 84) but achieve a higher AlignScore. The GPT-5 model sits at the upper-left extreme. Interestingly, the Vanilla baseline is an outlier above the trend line. It achieves both a strong BERTScore and an above-average AlignScore.

### 5.2 RQ2: When do reasoning strategies excel across different datasets and context?

**III. SC and IR achieve the best summary quality across most settings.** As discussed in RQ1, explicit reasoning methods perform best on reference-based metrics. Figures 2a and 2b show that SC




Figure 3: BERTScore vs. AlignScore trade-off analysis. The linear fit ( $r = -0.685$ ,  $p = 0.014$ ) reveals a significant negative correlation between reference-based quality and factual faithfulness.

and *IR* consistently achieve the highest ROUGE and BERTScore in the 0-shot setting. This overall strength reflects complementary, dataset-specific advantages: *SC* leads in short-form summarization (e.g., CNN/DM, SAMSum), while *IR* performs best on long-form datasets (ArXiv, Multi-News). On table-to-text (SciGen), the differences are smaller, with *SC*, *IR*, and *E2A* delivering comparable results. As noted in RQ1, this advantage diminishes in the 2-shot setting, where *Vanilla* becomes highly competitive. In contrast, LRMs (*o1*, *o3*, *GPT-5*) deliver lower scores on reference-based metrics.

#### IV. LRMs achieve the highest factual faithfulness, especially on long-form summarization.

As shown in Figures 2c and 2d, in both 0-shot and 2-shot settings, *GPT-5* achieves the highest average SummaC and AlignScore. Both *o1* and *o3* also do better than explicit reasoning methods. This advantage is clearest in long-form datasets, while *Vanilla* stays competitive only in short-form tasks.

In table-to-text summarization (SciGen), AlignScore approaches 98% to 100% across all methods. We infer that factual alignment is easily met because of the structured inputs. Explicit reasoning methods, like *SC* and *IR*, lose faithfulness when provided with 2-shot examples. In contrast, LRMs either stay the same or show slight improvement. Overall, LRMs with implicit reasoning produce the most faithful summaries.

**V. LLM as a judge overestimates faithfulness compared to humans.** Following Song et al. (2024), we use GEval to evaluate three metrics: Completeness, Conciseness, and Faithfulness. Overall illustrates the average of the above met-

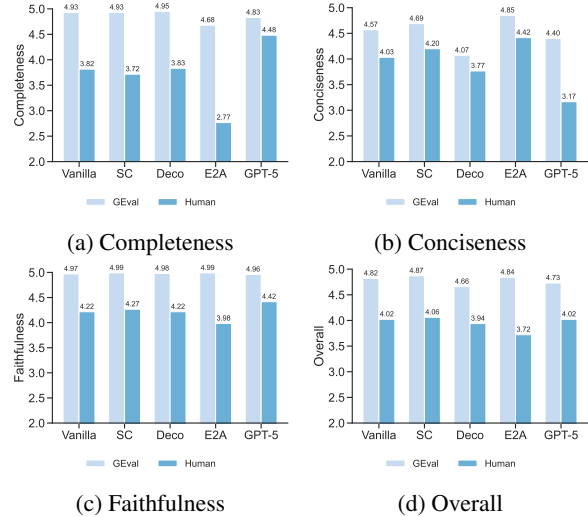


Figure 4: Comparison between human evaluation and GEval scores across five systems (Deco, E2A, GPT-5, SC, Vanilla). Each subplot reports one metric.




Figure 5: Higher abstractive-ness correlates negatively with both reference-based quality (ROUGE) and factual faithfulness (SummaC).

rics. As shown in Figure 4, GEval assigns nearly perfect faithfulness scores (4.96 to 4.99) to all systems. In contrast, human ratings are lower and more varied, from 3.98 to 4.42 (See Appendix I for human evaluation setup). GEval considers GPT-5 and E2A nearly the same and even slightly favors E2A (4.96 vs. 4.99), whereas humans score GPT-5 much higher in faithfulness (4.42 vs. 3.98). Human evaluations also reveal a clear trade-off: GPT-5 is preferred when faithfulness and completeness are most important (4.42, 4.48) but ranks lowest in conciseness (3.17), while E2A is the most concise (4.42) but loses important content and accuracy (2.77 completeness, 3.98 faithfulness). GEval captures part of this trend by favoring SC and penalizing Deco and GPT-5 for their lack of conciseness. Our results suggest that faithfulness evaluation in summarization cannot rely solely on the LLM.

#### 5.3 RQ3: How do different reasoning methods impact summarization results?

<b>Document</b>	[...] If you are going to be eating the pieces raw, they can be a little longer, but <b>if you’re cooking with them, 2–3 inches is best</b> . Each chunk does not have to be the exact same length, just divide the vegetable evenly. The general rule for juliennes is that you’d be able to eat them in one bite, but you can cut them into longer chunks if you want. When cutting juliennes, be careful not to cut yourself as you slice. <b>Flip the slices</b> you’ve just cut so that they are piled on top of each other. Then cut them into <b>sticks that are <math>\frac{1}{8}</math> in (0.3 cm) to <math>\frac{1}{16}</math> in (0.15 cm) thick</b> . The goal with juliennes is to get them all to a fairly uniform size.
<b>Human Ref.</b>	[...] Cut rounded vegetables into a rectangle. Chop the vegetable into <b>chunks that are 2–3 inches long</b> . Slice the vegetable down the length. <b>Stack 2–3 slices</b> and cut again.
<b>Vanilla</b>	[...] Slice vegetables into rectangular chunks, discarding or saving scraps. For julienne cuts, divide vegetables evenly; <b>raw pieces can be longer, but cooked ones should be 2–3 inches</b> . Cut slices into <b>sticks <math>\frac{1}{8}</math> in (0.3 cm) to <math>\frac{1}{16}</math> in (0.15 cm) thick, stacking slices</b> for uniformity, and always be careful not to cut yourself. Aim for uniform, bite-sized pieces.
<b>SC</b>	[...] Cut off stems and bottoms as needed, slice vegetables into rectangular chunks, and either discard or use scraps. For raw pieces, longer cuts are acceptable; <b>for cooking, 2–3 inch lengths are best</b> . Aim for uniform, bite-sized <b>sticks about <math>\frac{1}{8}</math> to <math>\frac{1}{16}</math> inch thick, stacking slices</b> before cutting, and always be careful not to cut yourself.
<b>Deco</b>	[...] Prepare vegetables by peeling them if desired, removing unwanted parts like stems, bottoms, or potato eyes, and cutting them into manageable rectangular chunks based on your intended use and preference. For julienning, <b>stack the slices</b> and cut them into <b>uniform sticks about <math>\frac{1}{8}</math> to <math>\frac{1}{16}</math> inch thick and 2–3 inches long</b> , aiming for consistency in size to ensure even cooking and presentation.
<b>E2A</b>	[...] Whether or not you peel vegetables depends on how you plan to eat them. Slice off one side at a time to create rectangular chunks, aiming for pieces that are generally bite-sized, though longer chunks are acceptable if preferred. <b>The key is to cut the juliennes to a fairly uniform size.</b>
<b>GPT-5</b>	[...] <ul style="list-style-type: none"> <li>- Square the vegetable by slicing off sides to form a rectangular chunk; discard or use trimmings. This is especially helpful for potatoes, zucchini, and squash.</li> <li>- Cut the chunk into even slabs, <b>stack them</b>, then slice into <b>uniform sticks about <math>\frac{1}{8}</math>–<math>\frac{1}{16}</math> inch (<math>\approx 0.3</math>–<math>0.15</math> cm) thick</b>. Aim for one-bite, fairly uniform pieces; be careful not to cut yourself.</li> <li>- Length: <b>for cooking, 2–3 inches is best; raw pieces can be a little longer</b>. Divide the vegetable into evenly sized chunks.</li> </ul>

Table 2: Case study from WikiHow comparing system summaries with the source document and a human reference. **Blue text** marks source-faithful procedural details, especially the conditional length constraint and the stacking/cutting steps. **Red text** marks errors, including over-generalization in **Deco** and critical omission in **E2A**.

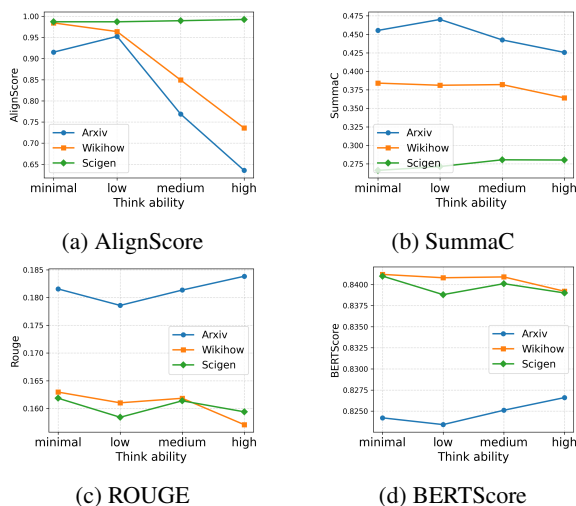


Figure 6: Effect of GPT-5 think ability (from minimal to high) on 4 evaluation metrics across 3 datasets.

**VI. Summary abstractiveness is negatively correlated with reference-based quality and factual faithfulness.** As outlined in Section 3, we group systems into five reasoning paradigms and

measure their average abstractiveness (Yuan and Zhang, 2025). In Figure 5, *Vanilla* and *Reflective* are the least abstract and achieve the highest ROUGE scores. *Augmentation* methods fall in the middle on both metrics. *Organization* methods are the most abstract and perform the worst on both metrics. LRMs break this trend: they are relatively extractive but trade lower ROUGE for higher SummaC. This suggests that abstractiveness interacts with model family and internal reasoning style.

We further conduct a qualitative case study to investigate specific model behaviors. As shown in Table 2, *Deco*’s chunk-based aggregation loses global conditional context, causing it to over-generalize a conditional rule and apply the “2–3 inches” cooking length to all julienne sticks, while *E2A*’s extraction bottleneck excludes the final cutting step, resulting in an incomplete procedure. By comparison, *Vanilla* and GPT-5 largely preserve the original conditional structure. These examples suggest that summarization requires accurate compression in-

stead of creative reasoning. Reasoning pipelines that fragment or filter context risk driving models to fill non-existent logical gaps or omit important information. The complete case study table is provided in Appendix H.

### VII. Risk of Over-Thinking: Increased reasoning reduces faithfulness without improving quality.

We analyze the effect of a hyperparameter named “think ability”, which controls the depth of internal reasoning within *GPT-5*. We select one representative dataset from each of the three groups: short-form (*WikiHow*), long-form (*ArXiv*), and table-to-text (*SciGen*). As shown in Figure 6, factual faithfulness consistently declines as think ability increases: *AlignScore* shows the steepest drop, while *SummaC* exhibits a similar but milder trend. In contrast, reference-based quality metrics (*ROUGE* and *BERTScore*) remain stable. This observed drop suggests that the LRM can “over-think” the source instead of maintaining faithful compression. Overall, too much internal reasoning poses a clear risk. It undermines factual grounding without improving summary quality.

### VIII. Controlled validation beyond GPT: More internal reasoning does not always help summarization.

Our main benchmark focuses on the GPT family. To test whether this pattern also appears beyond GPT, we conduct an additional controlled study on *Gemini 2.5 Flash*. Specifically, we vary its internal *thinking budget* (tb) while keeping the summarization prompt fixed to *Vanilla*. We use the same evaluation protocol and metrics as in Section 4. Table 3 reports results on three representative datasets: short-form *CNN/DM*, long-form *ArXiv*, and table-to-text *SciGen*.

We find that increasing the thinking budget does not consistently improve summarization performance. *BERTScore* stays nearly unchanged across different budgets. For *ROUGE*, the clearest pattern appears on *ArXiv*, where performance drops from 0.2284 at  $tb=0$  to 0.1954 at  $tb=1024$ . On *CNN/DM* and *SciGen*, the best *ROUGE* is obtained at the middle setting  $tb=256$ , rather than at the largest budget. The faithfulness metrics show a similarly mixed pattern across datasets. Overall, these results support our main conclusion that internal reasoning is highly context-dependent in summarization, and that simply allocating more reasoning

budget does not reliably lead to better summaries.

(a) Reference similarity.				
Metric	tb	CNN/DM	ArXiv	SciGen
ROUGE	0	0.2374	<b>0.2284</b>	0.1737
	256	<b>0.2423</b>	0.2278	<b>0.1769</b>
	1024	0.2366	0.1954	0.1716
BERTScore	0	0.8713	0.8399	0.8467
	256	0.8712	<b>0.8410</b>	0.8489
	1024	<b>0.8721</b>	0.8402	<b>0.8501</b>
(b) Factual faithfulness.				
Metric	tb	CNN/DM	ArXiv	SciGen
SummaC	0	0.3224	0.4831	<b>0.3074</b>
	256	<b>0.3275</b>	0.4756	0.2941
	1024	0.3180	<b>0.5317</b>	0.2957
AlignScore	0	0.4916	0.8779	0.9886
	256	<b>0.6156</b>	<b>0.9761</b>	<b>0.9940</b>
	1024	0.5280	0.7472	0.9871

Table 3: Gemini 2.5 Flash with different thinking budgets (tb) under the *Vanilla* prompt. Bold marks the best tb for each dataset and metric.

## 6 Conclusion

This paper presents the first large-scale comparative study of reasoning paradigms in LLM-based summarization. We evaluate eight prompting strategies and three LRMs across eight diverse datasets. Our findings challenge the assumption that reasoning always improves summarization. We uncover a trade-off: explicit prompting enhances quality but reduces faithfulness, whereas implicit reasoning in LRMs generally shows the opposite. Crucially, we find that excessive internal reasoning triggers creative gap-filling that harms factual consistency. Overall, our results suggest that reasoning in summarization is highly context-dependent: stronger reasoning does not always lead to better summaries, and more complex reasoning can reduce faithfulness without improving final summary quality.

## Acknowledgement

This work was supported in part by National Science Foundation awards 2149133, 2004014, and 2138259, as well as by computing resources provided through the NSF ACCESS projects CIS250940 and CIS24061.

## Limitations

In our evaluation, we used 100-instance subsets per dataset rather than the full test sets. This choice was made to support a broad comparison across eight diverse datasets, eight reasoning paradigms, and three LRMs. Previous research has similarly adopted sampled evaluation for large-scale LLM studies (Goyal et al., 2023; Zhang et al., 2024c). While this allows for extensive cross-domain analysis, future work may further verify our findings on full test splits.

Our study mainly focuses on English summarization. Although the selected datasets span diverse domains, including news, dialogue, social media, structured table-to-text (Guo et al., 2025), and scientific papers, extending our analysis to multilingual settings or specialized vertical domains (e.g., legal (Zhang et al., 2025b), biomedical (Yuan et al., 2024; Ou et al., 2025), or clinical (Shen et al., 2026)) remains an important direction for future research.

We conduct our experiments using the GPT model family (GPT-4.1, o1, o3, GPT-5). This ensures a consistent baseline capability, allowing us to isolate the impact of different reasoning strategies without confounding factors from varying model architectures or training distributions. However, relying on closed-source APIs restricts fine-grained inspection of internal reasoning processes (e.g., attention weights). Consequently, our analysis of "internal reasoning in the LRM" relies on observing output behaviors rather than internal mechanisms (Lin et al., 2026; Zhu et al., 2026). Verifying these patterns on open-weight models (Li et al., 2025) would be a valuable extension.

Regarding evaluation, we employ a combination of human and automated metrics. While human evaluation serves to validate the trends observed in automated scores, scaling manual annotation to the full extent of such a large benchmark remains challenging. Furthermore, we do not conduct a fine-grained latency analysis. Some resource-heavy methods (e.g., QAG, SC, IR) need multiple generations, which may inherently trade inference cost for quality. Optimizing this efficiency trade-off (Nie et al., 2026; Liu et al., 2026) is a practical consideration for deployment that warrants future engineering-focused studies.

## References

- Griffin Adams, Alexander R Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: Gpt-4 summarization with chain of density prompting. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing, 4th New Frontier Summarization Workshop*, page 68.
- Yapei Chang, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2023. Boookscore: A systematic exploration of book-length summarization in the era of llms. In *International Conference on Learning Representations*.
- Xi Chen, Wei Xue, and Yike Guo. 2026a. [Actormind: Emulating human actor reasoning for speech role-playing](#). *Preprint*, arXiv:2604.11103.
- Xiaoyong Chen, Zhiqiang Chen, and Shi Cheng. 2025. Cothssum: Structured long-document summarization via chain-of-thought reasoning and hierarchical segmentation. *Journal of King Saud University Computer and Information Sciences*, 37(4):40.
- Xinrong Chen, Xu Chu, Yingmin Qiu, Hengyuan Zhang, Jing Xiong, Shiyu Tang, Shuai Liu, Shaokang Yang, Cheng Yang, Hayden Kwok-Hay So, et al. 2026b. Residual decoding: Mitigating hallucinations in large vision-language models via history-aware residual guidance. *arXiv preprint arXiv:2602.01047*.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2023. Universal self-consistency for large language model generation. *arXiv preprint arXiv:2311.17311*.
- Xu Chu, Zhijie Tan, Hanlin Xue, Guanyu Wang, Tong Mo, and Weiping Li. 2025. Domain1s: Guiding llm reasoning for explainable answers in high-stakes domains. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 3275–3293.
- 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.
- Shizhe Diao, Pengcheng Wang, Yong Lin, Rui Pan, Xiang Liu, and Tong Zhang. 2024. Active prompting with chain-of-thought for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1330–1350.
- Haonan Dong, Kehan Jiang, Haoran Ye, Wenhao Zhu, Zhaolu Kang, and Guojie Song. 2026. [Neureasoner: Towards explainable, controllable, and uni-](#)

- fied reasoning via mixture-of-neurons. *Preprint*, arXiv:2604.02972.
- Haonan Dong, Wenhao Zhu, Guojie Song, and Liang Wang. 2025. [AuroRA: Breaking low-rank bottleneck of loRA with nonlinear mapping](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- 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.
- Esin Durmus, He He, and Mona Diab. 2020. Feqa: A question answering evaluation framework for faithfulness assessment in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5055–5070.
- Alexander R. Fabbri, Greg Durrett, Justin Farcas, and Vivian Ng. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2023. [News summarization and evaluation in the era of gpt-3](#). *Preprint*, arXiv:2209.12356.
- Yu Guo, Dong Jin, Shenghao Ye, Shuangwu Chen, Jianyang Jianyang, and Xiaobin Tan. 2025. Sqlforge: Synthesizing reliable and diverse data to enhance text-to-sql reasoning in llms. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8441–8452.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 1693–1701. Curran Associates, Inc.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065.
- Kehan Jiang, Haonan Dong, Zhaolu Kang, Zhengzhou Zhu, and Guojie Song. 2026. [Foe: Forest of errors makes the first solution the best in large reasoning models](#). *Preprint*, arXiv:2604.02967.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2023. [Decomposed prompting: A modular approach for solving complex tasks](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. Abstractive summarization of Reddit posts with multi-level memory networks. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2519–2530, Minneapolis, Minnesota. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Moin Koupae and L.V. Ganesan. 2018. WikiHow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*.
- Wojciech Kryściński, Roman Paulus, Caiming Xiong, and Richard Socher. 2021. BookSum: A collection of datasets for long-form narrative summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8969–8981, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Philippe Laban, Tobias Hsi, and Arman Cohan. 2022. SummaC: Re-visiting NLI-based models for inconsistency detection in summarization. In *Transactions of the Association for Computational Linguistics*, volume 10, pages 163–180, Cambridge, MA. MIT Press.
- Mengdi Li, Jiaye Lin, Xufeng Zhao, Wenhao Lu, Peilin Zhao, Stefan Wermter, and Di Wang. 2025. Curriculum-rlaif: Curriculum alignment with reinforcement learning from ai feedback. *arXiv preprint arXiv:2505.20075*.
- Miao Li, Jey Han Lau, and Eduard Hovy. 2024a. [A sentiment consolidation framework for meta-review generation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10158–10177, Bangkok, Thailand. Association for Computational Linguistics.
- Taiji Li, Zhi Li, and Yin Zhang. 2024b. Improving faithfulness of large language models in summarization via sliding generation and self-consistency. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8804–8817.
- Zixu Li, Yupeng Hu, Zhiwei Chen, Qinlei Huang, Guozhi Qiu, Zhiheng Fu, and Meng Liu. 2026a. [Retrack: Evidence-driven dual-stream directional](#)

- anchor calibration network for composed video retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 23373–23381.
- Zixu Li, Yupeng Hu, Zhiwei Chen, Shiqi Zhang, Qinlei Huang, Zhiheng Fu, and Yinwei Wei. 2026b. Habit: Chrono-synergia robust progressive learning framework for composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 6762–6770.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Honglin Lin, Zheng Liu, Yun Zhu, Chonghan Qin, Juekai Lin, Xiaoran Shang, Conghui He, Wentao Zhang, and Lijun Wu. 2026. Mmfinereason: Closing the multimodal reasoning gap via open data-centric methods. *arXiv preprint arXiv:2601.21821*.
- Jiaye Lin, Yifu Guo, Yuzhen Han, Sen Hu, Ziyi Ni, Licheng Wang, Mingguang Chen, Hongzhang Liu, Ronghao Chen, Yangfan He, et al. 2025. Se-agent: Self-evolution trajectory optimization in multi-step reasoning with llm-based agents. *arXiv preprint arXiv:2508.02085*.
- Guoming Ling, Zhongzhan Huang, Yupei Lin, Junxin Li, Shanshan Zhong, Hefeng Wu, and Liang Lin. 2026. [Neural chain-of-thought search: Searching the optimal reasoning path to enhance large language models](#). *Preprint*, arXiv:2601.11340.
- Junyuan Liu, Zhengyan Shi, and Aldo Lipani. 2024. Summequal: Summarization evaluation via question answering using large language models. In *Proceedings of the 2nd Workshop on Natural Language Reasoning and Structured Explanations (@ ACL 2024)*, pages 46–55.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, and James Glass. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Zheng Liu, Honglin Lin, Chonghan Qin, Xiaoyang Wang, Xin Gao, Yu Li, Mengzhang Cai, Yun Zhu, Zhanping Zhong, Qizhi Pei, et al. 2026. Chart-verse: Scaling chart reasoning via reliable programmatic synthesis from scratch. *arXiv preprint arXiv:2601.13606*.
- Nafise Sadat Moosavi, Andreas Ghasemi, Heidar Ghasemzade, Ranalli Kazemi, Horacio Saggion, Francesco Ruggiano, and James Henderson. 2021. SciGen: A dataset for reasoning-aware text generation from scientific tables. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6977–6985, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Shuaiyi Nie, Siyu Ding, Wenyuan Zhang, Linhao Yu, Tianmeng Yang, Yao Chen, Tingwen Liu, Weichong Yin, Yu Sun, and Hua Wu. 2026. Attnpo: Attention-guided process supervision for efficient reasoning. *arXiv preprint arXiv:2602.09953*.
- OpenAI, :, Aaron Jaech, Adam Kalai, and Adam Lerer et al. 2024a. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- OpenAI, Josh Achiam, Steven Adler, and Sandhini Agarwal et al. 2024b. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Justice Ou, Tinglin Huang, Yilun Zhao, Ziyang Yu, Peiqing Lu, and Rex Ying. 2025. [Experience retrieval-augmentation with electronic health records enables accurate discharge qa](#). *Preprint*, arXiv:2503.17933.
- Litu Ou and Mirella Lapata. 2025. Context-aware hierarchical merging for long document summarization. *arXiv preprint arXiv:2502.00977*.
- Yifei Shen, Yilun Zhao, Justice Ou, Tinglin Huang, and Arman Cohan. 2026. [Patient-similarity cohort reasoning in clinical text-to-SQL](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1367–1412, Rabat, Morocco. Association for Computational Linguistics.
- Neelabh Sinha. 2025. Qa-prompting: Improving summarization with large language models using question-answering. *arXiv preprint arXiv:2505.14347*.
- Hwanjun Song, Hang Su, Igor Shalyminov, Jason Cai, and Saab Mansour. 2024. Finesure: Fine-grained summarization evaluation using llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 906–922.
- Simeng Sun, Yang Liu, Shuohang Wang, Dan Iter, Chenguang Zhu, and Mohit Iyyer. 2024. Pearl: Prompting large language models to plan and execute actions over long documents. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 469–486.
- Zelin Tan, Hejia Geng, Xiaohang Yu, Mulei Zhang, Guancheng Wan, Yifan Zhou, Qiang He, Xiangyuan Xue, Heng Zhou, Yutao Fan, Zhongzhi Li, Zaibin Zhang, Guibin Zhang, Chen Zhang, Zhenfei Yin, Philip Torr, and Lei Bai. 2025. [Scaling behaviors of llm reinforcement learning post-training: An empirical study in mathematical reasoning](#). *Preprint*, arXiv:2509.25300.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, et al. 2025. Towards large reasoning models: A survey of reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Haohan Yuan, Sukhwa Hong, and Haopeng Zhang. 2025. Strucsum: Graph-structured reasoning for long document extractive summarization with llms. *arXiv preprint arXiv:2505.22950*.
- Haohan Yuan, Siu Cheung Hui, and Haopeng Zhang. 2024. A structure-aware generative model for biomedical event extraction. *arXiv preprint arXiv:2408.06583*.
- Haohan Yuan and Haopeng Zhang. 2025. Domainsum: A hierarchical benchmark for fine-grained domain shift in abstractive text summarization. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2219–2231.
- Taewon Yun, Jihwan Oh, Hyangsuk Min, Yuho Lee, Jihwan Bang, Jason Cai, and Hwanjun Song. 2025. [Refeed: Multi-dimensional summarization refinement with reflective reasoning on feedback](#). *Preprint*, arXiv:2503.21332.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. Alignscore: Evaluating factual consistency with a unified alignment function. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348.
- Haopeng Zhang. 2025. [Bridging multimodal and video summarization: A unified survey](#). In *Proceedings of The 5th New Frontiers in Summarization Workshop*, pages 157–171, Hybrid. Association for Computational Linguistics.
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023a. Extractive summarization via chatgpt for faithful summary generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3270–3278.
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023b. Summit: Iterative text summarization via chatgpt. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10644–10657.
- Haopeng Zhang, Philip S Yu, and Jiawei Zhang. 2025a. A systematic survey of text summarization: From statistical methods to large language models. *ACM Computing Surveys*, 57(11):1–41.
- Kepu Zhang, Weijie Yu, Sunhao Dai, and Jun Xu. 2025b. Citalaw: Enhancing llm with citations in legal domain. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11183–11196.
- Longyin Zhang, Bowei Zou, Jacintha Yi, and Aiti Aw. 2024a. Comprehensive abstractive comment summarization with dynamic clustering and chain of thought. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 2884–2896.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Yongxin Tong, Hongwei Zheng, and Zhiming Zheng. 2026a. Less is more: Compact clue selection for efficient retrieval-augmented generation reasoning. In *Proceedings of the ACM Web Conference 2026*, pages 1971–1982.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Hongwei Zheng, and Zhiming Zheng. 2024b. Adacomp: Extractive context compression with adaptive predictor for retrieval-augmented large language models. *arXiv preprint arXiv:2409.01579*.
- Qianchi Zhang, Hainan Zhang, Liang Pang, Hongwei Zheng, and Zhiming Zheng. 2026b. Stable-rag: Mitigating retrieval-permutation-induced hallucinations in retrieval-augmented generation. *arXiv preprint arXiv:2601.02993*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with BERT](#). In *International Conference on Learning Representations*.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024c. [Benchmarking large language models for news summarization](#). *Transactions of the Association for Computational Linguistics*, 12:39–57.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026c. [Logical phase transitions: Understanding collapse in llm logical reasoning](#). *Preprint*, arXiv:2601.02902.
- Yunyao Zhang, Yihao Ai, Zuocheng Ying, Qirui Mi, Junqing Yu, Wei Yang, and Zikai Song. 2026d. [Coupling macro dynamics and micro states for long-horizon social simulation](#). *Preprint*, arXiv:2604.05516.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Zikai Song. 2026e. [Semantic-aware logical reasoning via a semiotic framework](#). *Preprint*, arXiv:2509.24765.
- Wei Zhu, Jian Zhang, Lixing Yu, Kun Yue, and Zhiwen Tang. 2026. [Dissecting failure dynamics in large language model reasoning](#). *Preprint*, arXiv:2604.14528.

## A Detailed Implementation Setup

### A.1 Dataset Sampling and Reproducibility

Conducting a large-scale evaluation involving computationally intensive Large Reasoning Models (LRMs) such as o1 and GPT-5, alongside multi-step reasoning strategies (e.g., Iterative Refinement), presents significant cost and latency challenges. To manage these constraints while ensuring statistically meaningful comparisons, we performed our evaluation on a fixed subset of the data.

Specifically, for each of the eight datasets we used, we randomly sampled 100 instances from the official test split. This sampling process was controlled by a unified global random seed (seed=42) to ensure reproducibility. Consequently, our final evaluation suite consists of 800 distinct document-summary pairs.

### A.2 Software and Environment

All experiments were conducted using Python. We utilized the standard HuggingFace datasets library for data loading and processing. The random sampling was implemented using Python’s native random module initialized with the aforementioned seed. All model interactions were performed via the Azure OpenAI API, with generation parameters fixed (temperature=0) as detailed in Section 4.2 to ensure deterministic outputs.

## B Detailed Main Results

We report detailed results for all methods and datasets, with ROUGE, BERTScore, SummaC, and AlignScore scores across 0-shot and 2-shot settings. Tables 8, 9, 10, and 11 present these full results.

### C Performance Results Grouped by Datasets

Here we reorganize the results by dataset to better understand how reasoning strategies behave across short-form, long-form, multi-document, and table-to-text summarization. As shown in Tables 4, 5, 6, and 7, this grouped view highlights domain-specific behavior more clearly.

### D Detailed Human Evaluation Results

Here, we list the detailed Human Evaluation Scores (Completeness, Conciseness, Faithfulness) of selected systems on three datasets. Table 12 reports the full human-annotated comparisons.

System	Short		Long		Table		Avg	
	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS
<i>LLM (GPT-4.1)</i>								
Vanilla	19.43	85.29	19.42	84.98	18.60	85.22	19.24	85.17
COT	17.00	85.27	18.40	84.56	17.63	84.16	17.61	84.78
SC	20.12	86.85	20.09	85.37	19.12	86.01	19.89	86.17
Cite	19.41	86.77	19.06	85.23	17.95	85.93	18.97	86.07
Deco	16.55	85.28	18.44	84.75	17.70	85.39	17.43	85.13
QAG	17.78	85.36	20.15	85.04	18.27	85.04	18.68	85.18
Plan	16.34	85.22	17.50	84.55	16.80	84.98	16.83	84.94
IR	19.42	85.97	20.68	85.44	19.10	85.52	19.77	85.70
E2A	20.11	86.77	19.74	85.24	18.76	85.92	19.69	86.07
<i>Large Reasoning Models (LRM)</i>								
o1	18.35	84.04	18.41	84.27	16.60	82.46	17.98	83.77
o3	16.61	83.74	17.43	83.85	15.08	83.70	16.54	83.77
GPT-5	17.34	81.66	17.66	83.47	15.26	82.75	16.99	82.51

Table 4: Zero-shot ROUGE (R) and BS scores grouped by dataset type: Short, Long, and Table. Avg = per-system mean.

System	Short		Long		Table		Avg	
	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS
<i>LLM (GPT-4.1)</i>								
Vanilla	22.56	85.57	20.95	83.81	19.96	86.05	21.63	84.97
COT	23.05	85.89	21.33	85.40	19.45	85.11	21.95	85.61
SC	21.65	86.24	20.83	85.26	19.13	86.16	21.03	85.86
Cite	20.67	83.95	19.69	84.03	19.43	86.39	20.15	84.28
Deco	18.64	84.48	18.16	82.44	18.84	85.78	18.49	83.88
QA	19.87	85.83	20.41	84.88	18.64	85.03	19.92	85.37
Plan	18.00	85.92	17.97	84.44	18.43	85.28	18.04	85.28
IR	22.18	86.88	21.32	85.62	19.94	86.34	21.58	86.34
E2A	21.79	86.24	20.06	85.42	19.60	86.27	20.87	85.93
<i>Large Reasoning Models (LRM)</i>								
o1	20.07	84.84	18.47	84.11	17.66	85.50	19.17	84.65
o3	17.45	81.99	17.37	82.88	15.51	84.28	17.18	82.61
GPT-5	19.07	83.53	18.36	83.60	15.86	83.83	18.40	83.59

Table 5: 2-shot ROUGE (R) and BS scores grouped by dataset type: Short, Long, and Table. Avg = per-system mean.

System	Short		Long		Table		Avg	
	SumC	AS	SumC	AS	SumC	AS	SumC	AS
<i>LLM (GPT-4.1)</i>								
Vanilla	39.69	60.55	26.48	80.49	34.83	93.78	34.13	72.18
COT	29.98	68.35	38.64	38.09	28.22	99.45	33.01	60.89
SC	35.82	38.40	23.73	79.27	33.19	88.79	30.96	60.03
Cite	38.35	41.69	27.34	81.60	33.10	86.38	33.57	62.24
Deco	36.06	55.35	23.97	80.70	29.31	76.49	30.68	67.25
QAG	28.84	65.48	38.90	45.32	24.84	98.63	32.11	62.06
Plan	34.08	39.01	23.15	80.73	30.28	83.90	29.51	60.27
IR	36.00	38.85	25.52	80.86	32.63	85.05	31.65	60.13
E2A	30.73	68.82	41.05	40.78	24.59	97.13	33.83	61.72
<i>Large Reasoning Models (LRM)</i>								
o1	29.89	68.92	40.36	50.86	29.09	97.30	33.71	65.70
o3	38.35	58.74	24.65	78.39	35.34	86.93	32.83	69.63
GPT-5	38.97	61.01	27.39	82.02	37.91	92.73	34.50	72.86

Table 6: Zero-shot Faithfulness grouped by dataset type: Short, Long, and Table. Metrics: SummaC (SumC) and AlignScore (AS). Avg = per-system mean.

### E Detailed GEval Results

Here, we list the detailed GEval Scores (Completeness, Conciseness, Faithfulness) for all methods across all datasets. Table 13 summarizes all GEval dimensions.

### F Compression Ratio and Abstractive-ness

We also provide the Compression Ratio and Abstractive-ness of the summaries. Table 14 shows the complete results.

### G Reasoning Strategy Selection

Based on our comprehensive evaluation, we summarize the recommended reasoning strategies for each dataset and domain. Table 15 highlights the

System	Short		Long		Table		Avg	
	SumC	AS	SumC	AS	SumC	AS	SumC	AS
<i>LLM (GPT-4.1)</i>								
Vanilla	33.81	62.66	39.79	37.17	27.59	98.50	35.27	57.58
COT	29.98	68.35	38.64	38.09	28.22	99.45	33.01	60.89
SC	28.11	63.41	38.16	37.90	24.61	98.44	31.44	58.22
Cite	28.77	66.94	38.88	37.24	24.66	97.84	32.05	59.66
Deco	26.62	67.48	37.16	42.38	25.04	98.59	30.37	61.95
QA-G	29.69	66.36	37.06	45.66	25.62	98.91	31.95	62.67
Plan	26.58	66.06	36.36	40.30	23.46	99.45	29.85	60.57
IR	28.02	66.35	37.41	37.33	24.47	98.70	31.10	59.51
E2A	30.10	65.06	39.35	39.04	26.46	98.71	33.11	59.51
<i>Large Reasoning Models (LRM)</i>								
o1	30.42	70.50	39.78	47.62	26.46	97.95	32.44	65.35
o3	28.74	54.95	42.45	63.10	23.84	98.25	33.27	63.42
GPT-5	33.93	61.53	41.70	71.61	26.53	99.39	35.92	70.04

Table 7: 2-shot Faithfulness grouped by dataset type: Short, Long, and Table. Metrics: SummaC (SumC) and AlignScore (AS). Avg = per-system mean.

most effective method for maximizing either summary quality or factual faithfulness under different settings.

## H Case Study

To study the potential hallucination caused by the reasoning (Chen et al., 2026b; Zhang et al., 2026b), here we show the full example of the generated summaries in Table 16 from the Wikihow dataset, together with the golden reference.

## I Human Evaluation Guidelines

Our human evaluation used two PhD students, each independently rating summaries using the rubric in Table 17 on a 1–5 scale for Conciseness, Completeness, and Faithfulness. We report the average of their scores across examples. Here we list the full human evaluation guidelines used for our manual annotation study, including rating scales and definitions.

## J Emerging Reasoning Directions for Future Adaptation

Beyond the methods benchmarked in this work, several recent directions may be adapted to abstractive summarization in future studies. Integrating Retrieval-Augmented Generation (RAG) (Zhang et al., 2026a) with reasoning-enhanced summarization represents a promising future direction, wherein relevant external passages can be dynamically retrieved to supplement key content segments with additional factual detail (Zhang et al., 2024b). Moreover, recent analyses of logical phase transitions indicate that reasoning behavior can degrade abruptly under certain regimes, which is relevant to stability-aware summarization design (Zhang et al., 2026c).

At the reasoning-policy level, neuron-level mixture routing has been explored as a way to make

reasoning trajectories more explainable and controllable (Dong et al., 2026), while error-driven refinement strategies provide another perspective on improving first-pass solution quality (Jiang et al., 2026). Reasoning-path search approaches are also promising for selecting stronger inference traces before generation (Ling et al., 2026; Chen et al., 2026a). Semantic-aware frameworks for logical reasoning may further help control intermediate reasoning states during summary construction (Zhang et al., 2026e).

We also note that progress in parameter-efficient adaptation may offer practical implementation benefits when transferring new reasoning modules to summarization settings (Dong et al., 2025). Meanwhile, advances in robust cross-modal retrieval and alignment can provide transferable architectural ideas, even when the source tasks are not summarization-specific (Li et al., 2026a,b). More broadly, work on long-horizon multi-agent dynamics suggests another possible route for modeling planner-decoder interactions in iterative summarization pipelines (Zhang et al., 2026d).

## K Prompts

To ensure ethical transparency and reproducibility, we disclose the prompts we used to deploy each reasoning strategy. The full templates are shown in Figures 7, 8, 9, 10, 11, 12, 13, 14, and 15.

System	CNN/DM		SAMSum		Reddit		WikiHow		ArXiv		Multi-News		BookSum		SciGen		Average	
	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS
<b>LLM (GPT-4.1)</b>																		
Vanilla	21.27	84.24	26.88	88.85	12.82	84.22	16.75	83.85	20.97	83.76	19.79	85.41	17.49	85.76	17.96	85.28	19.24	85.17
COT	20.88	84.88	21.03	88.11	11.36	83.85	14.72	84.22	20.39	83.59	19.95	85.44	14.87	84.65	16.55	84.65	17.47	84.92
SC	22.23	86.94	26.67	90.01	14.82	85.55	16.75	84.89	23.12	84.84	19.47	85.50	17.68	85.76	18.35	85.85	19.89	86.17
Cite	22.76	87.16	24.25	89.51	13.04	85.11	17.58	85.30	22.95	84.79	17.89	85.26	16.34	85.64	16.93	85.79	18.97	86.07
Deco	19.47	86.47	21.62	86.85	10.58	83.39	14.51	84.40	19.91	83.51	18.26	85.09	17.15	85.66	17.96	85.65	17.43	85.13
QAG	20.03	84.60	24.78	89.30	11.48	83.44	14.83	84.12	22.23	84.20	20.72	85.44	17.49	85.47	17.87	84.89	18.68	85.18
Plan	18.94	84.70	19.57	88.39	12.43	83.65	14.43	84.15	19.76	83.87	18.33	85.09	14.40	84.69	16.76	85.01	16.83	84.94
IR	22.68	86.16	25.10	88.79	13.68	84.27	16.22	84.67	23.65	84.85	20.67	85.64	17.71	85.84	18.42	85.34	19.77	85.70
E2A	22.86	87.01	26.51	89.85	13.87	85.36	17.21	84.85	22.68	84.52	19.05	85.43	17.49	85.77	17.84	85.78	19.69	86.07
<b>Large Reasoning Models (LRM)</b>																		
o1	21.12	85.01	23.50	87.59	12.74	80.74	16.04	82.82	20.15	83.95	17.97	84.17	17.12	84.69	15.21	81.16	17.98	83.77
o3	18.54	82.54	22.02	87.44	11.01	81.03	14.87	83.96	16.61	81.71	17.99	84.67	17.68	85.18	13.62	83.64	16.54	83.77
GPT-5	19.64	82.41	23.52	84.60	12.45	80.19	13.75	79.46	16.79	81.98	18.81	83.67	17.39	84.75	13.54	82.99	16.99	82.51

Table 8: Summary Coverage across 8 datasets with ROUGE (average of ROUGE 1/2/L) and BERTScore (BS). The last two columns show per-system averages across all 8 datasets.

System	CNN/DM		SAMSum		Reddit		WikiHow		ArXiv		Multi-News		BookSum		SciGen		Average	
	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS	ROUGE	BS
<b>LLM (GPT-4.1)</b>																		
Vanilla	23.39	83.82	31.58	90.37	16.05	82.56	19.23	85.52	23.29	83.83	20.22	82.27	19.33	85.32	19.96	86.05	21.63	84.97
COT	21.68	84.32	29.62	89.41	17.23	84.37	19.20	85.46	23.64	84.35	21.11	85.71	19.24	86.14	19.45	85.11	21.40	85.61
SC	21.96	84.58	29.87	90.57	16.07	84.97	18.71	84.83	23.97	85.03	19.51	84.77	19.00	85.99	19.13	86.16	21.03	85.86
Cite	21.58	84.66	27.49	90.11	14.94	76.96	18.66	84.05	23.55	84.13	18.51	83.86	17.01	84.09	19.43	86.39	20.15	84.28
Deco	19.96	84.96	25.27	86.02	12.18	81.69	17.16	85.26	19.50	79.80	17.19	81.68	17.80	85.83	18.84	85.78	18.49	83.88
QAG	20.91	84.09	27.73	89.89	14.59	84.95	16.25	84.39	22.41	84.30	20.52	84.63	18.31	85.70	18.64	85.03	19.92	85.37
Plan	19.89	85.75	23.71	89.34	12.97	83.99	15.42	84.59	19.79	83.15	18.41	85.15	15.71	85.02	18.43	85.28	18.04	85.28
IR	22.80	86.27	29.97	90.47	17.56	85.58	18.40	85.20	23.93	85.03	20.72	85.72	19.32	86.11	19.94	86.34	21.58	86.34
E2A	21.66	85.19	28.44	89.33	18.72	85.02	18.32	85.40	22.90	84.83	18.93	85.50	18.34	85.92	19.60	86.27	20.87	85.93
<b>Large Reasoning Models (LRM)</b>																		
o1	21.54	83.56	25.50	89.02	15.84	83.25	17.38	83.52	20.88	84.14	17.67	83.53	16.85	84.66	17.66	85.50	19.17	84.65
o3	18.28	82.47	22.67	86.84	12.55	74.04	16.31	84.61	16.30	81.44	18.44	84.69	17.36	82.51	15.51	84.28	17.18	82.61
GPT-5	19.80	82.41	25.58	87.67	14.49	79.96	16.41	84.06	17.22	82.11	19.63	84.64	18.23	84.05	15.86	83.83	18.40	83.59

Table 9: 2-shot performance across 8 datasets with ROUGE (average of ROUGE 1/2/L) and BERTScore (BS). The last two columns show per-system averages across all 8 datasets.

System	CNN/DM		SAMSum		Reddit		WikiHow		ArXiv		Multi-News		BookSum		SciGen		Avg	
	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS
<b>LLM (GPT-4.1)</b>																		
Vanilla	45.49	83.49	41.08	72.44	33.77	45.56	38.40	40.70	25.24	56.25	26.84	86.84	27.35	98.38	34.83	93.78	34.13	72.18
COT	30.94	42.67	27.41	86.62	27.50	49.92	34.06	94.17	40.54	7.46	37.07	39.78	38.31	67.03	28.22	99.45	33.01	60.89
SC	40.44	6.39	38.65	67.06	28.47	40.43	35.72	39.73	24.25	53.12	23.64	88.77	23.30	95.93	33.19	88.79	30.96	60.03
Cite	43.91	11.92	39.36	67.54	31.75	44.63	38.37	42.67	31.97	60.19	25.25	87.63	24.81	96.99	33.10	86.38	33.57	62.24
Deco	43.67	66.06	37.50	73.20	26.84	37.86	36.22	44.27	24.61	55.42	23.35	88.03	23.94	98.64	29.31	76.49	30.68	67.25
QAG	31.59	40.70	24.43	88.32	26.96	56.02	32.36	76.86	40.90	29.66	36.44	33.61	39.37	72.69	24.84	98.63	32.11	62.06
Plan	38.91	10.86	38.83	72.78	25.98	32.38	32.61	40.02	23.85	56.64	23.40	87.72	22.21	97.83	30.28	83.90	29.51	60.27
IR	40.99	8.53	38.66	69.23	29.45	40.04	34.90	37.58	25.17	54.72	24.61	89.57	26.77	98.30	32.63	85.05	31.65	60.13
E2A	33.58	43.24	25.96	90.14	26.51	53.41	36.87	88.49	45.74	12.12	37.56	42.19	39.84	68.04	24.59	97.13	33.83	61.72
<b>Large Reasoning Models (LRM)</b>																		
o1	31.24	39.83	27.61	86.52	26.02	54.86	34.67	94.46	44.23	31.71	38.02	51.92	38.82	68.96	29.09	97.30	33.71	65.70
o3	52.85	88.31	37.98	69.22	27.94	30.34	34.62	47.07	25.05	48.63	23.29	88.07	25.60	98.46	35.34	86.93	32.83	69.63
GPT-5	48.59	91.57	39.60	74.82	30.56	35.32	37.14	42.34	27.38	56.30	27.47	90.40	27.32	99.37	37.91	92.73	34.50	72.86

Table 10: Faithfulness across 8 datasets with SummaC (SumC) and AlignScore (AS). The last two columns show per-system averages across the 8 datasets.

System	CNN/DM		SAMSum		Reddit		WikiHow		ArXiv		Multi-News		BookSum		SciGen		Avg	
	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS	SumC	AS
<b>LLM (GPT-4.1)</b>																		
Vanilla	35.56	47.12	29.03	85.05	32.21	29.35	38.43	89.11	43.36	26.07	36.96	18.87	39.04	66.56	27.59	98.50	35.27	57.58
COT	30.94	42.67	27.41	86.62	27.50	49.92	34.06	94.17	40.54	7.46	37.07	39.78	38.31	67.03	28.22	99.45	33.01	60.89
SC	29.43	40.67	25.14	86.47	25.33	50.80	32.52	75.70	40.68	8.00	35.39	40.53	38.42	65.17	24.61	98.44	31.44	58.22
Cite	31.51	38.81	24.68	86.83	24.17	50.10	34.73	92.01	41.79	5.99	36.05	38.20	38.81	67.53	24.66	97.84	32.05	59.66
Deco	26.90	38.20	24.36	88.02	24.26	56.46	30.96	87.23	40.37	17.88	33.68	37.27	37.42	71.98	25.04	98.59	30.37	61.95
QAG	32.00	37.25	25.34	86.88	27.10	55.43	34.32	85.87	40.37	26.94	35.86	40.98	34.95	69.07	25.62	98.91	31.95	62.67
Plan	26.72	35.59	24.30	87.17	24.31	55.33	30.97	86.16	39.09	9.52	33.01	41.17	36.98	70.21	23.46	99.45	29.85	60.57
IR	29.62	40.42	25.11	88.71	24.59	44.35	32.76	91.91	39.87	5.65	34.79	40.57	37.57	65.76	24.47	98.70	31.10	59.51
E2A	32.85	37.28	26.56	87.62	26.53	42.49	34.47	92.83	42.96	8.68	37.31	43.38	37.77	65.07	26.46	98.71	33.11	59.51
<b>Large Reasoning Models (LRM)</b>																		
o1	32.72	51.27	27.41	86.62	27.50	49.92	34.06	94.17	44.24	27.61	37.08	49.02	38.01	66.22	26.46	97.95	33.44	65.35
o3	27.72	42.15	25.59	85.87	26.54	35.54	35.10	56.23	53.58	85.22	34.12	37.48	39.65	66.61	23.84	98.25	33.27	63.42
GPT-5	32.69	62.82	29.58	87.44	34.39	21.41	39.06	74.44	45.90	85.55	38.50	50.65	40.71	78.62	26.53	99.39	35.92	70.04

Table 11: 2-shot Faithfulness across 8 datasets with SummaC (SumC) and AlignScore (AS). The last two columns show per-system averages across the 8 datasets.

System	arXiv				SciGen				WikiHow				Average			
	Con	Com	Fai	Avg	Con	Com	Fai	Avg	Con	Com	Fai	Avg	Con	Com	Fai	Avg
Vanilla	3.75	4.05	4.15	3.98	4.20	3.60	3.90	3.90	4.15	3.80	4.60	4.18	4.03	3.82	4.22	4.02
SC	4.05	3.75	4.20	4.00	4.30	3.55	4.10	3.98	4.25	3.85	4.50	4.20	4.20	3.72	4.27	4.06
Deco	3.40	4.05	4.20	3.88	3.95	3.40	4.30	3.88	3.95	4.05	4.15	4.05	3.77	3.83	4.22	3.94
E2A	4.50	2.60	3.70	3.60	4.30	2.60	4.00	3.63	4.45	3.10	4.25	3.93	4.42	2.77	3.98	3.72
GPT-5	3.00	4.50	4.50	4.00	3.20	4.50	4.30	4.00	3.30	4.45	4.45	4.07	3.17	4.48	4.42	4.02

Table 12: Human evaluation results across datasets (arXiv, SciGen, WikiHow). Highest values per column highlighted in red.

(a) GEval-based Dimensions (1/2): CNN/DM, SAMSum, Reddit, WikiHow

System	CNN/DM			SAMSum			Reddit			WikiHow		
	Comp	Conc	Fact	Comp	Conc	Fact	Comp	Conc	Fact	Comp	Conc	Fact
<b>LLM (GPT-4.1)</b>												
Vanilla	4.82	4.60	4.97	4.96	4.69	4.99	4.92	4.69	4.95	4.94	4.64	4.98
COT	4.90	4.23	4.98	4.91	4.31	4.97	4.87	4.28	4.94	4.98	4.24	4.98
SC	4.89	4.63	5.00	4.92	4.73	5.00	4.91	4.83	4.99	4.96	4.70	4.99
Cite	4.77	4.89	5.00	4.93	4.44	4.98	4.55	4.69	5.00	4.92	4.94	4.99
Deco	4.98	4.21	5.00	4.94	3.85	4.98	4.90	3.91	4.94	4.98	4.15	4.99
QA-G	5.00	4.07	5.00	4.96	3.77	5.00	4.82	4.07	4.99	4.98	3.90	4.99
Plan	4.88	4.14	5.00	4.98	3.54	4.99	4.80	4.00	4.97	5.00	4.24	4.98
IR	4.97	4.55	5.00	4.96	4.46	5.00	4.94	4.65	4.96	4.98	4.51	4.98
E2A	4.89	4.88	5.00	4.96	4.85	5.00	4.66	4.88	5.00	4.86	4.78	4.99
<b>Large Reasoning Models (LRM)</b>												
o1	4.76	4.77	4.93	4.82	4.61	4.94	4.75	4.81	4.92	4.91	4.88	4.96
o3	4.70	4.74	4.93	4.88	4.81	4.96	4.43	4.64	4.91	4.99	4.88	4.99
GPT-5	4.79	4.47	5.00	4.77	4.67	4.94	4.55	4.52	4.88	4.82	4.37	4.96

(b) GEval-based Dimensions (2/2): ArXiv, Multi-News, BookSum, SciGen

System	ArXiv			Multi-News			BookSum			SciGen			Average			
	Comp	Conc	Fact	Comp	Conc	Fact	Comp	Conc	Fact	Comp	Conc	Fact	Comp	Conc	Fact	Avg
<b>LLM (GPT-4.1)</b>																
Vanilla	5.00	4.29	5.00	4.94	4.56	4.99	4.92	4.47	4.94	4.97	4.61	4.97	4.93	4.57	4.97	4.82
COT	5.00	4.05	5.00	4.98	4.24	5.00	5.00	3.17	4.98	4.99	3.93	4.98	4.95	4.06	4.98	4.66
SC	4.99	4.68	5.00	4.95	4.59	4.98	4.98	4.63	4.97	4.84	4.76	4.99	4.93	4.69	4.99	4.87
Cite	4.81	4.86	4.98	4.83	4.84	4.98	4.93	4.88	4.97	4.71	4.80	4.98	4.81	4.79	4.98	4.86
Deco	4.98	3.91	4.97	4.92	4.10	4.98	5.00	4.23	5.00	4.86	4.20	4.96	4.95	4.07	4.98	4.66
QA-G	3.83	3.56	4.97	1.93	3.97	4.98	2.03	2.50	4.85	4.38	3.92	4.99	3.99	3.72	4.97	4.23
Plan	4.98	4.38	5.00	4.91	4.14	4.95	4.95	4.04	4.99	4.99	3.69	4.97	4.93	4.02	4.98	4.65
IR	5.00	4.69	5.00	4.96	4.51	4.97	5.00	4.53	5.00	4.91	4.30	4.99	4.97	4.53	4.99	4.83
E2A	4.67	4.92	4.99	4.68	4.82	4.99	4.24	4.83	4.95	4.47	4.87	4.99	4.68	4.85	4.99	4.84
<b>Large Reasoning Models (LRM)</b>																
o1	4.97	4.88	5.00	4.80	4.63	4.92	4.94	4.87	4.98	4.88	4.62	4.92	4.85	4.76	4.95	4.87
o3	5.00	4.87	5.00	4.86	4.67	4.97	4.97	4.73	5.00	4.98	4.40	4.98	4.85	4.72	4.97	4.84
GPT-5	5.00	4.33	5.00	4.79	4.20	4.98	4.93	4.25	4.99	4.97	4.35	4.96	4.83	4.40	4.96	4.73

Table 13: Performance across 8 datasets with GEval dimensions: Completeness (Comp), Conciseness (Conc), and Factualness (Fact).

System	CNN/DM		SAMSum		Reddit		WikiHow		ArXiv		Multi-News		BookSum		SciGen		Average	
	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.	CR	Abs.
<b>LLM (GPT-4.1)</b>																		
Vanilla	21.1%	52.84%	63.65%	72.7%	25.62%	67.86%	31.22%	58.59%	6.54%	51.98%	11.29%	52.08%	8.68%	67.94%	63.2%	66.46%	28.91%	61.31%
COT	26.56%	54.67%	73.41%	81.72%	28.73%	73.12%	37.27%	65.02%	7.28%	52.8%	14.22%	53.86%	14.78%	74.05%	100.97%	70.02%	37.9%	65.66%
SC	20.73%	54.68%	61.04%	74.7%	23.41%	65.13%	30.88%	60.85%	4.66%	48.25%	10.83%	53.24%	8.38%	68.35%	48.64%	69.38%	26.07%	61.82%
Cite	17.85%	57.11%	72.81%	79.33%	22.32%	63.83%	26.13%	64.64%	3.45%	47.3%	9.06%	54.11%	6.91%	68.91%	50.86%	71.7%	26.17%	63.37%
Deco	24.12%	61.15%	89.88%	80.02%	33.74%	72.18%	34.93%	67.82%	7.44%	50.53%	13.34%	57.44%	9.16%	70.62%	68.24%	73.98%	35.11%	66.72%
QAG	30.03%	53.34%	83.15%	76.68%	38.42%	68.45%	42.8%	59.04%	6.89%	49.33%	16.44%	51.76%	12.62%	66.19%	102.12%	65.92%	41.56%	61.34%
Plan	24.52%	60.29%	91.9%	82.29%	34.43%	67.63%	37.72%	65.66%	5.02%	55.66%	12.84%	57.55%	9.3%	72.38%	107.32%	76.19%	40.38%	67.21%
IR	23.35%	51.65%	69.54%	75.07%	28.4%	65.05%	33.96%	58.61%	4.77%	47.03%	12.03%	50.62%	8.9%	66.74%	72.92%	64.71%	31.78%	59.98%
E2A	17.66%	53.55%	60.96%	75.12%	21.18%	67.36%	23.98%	59.41%	4.05%	48.44%	9.56%	52.6%	7.56%	68.47%	42.89%	66.79%	23.48%	61.47%
<b>Large Reasoning Models (LRM)</b>																		
o1	19.2%	54.62%	67.75%	75.34%	24.08%	69.56%	28.83%	60.78%	4.03%	54.78%	10.51%	53.03%	7.09%	70.16%	67.9%	65.89%	28.67%	63.02%
o3	21.18%	59.63%	63.37%	76.85%	24.78%	70.8%	32.19%	66.1%	7.27%	59.96%	12.29%	55.92%	9.37%	69.49%	88.39%	65.11%	32.35%	65.48%
GPT-5	25.72%	52.35%	69.36%	70.8%	27.18%	64.86%	39.3%	57.59%	10.62%	56.35%	15.35%	51.08%	12.13%	65.9%	89.66%	56.77%	36.17%	59.46%

Table 14: Compression Ratio (CR) and Abstractiveness (Abs.) of the summaries generated by 0-shot experiments are reported across eight datasets, with all values presented as percentages rounded to two decimal places. The last two columns show the average results of each system over the eight datasets.

Dataset	Domain	0-shot Setting		2-shot Setting	
		Quality	Faithfulness	Quality	Faithfulness
<i>Short-Form Documents (SDS)</i>					
CNN/DM	News	Cite	GPT-5	IR	Vanilla
SAMSum	Dialogue	SC	GPT-5	SC	GPT-5
Reddit	Social Media	SC	Vanilla	IR	QAG
WikiHow	Knowledge Base	Cite	E2A	Vanilla	CoT
<i>Long-Form Documents (MDS / LNS)</i>					
ArXiv	Scientific Paper	IR	Cite	SC	GPT-5
Multi-News	News	IR	GPT-5	IR	GPT-5
BookSum	Novels	IR	GPT-5	CoT	GPT-5
<i>Table-to-Text (TTS)</i>					
SciGen	Scientific Paper	SC	GPT-5	IR	CoT

Table 15: Recommended reasoning strategies across datasets and domains. Methods are selected based on the **average ranking** of constituent metrics: *Quality* considers both ROUGE and BERTScore, while *Faithfulness* considers SummaC and AlignScore. The best method (lowest average rank) is shown for each setting.

<b>Document</b>	Clean the knife and cutting board with hot water and soap, and be sure to rinse both off thoroughly. Your vegetables don't need to be washed with soap, but rinse with hot water and consider scrubbing with a cloth or brush. Use a sharp, non-serrated knife for cutting. A chopping knife works well but is not required. It's important to only use a sharp knife so you can make clean cuts. Potatoes and carrots are often peeled before being julienned. It is up to you depending on how you plan to eat them. If you don't peel your potatoes, at least make sure buds or eyes are removed from the skin. Use a peeler or a paring knife to remove the skin. Cut off the stem and bottoms when necessary. Slice off one side at a time so you are left with a rectangular chunk of vegetable. You can throw the scraps away, or you can slice them up and mix them in with the uniform pieces you'll cut later. This is especially helpful for potatoes, zucchini, and squash. If you are going to be eating the pieces raw, they can be a little longer, but <b>if you're cooking with them, 2–3 inches is best</b> . Each chunk does not have to be the exact same length, just divide the vegetable evenly. The general rule for juliennes is that you'd be able to eat them in one bite, but you can cut them into longer chunks if you want. When cutting juliennes, $\frac{1}{8}$ in (0.3 cm) to 1 be careful not to cut yourself as you slice. <b>Flip the slices</b> you've just cut so that they are piled on top of each other. Then cut them into <b>sticks that are <math>\frac{1}{8}</math> in (0.3 cm) to <math>\frac{1}{16}</math> in (0.15 cm) thick</b> , again. The goal with juliennes is to get them all to a fairly uniform size.
<b>Reference</b>	Wash the knife, cutting board, and vegetables. Peel the vegetable if you want to. Cut rounded vegetables into a rectangle. chop the vegetable into <b>chunks that are 2-3 inches ( 5-8 cm ) long</b> . Slice the vegetable down the length. <b>stack 2-3 slices</b> and cut again .
<b>Vanilla</b>	Clean the knife and cutting board with hot water and soap, rinsing thoroughly. Rinse vegetables with hot water, scrubbing if needed; do not use soap. Use a sharp, non-serrated knife, or optionally a chopping knife, for clean cuts. Peel potatoes and carrots if desired, and remove eyes from potatoes. Use a peeler or paring knife to remove skin and cut off stems and bottoms as needed. Slice vegetables into rectangular chunks, discarding or saving scraps. For julienne cuts, divide vegetables evenly; <b>raw pieces can be longer, but cooked ones should be 2–3 inches</b> . Cut slices into <b>sticks <math>\frac{1}{8}</math> in (0.3 cm) to <math>\frac{1}{16}</math> in (0.15 cm) thick, stacking slices</b> for uniformity, and always be careful not to cut yourself. Aim for uniform, bite-sized pieces.
<b>SC</b>	To julienne vegetables, clean your knife and cutting board thoroughly with hot water and soap, and rinse well. Rinse vegetables with hot water and scrub if needed, but do not use soap. Use a sharp, non-serrated knife for clean cuts; peeling potatoes and carrots is optional, but remove any eyes or buds if not peeling. Cut off stems and bottoms as needed, slice vegetables into rectangular chunks, and either discard or use scraps. For raw pieces, longer cuts are acceptable; <b>for cooking, 2–3 inch lengths are best</b> . Aim for uniform, bite-sized <b>sticks about <math>\frac{1}{8}</math> to <math>\frac{1}{16}</math> inch thick, stacking slices</b> before cutting, and always be careful not to cut yourself.
<b>Deco</b>	Begin by thoroughly cleaning your knife and cutting board with hot water and soap, and rinsing vegetables with hot water to ensure hygiene. Use a sharp, non-serrated knife for safe, clean cuts. Prepare vegetables by peeling them if desired, removing unwanted parts like stems, bottoms, or potato eyes, and cutting them into manageable rectangular chunks based on your intended use and preference. For julienning, <b>stack the slices</b> and cut them into <b>uniform sticks about <math>\frac{1}{8}</math> to <math>\frac{1}{16}</math> inch thick and 2–3 inches long</b> , aiming for consistency in size to ensure even cooking and presentation.
<b>E2A</b>	To julienne vegetables, first clean your knife and cutting board thoroughly. Use a sharp, non-serrated knife for precise cuts. Whether or not you peel vegetables depends on how you plan to eat them. Slice off one side at a time to create rectangular chunks, aiming for pieces that are generally bite-sized, though longer chunks are acceptable if preferred. <b>The key is to cut the juliennes to a fairly uniform size.</b>
<b>GPT-5</b>	<ul style="list-style-type: none"> <li>- Clean knife and cutting board with hot, soapy water; rinse thoroughly. Rinse vegetables with hot water (no soap); scrub with a cloth/brush if needed.</li> <li>- Use a sharp, non-serrated knife (a chopping knife works). Sharpness ensures clean cuts.</li> <li>- Peel potatoes and carrots if desired; always remove potato eyes. Use a peeler or paring knife. Trim stems and bottoms as needed.</li> <li>- Square the vegetable by slicing off sides to form a rectangular chunk; discard or use trimmings. This is especially helpful for potatoes, zucchini, and squash.</li> <li>- Cut the chunk into even slabs, <b>stack them</b>, then slice into <b>uniform sticks about <math>\frac{1}{8}</math>–<math>\frac{1}{16}</math> inch (<math>\approx</math> 0.3–0.15 cm) thick</b>. Aim for one-bite, fairly uniform pieces; be careful not to cut yourself.</li> <li>- Length: <b>for cooking, 2–3 inches is best; raw pieces can be a little longer</b>. Divide the vegetable into evenly sized chunks.</li> </ul>

Table 16: Case study comparing summaries generated by different reasoning strategies against a human reference. **Blue text** in summaries shows their better faithfulness to the source, as they manage to keep these important procedural details that the reference missed. **red text** marks hallucinations or omissions, such as **Deco**'s failure to recognize conditional branches (generalizing cooking length to all cases) or **E2A**'s complete omission of the final cutting step.

<b>Criterion</b>	<b>Score</b>	<b>Description</b>
<b>Conciseness</b>	5	The summary is highly concise, containing only essential information with no redundancy or irrelevant content.
	4	Mostly concise with minor redundancy; still clear and compact overall.
	3	Contains some unnecessary details or minor repetition that slightly affects brevity.
	2	Wordy or includes multiple redundant or off-topic parts; could be shortened significantly.
	1	Very verbose or includes large amounts of irrelevant content; lacks focus and compactness.
<b>Completeness</b>	5	Fully covers all key points and information in the source document; no important content omitted.
	4	Covers most key information, with only minor omissions that do not affect overall understanding.
	3	Misses some moderately important points; the reader may lose some aspects of the original meaning.
	2	Covers only a portion of the important content; several key ideas missing.
	1	Severely incomplete; fails to convey most of the main ideas of the source.
<b>Faithfulness</b>	5	Entirely faithful to the source; no hallucination, distortion, or factual errors.
	4	Mostly faithful, with only minor wording deviations or small factual imprecision.
	3	Some inaccuracies or inferred details not explicitly supported by the source.
	2	Several incorrect or misleading statements that distort the original meaning.
	1	Largely unfaithful; contains major fabrications or contradictions with the source content.

Table 17: Human evaluation guidelines for summarization quality. Scores range from 1 to 5, with 0.5-point increments allowed for finer judgment.

## Vanilla and CoT

Vanilla Prompt

**Instruction:**  
You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text.

**Formatting Constraint:**  
Then write ONLY the final summary text (no lists, no bullets, no metadata). Aim to be as concise as possible while fully covering the key points.

**[Begin Document]**  
{*Input Document Content Here*}  
**[End Document]**

SUMMARY:

Figure 7: The Vanilla prompt used for the baseline setting. It provides general summarization constraints without specific reasoning instructions.

Chain-of-Thought (CoT) Prompt (News Domain)

**Instruction:**  
You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text.

**Reasoning Guidance (Injected by Domain):**  
First, internally (do NOT output lists), identify the following to guide your reasoning:

- important entities (people, organizations, places)
- important dates or time expressions
- key events or actions
- the outcome or conclusion of these events

*Writing guidance: Keep chronology clear; avoid speculation; reflect the article's scope and tone.*

**Formatting Constraint:**  
Then write ONLY the final summary text (no lists, no bullets, no metadata). Aim to be as concise as possible while preserving essentials.

**[Begin Document]**  
{*Input Document Content Here*}  
**[End Document]**

SUMMARY:

Figure 8: The Chain-of-Thought (CoT) prompt template adapted for the News domain. The model identifies key entities and timelines internally before generating the summary.

## Self-Consistency (SC)

Self-Consistency (SC) Step 1: Candidate Generation

**Instruction:**  
You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text.

**Formatting Constraint:**  
Then write ONLY the final summary text (no lists, no bullets, no metadata). Aim to be as concise as possible while fully covering the key points.

**[Begin Document]**  
{*Input Document Content Here*}  
**[End Document]**

SUMMARY:

## Self-Consistency (SC) Step 2: Evaluation and Selection

### Instruction:

You are a rigorous summarization evaluator. Given a DOCUMENT and three candidate abstractive summaries (A/B/C), score each summary and select ONE best candidate.

### Evaluation Rubric:

For each candidate, assign sub-scores based on the following weights:

- **Faithfulness (0.35)**: no hallucination; consistent with DOCUMENT.
- **Coverage (0.35)**: captures all core ideas; no major omissions.
- **Coherence (0.15) / Concision (0.15)**.

### Output Constraint:

Return ONLY a JSON object with the following keys:

- scores: object with keys A/B/C.
- winner: one of 'A', 'B', 'C'.
- reason: explanation ( $\leq 50$  words).
- final\_summary: EXACT text of the winning candidate.

### Input Context:

DOCUMENT: {...} CANDIDATE\_A/B/C: {...}

Figure 9: Prompts used for the Self-Consistency (SC) method. **Top:** The prompt for generating candidate summaries (Step 1). **Bottom:** The prompt for the evaluator to score and select the best summary (Step 2).

## Cited Summarization (Cite)

### Cited Summarization (Cite) Prompt

#### Instruction:

You are a faithful abstractive summarizer and evidence aligner. Write an abstractive summary for the DOCUMENT and, for each summary sentence, cite up to 3 supporting sentences from the provided SENTENCES array and explain briefly why it is important.

#### Guidelines:

- **Faithfulness**: Do not invent content not supported by DOCUMENT.
- **Citations**: Every summary sentence must have 1–3 supports from the SENTENCES list.
- **Paraphrase**: Do not copy support sentences verbatim (short quotes allowed).
- **Concision**: Use the minimum length that still covers key ideas.

#### Output Constraint (JSON):

Output ONLY a JSON object matching the schema:

- summary\_text: string (the final abstractive summary).
- alignments: array of objects linking summary parts to evidence:
  - summary\_id: integer (matching the summary sentence).
  - support: array of indices from the SENTENCES list.
  - importance\_reason: string (explanation).
  - support\_strength: integer [1, 5].

#### Input Context:

SENTENCES (0-based array): ["Sentence 1...", "Sentence 2...", ...]  
DOCUMENT: {*Input Document Content Here*}

Figure 10: The prompt for Cited Summarization (Cite). It requires the model to generate the summary and simultaneously provide evidence alignments (indices from the source sentences) in a structured JSON format.

## Decomposition (Deco)

### Decomposition (Deco) Step 1: Chunking & Local Summarization

**Instruction:**

You are a planner and summarizer. Decompose the document into coherent contiguous chunks, then write a concise abstractive summary and a set of short supporting contexts for each chunk.

**Task Guidelines:**

- **Decompose:** Split the document into  $M$  contiguous chunks (adaptively, typically 5–20 sentences each).
- **Local Summary:** For each chunk, write an abstractive summary.
- **Contexts:** Provide 3–6 short supporting contexts (key facts needed for coherence) for each chunk.

**Output Constraint (JSON):**

Return ONLY a JSON object matching the schema:

- `doc_stats`: total sentences and chunk count.
- `chunks`: array of objects, each containing:
  - `id` and `span` (start/end indices).
  - `summary`: string (local summary).
  - `contexts`: array of strings (supporting facts).

**[Begin Document]**

{*Input Document Content Here*}

**[End Document]**

### Decomposition (Deco) Step 2: Merging

**Instruction:**

You are a merger that produces a single high-quality abstractive summary. Merge the given chunk summaries into one single concise summary that covers all key information.

**Reasoning Guidance:**

Use the provided *chunk summaries* as the primary gist. Use the *contexts* to proofread, resolve conflicts, keep terminology consistent, and maintain a clear overall chronology.

**Formatting Constraint:**

Output ONLY the final summary text (no lists, no JSON, no mention of 'document' or 'chunk').

**Input Context:**

Below are summaries of different parts:

{*List of Chunk Summaries (JSON from Step 1)*}

Below are the supporting contexts for the summaries:

{*List of Chunk Contexts (JSON from Step 1)*}

SUMMARY:

Figure 11: Prompts used for the Decomposition (Deco) method. **Top:** Step 1 prompt for decomposing the document and generating local summaries with contexts. **Bottom:** Step 2 prompt for merging the local outputs into a final coherent summary.

## Question-Answer Guided (QAG)

### QAG Step 1: Question Generation

**Instruction:**

You are a question planner for summarization. Read the document and write the exact questions you need to ask to extract its summary facts.

**Task Guidelines:**

Propose a focused set of concrete, answerable questions (4–8 adaptively). Aim for coverage first, then depth. Consider the following facets:

- topic, key\_pts, entities, timeline
- numbers, outcomes, challenges, insights

**Output Constraint (JSON):**

Return ONLY a JSON object with a questions array (keys: id, facet, question).

**[Begin Document]**

{*Input Document Content Here*}

**[End Document]**

### QAG Step 2: Answering

**Instruction:**

You answer the proposed questions using ONLY the given context. Answer each question precisely and concisely. Keep numbers/dates/entities exact; avoid speculation.

**Output Constraint (JSON):**

Return ONLY a JSON object with an answers array. Each item must contain:

- id and question (from input).
- answer: string.
- confidence: integer (1–5).

**Input Context:**

DOCUMENT: {*Input Document Content Here*}

QUESTIONS: {*Question List JSON from Step 1*}

### QAG Step 3: Summarization

**Instruction:**

You are a summarizer. Use the Q&A table as a structured guide to coverage, and the DOCUMENT as the factual ground truth.

**Requirements:**

- **Coverage:** reflect high-confidence answers across facets.
- **Faithfulness:** facts must be supported by the DOCUMENT; resolve contradictions using the DOCUMENT.
- **Coherence:** clear flow, consistent terminology.

**Input Context:**

Q&A TABLE: {*Answers JSON from Step 2*}

DOCUMENT: {*Input Document Content Here*}

SUMMARY:

Figure 12: Prompts used for the Question-Answer Guided (QAG) method. **Step 1** generates guiding questions based on specific facets. **Step 2** answers these questions using the source text. **Step 3** synthesizes the final summary using the Q&A pairs as a scaffold and the document for factual grounding.

## Plan-then-Write (Plan)

### Plan-then-Write (Plan) Step 1: Planning

**Instruction:**

You are a summarization planner for a single-document abstractive summary. Infer the document's domain/genre from content and design a minimal plan accordingly.

**Task Guidelines:**

- Identify the **domain** (e.g., news, scientific, narrative).
- Clarify the **goal** and intended **audience**.
- Decide a suitable writing **style** (e.g., 'concise, neutral').
- Provide 3–5 categories of **salient info** to cover.
- Provide an adaptive **length\_guidance** based on density/complexity.

**Output Constraint (JSON):**

Return ONLY a valid JSON object with keys: domain, goal, audience, style, salient\_info, length\_guidance.

**[Begin Document]**

SOURCE (truncated if long):

{*Input Document Content Here*}

### Plan-then-Write (Plan) Step 2: Execution

**Instruction:**

**Instruction:**

You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text. Follow the provided PLAN strictly. Produce an abstractive summary with adaptive length.

**Requirements (Derived from Plan):**

- **Domain context:** {*domain*}
- **Style:** {*style*}
- **Salient categories:** {*salient\_info*}
- **Length guidance:** {*length\_guidance*}

**Formatting Constraint:**

Do NOT output lists, bullets, headings, or metadata. Output the summary text only. Prefer the minimum length that fully covers the salient information.

**Input Context:**

DOCUMENT: {*Input Document Content Here*}

PLAN: {*Generated Plan JSON from Step 1*}

SUMMARY:

Figure 13: Prompts used for the Plan-then-Write (Plan) method. **Step 1** prompts the model to generate a structural plan including domain, style, and salient information. **Step 2** uses this plan to guide the generation of the final abstractive summary.

## Iterative Refinement (IR)

### Iterative Refinement (IR) Step 1: Drafting

**Instruction:**

You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text.

**Formatting Constraint:**

Then write ONLY the final summary text (no lists, no bullets, no metadata). Aim to be as concise as possible while fully covering the key points.

**[Begin Document]**

{*Input Document Content Here*}

**[End Document]**

SUMMARY:

Figure 14: Prompts used for the Iterative Refinement (IR) method. **Step 1** generates an initial draft. **Step 2** evaluates the draft and provides structured feedback (JSON). **Step 3** revises the summary based on the feedback. Steps 2 and 3 repeat until the stop condition is met.

### Iterative Refinement (IR) Step 2: Evaluating

**Instruction:**

You are a summary evaluator. Give a score (1–5) for the summary and provide precise revise suggestions (add, remove, rephrase, shorten, or none).

**Output Constraint (JSON):**

Return ONLY a JSON object with:

- score: integer 1–5.
- suggestions: array of objects (type, content, evidence).
- stop: boolean (true if score is 5 and no revision needed).

**Input Context:**

DOCUMENT: {*Input Document Content Here*}

CURRENT\_SUMMARY: {*Summary from Previous Step*}

### Iterative Refinement (IR) Step 3: Refining

**Instruction:**

Revise the CURRENT\_SUMMARY according to the EVALUATION suggestions.

**Rules:**

- Apply actionable suggestions faithfully (add/remove/rephrase/shorten).
- If any suggestion introduces hallucination, ignore it.
- Output ONLY the revised summary text.

**Input Context:**

DOCUMENT: {*Input Document Content Here*}

CURRENT\_SUMMARY: {*Summary from Previous Step*}

EVALUATION (JSON): {*Suggestions from Step 2*}

SUMMARY:

## Extract-to-Abstract (E2A)

### Extract-to-Abstract (E2A) Step 1: Extraction

**Instruction:**

You are an extractive selector. Your job is to pick the most salient sentences from a document.

**Selection Policy:**

- Split document into sentences ( $N$ ).
- Decide selection budget  $K$  based on  $N$  (e.g., if  $N \leq 6 \rightarrow K = \min(N, 3)$ ; if  $N > 60 \rightarrow K = 14$ ).
- Select  $K$  sentences that maximize salience and minimize redundancy.
- Keep each selected sentence **EXACTLY** as it appears (verbatim).

**Output Constraint (JSON):**

Return ONLY a valid JSON object with keys:

- stats: object with total\_sentences and selected\_budget.
- selected: array of objects (index, text).

**[Begin Document]**

{*Input Document Content Here*}

**[End Document]**

### Extract-to-Abstract (E2A) Step 2: Abstraction

**Instruction:**

You are a skilled analyst tasked with producing a concise and accurate summary from the given text. The summary must preserve all essential facts, maintain logical flow, and avoid introducing information not present in the text.

**Goal:**

- Preserve essential facts supported by EVIDENCE (and consistent with DOCUMENT).
- Maintain logical flow and coherence.
- Avoid copying sentences verbatim (paraphrase).

**Input Context:**

EVIDENCE: {*Extracted Sentences JSON from Step 1*}

DOCUMENT: {*Input Document Content Here*}

SUMMARY:

Figure 15: Prompts used for the Extract-to-Abstract (E2A) method. **Top:** Step 1 selects salient sentences based on an adaptive budget policy. **Bottom:** Step 2 generates the abstractive summary conditioned on both the extracted evidence and the full document context.