

UT-ACA: Uncertainty-Triggered Adaptive Context Allocation for Long-Context Inference

Lang Zhou^{1,2}, Shuxuan Li¹, Zhuohao Li^{1,2}, Shi Liu^{2,3},
Wei-Shi Zheng^{1,2}, Zhilin Zhao^{1,2*}

¹ School of Computer Science and Engineering, Sun Yat-sen University

² Shenzhen Loop Area Institute ³ Southern University of Science and Technology
zhoulang3@mail2.sysu.edu.cn, zhaozhlin@mail.sysu.edu.cn

Abstract

Long-context inference remains challenging for large language models due to attention dilution and out-of-distribution degradation. Context selection mitigates this limitation by attending to a subset of key-value cache entries, yet most methods allocate a fixed context budget throughout decoding despite highly non-uniform token-level contextual demands. To address this issue, we propose Uncertainty-Triggered Adaptive Context Allocation (UT-ACA), an inference-time framework that dynamically adjusts the context window based on token-wise uncertainty. UT-ACA learns an uncertainty detector that combines semantic embeddings with logit-based confidence while accounting for uncertainty accumulation across decoding steps. When insufficient evidence is indicated, UT-ACA selectively rolls back, expands the context window, and regenerates the token with additional support. Experiments show that UT-ACA substantially reduces average context usage while preserving generation quality in long-context settings. Code and dataset are available at <https://github.com/Tommy307/UT-ACA>.

1 Introduction

Large language models (LLMs) are expected to operate in long-context settings, such as long-document question answering, multi-document summarization, and evidence-grounded reasoning over extensive inputs (Zhang et al., 2025b; Wu et al., 2025c; Li et al., 2024a). In these scenarios, effective long-context handling is pivotal for robust generation, requiring reliable evidence retrieval and globally consistent generation. However, as context length increases, redundant or weakly relevant tokens dilute attention and hinder precise evidence retrieval. Furthermore, long contexts introduce out-of-distribution (OOD) dependencies and positional patterns that diverge from pretraining data (Ding




Figure 1: Comparison of context management strategies. The figure contrasts full-context decoding, fixed-size context selection, and the proposed adaptive context allocation strategy. Full-context decoding attends to all available KV entries, fixed selection uses the same retrieved budget at every step, and UT-ACA expands the window only when uncertainty indicates that additional evidence is needed.

et al., 2024; Peng et al., 2024; Li et al., 2024b), resulting in degraded calibration and unstable reasoning during length extrapolation.

Context selection has emerged as an effective strategy for mitigating long-context inference challenges by restricting attention to relevant Key-Value (KV) entries, thereby reducing attention dilution and computational overhead. Nevertheless, most existing approaches (Wu et al., 2025b; Hao et al., 2025; Tang et al., 2024) maintain a fixed context budget throughout decoding, which implicitly assumes that generation difficulty remains uniform across tokens. This assumption is often violated in practice: many tokens can be reliably predicted from local context, whereas others depend on long-range evidence dispersed across the prompt. Consequently, fixed-size context windows may be unnecessarily costly for easy generation steps yet inadequate for difficult ones, highlighting the need for adaptive token-wise context allocation during decoding.

*Corresponding author.

A natural implication is that context requirements vary at the token level and should be adjusted dynamically during decoding. Token-wise generation difficulty can be inferred from uncertainty signals available at inference time, enabling adaptive context control: decoding generally defaults to a compact window, and when high uncertainty or hallucination risk is detected, the process reverts, expands the context window, and regenerates the token with additional evidence. In this formulation, uncertainty estimation serves as an active decision signal for selective context enlargement rather than a passive diagnostic measure.

Accordingly, we propose **UT-ACA**, an inference-time framework that dynamically adjusts the context window based on token-wise uncertainty. As shown in Figure 1, UT-ACA generally decodes with a compact context window and expands it only when insufficient evidence or elevated hallucination risk is detected. Uncertainty is estimated via a learned dual-encoder detector combining logit margin and semantic embeddings from hidden states, a lightweight strategy that avoids the overhead of multi-sample decoding (Khairi et al., 2025; Li et al., 2025a; Kuhn et al., 2023). To train the detector, we construct a synthetic dataset of 10,000 training, 120 validation, and 100 test samples from LLM output logits and hidden states, and adopt a dual-encoder architecture that fuses semantic and confidence features with an LSTM-based temporal module. Experiments show that UT-ACA generalizes across diverse long-context settings, achieving comparable or improved generation quality while reducing context token consumption.

2 Preliminaries

At inference time, LLMs maintain a KV cache of past token representations across attention layers (Hooper et al., 2024; Liu et al., 2024). At each decoding step t , we define the *context window* $\mathcal{C}_t \subseteq \{(K_i, V_i)\}_{i=1}^{t-1}$ as the subset of KV entries participating in attention computation, and self-attention is computed as

$$\text{Attn}(Q_t, \mathcal{C}_t) = \text{softmax}\left(\frac{Q_t K_{\mathcal{C}_t}^\top}{\sqrt{d}}\right) V_{\mathcal{C}_t}. \quad (1)$$

In long-context settings, attending to the full KV cache is computationally expensive and often unnecessary (An et al., 2024), motivating adaptive strategies that dynamically adjust $|\mathcal{C}_t|$.

Rollback and Regeneration. Autoregressive generation is susceptible to error accumulation (Huang et al., 2024): uncertain or incorrect predictions at earlier steps propagate through the KV cache and affect subsequent decoding. We denote the decoding state at step t as $\xi_t = (x_{<t}, \mathcal{C}_t)$. When a token is generated under insufficient contextual evidence, a rollback operation restores the pre-generation snapshot for step (t) , after which the token is regenerated under an expanded context window.

Problem Formulation. Given a long-context input and an autoregressive LLM, this work studies how to adaptively allocate the context window $\mathcal{C}_t$ at each decoding step. Rather than committing to a fixed context budget throughout decoding, the goal is to default to a compact window when local evidence is sufficient, and to detect uncertain steps that warrant rollback, context expansion, and regeneration. In this way, the model aims to reduce average context token consumption while preserving generation quality in long-context inference.

3 UT-ACA Framework

We propose **UT-ACA**, an inference-time framework consisting of two core modules, the *uncertainty detector* and the *adaptive context window*, which jointly enable token-wise dynamic context allocation during decoding, as illustrated in Figure 2. At most decoding steps, UT-ACA performs a tentative generation with a compact context window and feeds both the output logits and hidden-state embeddings into the *uncertainty detector* to obtain the *Generation Difficulty Metric* (GDM). Conditioned on the GDM, the *adaptive context window* module either accepts the tentative token and proceeds with the compact window, or triggers rollback, context expansion, and regeneration with enriched contextual support.

3.1 Uncertainty Detector

To support uncertainty-aware decoding, we design a lightweight token-level uncertainty estimator that operates at each generation step. The estimator takes two complementary signals as input, namely the logit margin and a semantic embedding extracted from model hidden states. It then fuses these signals with a dual-encoder module and models uncertainty accumulation over time with a Long Short-Term Memory (LSTM) layer (Zhao et al., 2025b; Qiu et al., 2024), producing a three-way




Figure 2: Workflow of our framework. The figure illustrates how UT-ACA connects long-context input, token-level uncertainty estimation, and adaptive context allocation during decoding. (a) The prompt provides the long-context input and target instruction. (b) The uncertainty detector converts output logits and semantic embeddings into a Generation Difficulty Metric (GDM). (c) The adaptive context window accepts low-uncertainty tokens, but rolls back, expands the retrieved context, and regenerates the token when the GDM indicates insufficient support.

generation difficulty metric that is later used to trigger context reallocation.

Input signals. At decoding step t , let $\ell_t \in \mathbb{R}^{\mathcal{V}}$ denote the output logits, where $\mathcal{V}$ denotes the size of the LLM vocabulary. We adopt the logit margin m_t as a lightweight confidence signal

$$m_t = \ell_t^{(1)} - \ell_t^{(2)}, \quad (2)$$

where $\ell_t^{(1)}$ and $\ell_t^{(2)}$ are the largest and second largest logits. Margin alone is not a reliable uncertainty proxy because semantically similar candidates may receive logits with similar values, resulting in a comparatively small logit margin during generation. To handle this issue, we extract a token semantic embedding from the inner states of the LLM. Following prior work on hidden-state analysis (Ma et al., 2025; Ferrando et al., 2025; Singh et al., 2024; Chen et al., 2024), we use the output of the attention sub-layer in the final transformer layer as emb_t to represent token-level semantic concepts.

Dual-encoder fusion with temporal modeling.

We fuse the semantic embedding emb_t and the logit margin m_t with a dual-encoder module, which preserves the complementary roles of semantic complexity and score-based confidence while aligning them in a shared space for joint decision making. Concretely, the two branches map both signals into a d -dimensional representation and yield

$$\mathbf{z}_t = f_e(\text{emb}_t) + f_m(m_t), \quad (3)$$

where f_e and f_m denote the semantic and logit margin encoders, respectively. Element-wise addition is applied to generate a fused representation $\mathbf{z}_t$. Since uncertainty can propagate across autoregressive steps and compound after an early mistake, we apply an LSTM to explicitly model this temporal accumulation and aggregate historical evidence

$$\mathbf{h}_t, \mathbf{c}_t = \text{LSTM}[\mathbf{z}_t, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}], \quad (4)$$

where the memory state $\mathbf{c}_t$ captures past uncertainty and the hidden state $\mathbf{h}_t$ serves as the step-wise generation difficulty representation for downstream prediction.

Output as generation difficulty metric. We define the *generation difficulty metric* (GDM) as a three-way probability vector over the token-generation scenarios shown in Figure 3, including grounded correct generation (Correct Answer), unknown-style abstention (Unknown), and unsupported content generation (Hallucination). This three-way design is central to UT-ACA because it separates reliable generation from two distinct failure modes under insufficient contextual evidence. The prediction head maps $\mathbf{h}_t$ to the GDM:

$$\begin{aligned} \mathbf{p}_t &= \text{softmax}(g_\phi(\mathbf{h}_t)), \\ \mathbf{p}_t &= [p_{t,\text{hal}}, p_{t,\text{cor}}, p_{t,\text{unk}}], \\ \mathbf{p}_t &\in \mathbb{R}^3, \quad \sum_{k \in [3]} p_{t,k} = 1, \end{aligned} \quad (5)$$




Figure 3: Token-generation scenarios for the Generation Difficulty Metric. The figure categorizes grounded answers, abstentions, and hallucinations by whether the required information is supported by the context or by the intrinsic knowledge of the model.

where g_ϕ denotes the prediction head, and $p_{t,\text{cor}}$, $p_{t,\text{unk}}$, and $p_{t,\text{hal}}$ correspond to correct, unknown-style abstention, and hallucination, respectively.

A binary formulation is insufficient for this setting because it merges *Unknown* and *Hallucination* into a single negative class, despite their different confidence patterns. Unknown tokens often correspond to confident abstention phrases and can exhibit high logit margins comparable to correct predictions, whereas hallucinated tokens introduce unsupported content. This mismatch introduces label noise and weakens the ability of the detector to identify when additional context is needed. In practice, the downstream policy triggers context re-allocation whenever the probability mass assigned to non-grounded states becomes dominant:

$$\max(p_{t,\text{hal}}, p_{t,\text{unk}}) > p_{t,\text{cor}}, \quad (6)$$

which expands the context window and regenerates the current token for rectification.

3.2 Adaptive Context Window

Inspired by InfLLM (Xiao et al., 2024a), we adopt fixed-size blocks as the basic units of our adaptive context window. During prefilling, we partition the KV representations of the long context into blocks that contain the same number of tokens. For each block, we keep a small set of representative tokens to support efficient relevance estimation. During decoding, the query vector at step t retrieves the top- k most relevant blocks, and attention is computed only over the selected blocks to mitigate long-context OOD behaviors.

Algorithm 1 UT-ACA Framework

Require: Model $\mathcal{M}$, Detector $\mathcal{D}$, Input x_0

```

1: Initialize cache  $\mathbf{C}$ , context window size  $k \leftarrow K_{\max}$ ,  $t \leftarrow 0$ 
2: while  $x_t \neq \langle \text{EOS} \rangle$  do
3:    $\mathcal{S} \leftarrow \text{Snapshot}(\mathbf{C})$  ▷ Save state
4:   // Step1: Tentative Generation
5:    $\hat{x}, \hat{\mathbf{C}} \leftarrow \mathcal{M}(x_t, \mathbf{C}, \text{topk} = k)$ 
6:   // Step2: Uncertainty Check
7:    $\text{is\_unsafe} \leftarrow \mathcal{D}(\hat{x}, \text{activations})$ 
8:   if  $\text{is\_unsafe}$  and  $k < K_{\max}$  then
9:      $\mathbf{C} \leftarrow \text{Restore}(\mathcal{S})$  ▷ Roll back
10:     $k \leftarrow K_{\max}$  ▷ Expand
11:    Regenerate:
12:     $x_{t+1}, \mathbf{C} \leftarrow \mathcal{M}(x_t, \mathbf{C}, \text{topk} = k)$ 
13:  else
14:     $x_{t+1}, \mathbf{C} \leftarrow \hat{x}, \hat{\mathbf{C}}$ 
15:     $k \leftarrow \text{UpdatePolicy}(k)$  ▷ Shrink
16:  end if
17:   $t \leftarrow t + 1$ 
18: end while
```

Since token difficulty varies across steps, the required contextual evidence is non-uniform. UT-ACA initialized at $K_{\max}$ for the first token to establish contextual grounding and subsequently adjusted by the update policy. At each step, the output logits and hidden-state embeddings are fed into the uncertainty detector to obtain the GDM. When either of the non-grounded probabilities exceeds the grounded probability, UT-ACA rolls back the tentative token, expands the context window by retrieving more relevant blocks, and regenerates the current token. Formally, let $\mathcal{B}$ be the set of all blocks, $\mathbf{r}(B)$ denote the representative keys of block B , and $\mathbf{q}_t$ be the decoding query at step t . Block retrieval under a budget k is written as

$$\mathcal{S}_t(k) = \arg \max_{\substack{\mathcal{S} \subseteq \mathcal{B} \\ |\mathcal{S}|=k}} \sum_{B \in \mathcal{S}} \text{sim}[\mathbf{q}_t, \mathbf{r}(B)], \quad (7)$$

where $\text{sim}[\cdot, \cdot]$ denotes the dot-product similarity between the decoding query and the representative keys of each block, and the expansion condition is formally defined as $\max(p_{t,\text{hal}}, p_{t,\text{unk}}) > p_{t,\text{cor}}$. Algorithm 1 presents the detailed pseudocode for our UT-ACA framework.

4 Experiments

We evaluate UT-ACA from three perspectives, including uncertainty detection, long-context generation, and cross-domain generalization. We first

Question	Which member of the kingsguard is known as “The White Bull”?		
Correct Answer	Ser Gerold Hightower		
Prediction	Ser Barristan Selmy		
F1-Score	0.3	✗	Concept-Acc-Score
			0 ✓

Figure 4: Example comparison between F1 score and Conceptual Accuracy Score (CAS). F1 score gives partial credit for surface overlap, while CAS assigns zero because the prediction is semantically incorrect.

describe the synthetic biography dataset and evaluation metrics, then compare the learned uncertainty detector with logit-based heuristic baselines. Finally, we integrate the detector into the adaptive context allocation framework and evaluate its efficiency-quality trade-off against fixed-budget context-management methods.

4.1 Dataset and Metrics

Dataset. We construct a synthetic biography summarization dataset using a summary-first pipeline, where each instance contains a randomly generated attribute value for a fictitious person and is expanded from a ground-truth summary into a longer biography. Person names are unique across splits to prevent data leakage. The dataset includes 10,000 training samples, a validation set with six unseen attributes and 20 summaries each, and 100 Wikipedia-extended long-context test biographies of 171k–400k tokens. For token-level supervision, we constrain the context window during decoding, record output logits and hidden-state embeddings, and label each generated token as correct, unknown, or hallucinated using GPT-OSS-120B (OpenAI, 2025) through a constrained comparison against the explicit reference fields, followed by manual verification. Further details on dataset construction, uncertainty-label generation, and representative examples are provided in Appendix A.1, Appendix A.2, and Appendix A.3, respectively.

Metrics. For the uncertainty detection, we report mean accuracy (**mAcc**) and per-class recall for uncertain and certain samples (**Recall_N** and **Recall_P**). For generation quality, GPT-OSS-120B scores conceptual alignment between model outputs and ground truth, reported as Conceptual Accuracy Score (**mAcc_{conc}**). An illustrative comparison is shown in Figure 4. Efficiency is quantified by mean context token consumption (**mTokens**) and

average per-token decoding latency (**mTime_{tok}**).

4.2 Uncertainty Detection Experiments

We train and evaluate token-level uncertainty detectors on logits and hidden-state embeddings extracted from Llama-3.1-8B-Instruct (Llama Team, 2024), with token labels assigned by the GPT-OSS-120B reader described above. As summarized in Table 1, we compare UT-ACA against four uncertainty estimation baselines. TOP-10 STD and LOGIT MARGIN are unsupervised logit-only heuristics, using the standard deviation of the top-10 logits and the top-1 versus top-2 margin, respectively. We also include two supervised evidential baselines, both formulated as binary uncertainty classifiers by merging *Unknown* and *Hallucination* into a single uncertain class. EVIDENTIAL MASS uses the Dirichlet uncertainty mass, whereas EVIDENTIAL uses the Dirichlet mean probability assigned to the uncertain class.

Methods	mAcc (%)	Recall _N (%)	Recall _P (%)
TOP-10 STD	57.43	35.66	73.19
Logit Margin	55.54	48.09	60.93
Evidential Mass	58.29	55.87	60.04
Evidential	70.85	74.45	68.25
UT-ACA	83.43	89.71	81.27

Table 1: Token-level uncertainty detection performance. *Unknown* and *Hallucination* are grouped as uncertain, while *Correct* is treated as certain. mAcc denotes classification accuracy; Recall_N and Recall_P denote recall on uncertain and certain tokens, respectively.

The logit-only heuristics provide weak and imbalanced token-level discrimination: TOP-10 STD obtains higher certain-token recall but misses most uncertain tokens, while LOGIT MARGIN improves uncertain-token recall but remains close to chance-level mAcc. The evidential baselines benefit from supervised hidden-state features, with EVIDENTIAL improving mAcc to 70.85% and uncertain-token recall to 74.45%. However, UT-ACA remains substantially stronger than all baselines, reaching 83.43% mAcc and 89.71% recall on uncertain tokens. This gap shows that UT-ACA learns a more reliable token-level signal for detecting when the compact context window is insufficient, while preserving high recall on certain tokens.

4.3 Generation Experiments

We evaluate the generation performance of UT-ACA with context management baselines, and we

formalize the context window update policies for UT-ACA as follows: *Update: Set.n* resets the context block budget to n for the next token, whereas *Update: Sub.n* decreases the current budget by n , with a minimum budget of 1.

Methods	Settings	mTokens ↓	mAcc _{conc} (%) ↑
InfLLM	$K = 2$	32	91.37
	$K = 3$	48	99.83
UT-ACA	Update: Set.1	25	96.59
	Update: Sub.1	29	99.08

Table 2: Generation performance on the validation set. UT-ACA is compared with fixed-budget InfLLM in terms of average context tokens (mTokens) and conceptual accuracy (mAcc_{conc}), showing that adaptive allocation preserves high accuracy with fewer tokens.

Validation set. We compare two post-expansion update rules, *Update: Set.1* and *Update: Sub.1*. All methods use a fixed block size of 16 tokens and a maximum block budget of $K_{\max} = 3$.

As shown in Table 2, UT-ACA achieves 99.08% mAcc_{conc} with only 29 context tokens on average, approaching InfLLM’s peak performance (99.83% at 48 tokens) while reducing context usage by approximately 40%. These results indicate that uncertainty-triggered expansion selectively allocates additional context without unnecessarily increasing the average context budget.

Long-context test set. We further evaluate UT-ACA on a long-context test set consisting of biographies with 171k-400k tokens, with an average length of 252k tokens. These inputs exceed the nominal context lengths of both Llama-3.1-8B-Instruct and Qwen2-7B-Instruct (Yang et al., 2024), providing a challenging setting for context-management methods. For each backbone, we train a separate uncertainty detector using the same procedure as above. Table 3 reports context usage, per-token decoding latency, and conceptual accuracy across compression-, token-selection-, and block-retrieval-based baselines. For token-selection meth-

Models	Methods	Settings	mTokens ↓	mTime _{tok} (s) ↓	mAcc _{conc} (%) ↑
Llama-3.1-8B-it	†RetrievalAttn (Liu et al., 2025)	$K_{\text{tok}} = 2048$	2k	0.021	15.25
	†LLMLingua-2 (Pan et al., 2024)	$K_{\text{com}} = 25\%$	32k	-	18.08
	†SnapKV (Li et al., 2024c)	$K_{\text{tok}} = 1024$	1k	0.210	23.43
	TokenSelect (Wu et al., 2025b)	$K_{\text{tok}} = 1024$	1k	-	39.43
	OmniKV (Hao et al., 2025)	$K_{\text{tok}} = 6.7\%$	8k	-	40.47
	†InfLLM-V2 (Zhao et al., 2026)	$K_{\text{tok}} = 1024$	1k	-	40.29
	Ada-KV (Feng et al., 2025)	$K_{\text{tok}} = 2048$	2k	-	48.89
	FastKV (Jo et al., 2026)	$K_{\text{tok}} = 4096$	4k	0.212	51.40
	FreqKV (Kai et al., 2026)	$K_{\text{tok}} = 2050$	2k	0.253	56.94
	H ₂ O (Zhang et al., 2023)	$K_{\text{tok}} = 1024$	1k	0.123	65.81
	InfLLM (Xiao et al., 2024a)	$K = 8$	128	0.064	25.81
		$K = 16$	256	0.072	51.23
		$K = 32$	512	0.080	72.26
	UT-ACA (Update: Sub.16)	$K_{\text{max}} = 32$	123	0.069	45.23
$K_{\text{max}} = 48$		218	0.075	62.17	
$K_{\text{max}} = 64$		344	0.082	70.48	
$K_{\text{max}} = 96$		498	0.097	74.61	
Qwen2-7B-it	InfLLM (Xiao et al., 2024a)	$K = 8$	128	0.061	37.77
		$K = 16$	256	0.064	44.39
		$K = 32$	512	0.077	54.72
	UT-ACA (Update: Sub.8)	$K_{\text{max}} = 16$	119	0.059	37.08
		$K_{\text{max}} = 32$	302	0.066	48.33
		$K_{\text{max}} = 48$	432	0.079	56.64

Table 3: Context-management performance on the long-context test set, where Update: Sub. N denotes the policy of decreasing the block budget by N after each accepted token, with a minimum of 1. The table compares token selection, compression, fixed-budget retrieval, and UT-ACA across two backbones using context usage (mTokens), per-token latency (mTime_{tok}), and conceptual accuracy (mAcc_{conc}). Methods marked with [†] exclude samples with out-of-memory errors or incoherent outputs.

ods, K_{tok} denotes the number of tokens selected per decoding step, while K_{com} denotes the compression ratio for LLMLingua-2. InfLLM-V2 follows its original block size of 64 with $K = 16$ retrieved blocks. For a direct comparison between InfLLM and UT-ACA, both methods use a block size of 16, where K denotes the fixed block budget of InfLLM and K_{max} denotes the maximum block budget available to UT-ACA.

As shown in Table 3, UT-ACA consistently achieves a favorable efficiency-quality trade-off across both backbones. On Llama-3.1-8B-Instruct, UT-ACA with $K_{\text{max}} = 96$ attains 74.61% $\text{mAcc}_{\text{conc}}$ at 498 tokens, outperforming the best InfLLM (72.26% at 512 tokens); at $K_{\text{max}} = 64$, it reaches 70.48% with only 344 tokens, a 33% reduction at a 1.78% accuracy cost. On Qwen2-7B-Instruct, UT-ACA ($K_{\text{max}} = 48$) similarly exceeds InfLLM’s peak (56.64% vs. 54.72%, 432 vs. 512 tokens), confirming the consistent efficiency-quality advantage of adaptive allocation. Given the test-set size of 100 biographies, we treat these differences as observed gains rather than claims of statistical significance; the primary evidence is the consistent quality–efficiency trade-off across operating points and backbones.

4.4 Quality–Latency–Memory Trade-off

We further evaluate the quality–latency–memory trade-off on the long-context test set using Llama-3.1-8B-Instruct. We compare fixed-budget InfLLM with $K \in \{8, 16, 32\}$ against UT-ACA with $K_{\text{max}} \in \{32, 48, 64, 96\}$. All operating points use the same backbone, GPU, prompts, and test set. End-to-end (E2E) latency combines long-context prefilling, decoding, uncertainty detection, roll-back, and regeneration costs, while peak GPU memory is measured over the complete inference process.

As shown in Table 4, UT-ACA provides substantially stronger quality–context trade-offs with modest E2E latency overhead. At comparable context usage, $K_{\text{max}} = 32$ achieves 45.23% $\text{mAcc}_{\text{conc}}$ with 123 tokens, compared with 25.81% for InfLLM at $K = 8$ with 128 tokens. At $K_{\text{max}} = 48$, UT-ACA improves accuracy from 51.23% to 62.17% relative to InfLLM at $K = 16$, while reducing attended context by 14.8% (218 vs. 256 tokens) with only 2.10 s additional latency. Similarly, $K_{\text{max}} = 64$ remains within 1.78 points of InfLLM at $K = 32$ while reducing context usage by 32.8% (344 vs. 512 tokens), with only 1.14 s additional

Method	$\text{mAcc}_{\text{conc}}(\%) \uparrow$	$\text{mTok.} \downarrow$	E2E T. (s) $\downarrow$	GPU (GiB)
InfLLM	25.81	128	95.80	26.42
	51.23	256	97.46	26.42
	72.26	512	99.56	26.44
UT-ACA	45.23	123	98.42	26.44
	62.17	218	99.56	26.47
	70.48	344	100.70	26.51
	74.61	498	102.98	26.56

Table 4: Quality–latency–memory trade-off on the long-context test set. InfLLM rows correspond to $K = \{8, 16, 32\}$ and UT-ACA rows to $K_{\text{max}} = \{32, 48, 64, 96\}$, ordered from top to bottom. mTok. , E2E T., and GPU denote selected context tokens, end-to-end latency, and peak GPU memory, respectively.

latency. The largest setting, $K_{\text{max}} = 96$, achieves the highest accuracy of 74.61%, exceeding InfLLM at $K = 32$ by 2.35 points while still using slightly fewer selected tokens. Peak GPU memory remains nearly unchanged (26.42–26.56 GiB), as UT-ACA primarily adapts attended evidence rather than model-residency memory. Overall, these results demonstrate that uncertainty-triggered context allocation improves the quality–context trade-off with bounded system overhead, even on our relatively short-output biography summarization task; longer-generation settings may provide further benefits by amortizing context-management costs and correcting uncertain tokens before errors propagate across subsequent decoding steps.

4.5 Ablation Study on Test Set

We conduct an ablation study on the test set using Llama-3.1-8B-Instruct with $K_{\text{max}} = 64$ under the *Update: Sub.16* configuration. The contribution of each component in the uncertainty detector is evaluated through systematic ablation by removing and combining the three components.

LogM	SE	LSTM	$\text{mAcc}(\%)$	$\text{mAcc}_{\text{conc}}(\%)$
✓	✗	✗	26.82	-
✗	✓	✗	78.95	69.58
✓	✗	✓	53.71	-
✗	✓	✓	80.06	68.58
✓	✓	✗	80.19	69.42
✓	✓	✓	83.43	70.48

Table 5: Component ablation of the uncertainty detector. The table evaluates the contribution of the Logit Margin branch (LogM), the Semantic Embedding branch (SE), and the LSTM module by removing them in different combinations. mAcc measures detector accuracy, and $\text{mAcc}_{\text{conc}}$ measures downstream conceptual generation accuracy.

The results are reported in Table 5. Removing the Semantic Embedding branch results in substantially degraded uncertainty detection accuracy, which is insufficient to support reliable context reallocation. Consequently, generation performance under these settings is not reported. Overall, the full model that incorporates all components achieves the strongest uncertainty detection performance and consistently yields the highest generation quality across the evaluated configurations. These results suggest that semantic representations and temporal aggregation provide complementary signals that improve the stability of uncertainty estimation and enhance the reliability of the reallocation trigger mechanism.

4.6 Out-of-Domain Generalization

As presented in Table 6, we further evaluate UT-ACA on three out-of-domain long-context benchmarks, including ∞ Bench (Zhang et al., 2024), LongBench (Bai et al., 2024), and RULER (Hsieh et al., 2024), to examine whether the learned uncertainty detector generalizes beyond the synthetic biography domain. The detector is directly reused from the version trained on our synthetic dataset, without any benchmark-specific retraining. We follow the official evaluation protocol of each benchmark and compare InfLLM with fixed context budgets K against UT-ACA with corresponding maximum budgets $K_{\max}$. For ∞ Bench, we use a block size of 128 with $K, K_{\max} \in \{8, 16, 32\}$, and report the benchmark-provided scores for the summarization and QA tasks. For LongBench, we use a block size of 16 with $K, K_{\max} \in \{4, 8, 16\}$, eval-

uating narrativeqa with normalized token-level F1 and qmsum with ROUGE-L $\times 100$. For RULER, we evaluate QA-16k and QA-8k using a block size of 16 with $K, K_{\max} \in \{8, 16, 32\}$, and report the official string-match accuracy, which measures whether any normalized gold answer appears in the model prediction.

Across the evaluated operating points, the frozen uncertainty detector yields competitive quality-token Pareto efficiency despite the substantial distribution shift across benchmarks. In particular, UT-ACA consistently operates with fewer selected tokens than the corresponding fixed-budget InfLLM setting, while the associated quality loss is generally small and is often fully recovered as the maximum budget increases. On ∞ Bench, UT-ACA remains close to InfLLM on both summarization and QA while using fewer selected tokens at every budget; at the largest budget, it achieves scores of 28.43 and 21.32 with 14.4% and 9.4% fewer tokens, respectively. The efficiency gains are more pronounced on LongBench: UT-ACA improves narrativeqa from 15.57 to 16.92 using 109 rather than 256 tokens, and obtains 19.18 on qmsum using only 102 tokens, compared with InfLLM’s 19.13 at 256 tokens, corresponding to reductions of 57.4% and 60.2%. On RULER, aggressive pruning causes larger degradation at smaller budgets, but UT-ACA recovers comparable performance at the largest budget while retaining substantially less context, achieving 21.65 versus 21.70 on QA-16k with 385 versus 512 tokens and 26.85 versus 26.75 on QA-8k with 352 versus 512 tokens.

Overall, without any benchmark-specific adap-

Method	∞ Bench				LongBench				RULER			
	<i>Sum Task</i>		<i>QA Task</i>		<i>narrativeqa</i>		<i>qmsum</i>		<i>QA-16k</i>		<i>QA-8k</i>	
	mTok↓	Score↑	mTok↓	Score↑	mTok↓	Score↑	mTok↓	Score↑	mTok↓	Score↑	mTok↓	Score↑
H ₂ O	1024	11.25	1024	1.14	64	5.35	64	13.84	128	12.00	128	14.30
	2048	11.52	2048	0.93	128	4.68	128	14.25	256	11.00	256	16.60
	4096	10.59	4096	1.27	256	4.44	256	14.81	512	10.70	512	21.00
InfLLM	1024	27.24	1024	16.27	64	14.95	64	18.40	128	14.05	128	14.90
	2048	27.87	2048	18.38	128	15.75	128	18.62	256	18.75	256	24.30
	4096	28.58	4096	22.09	256	15.57	256	19.13	512	21.70	512	26.75
UT-ACA	609	27.46	896	16.58	39	15.01	24	17.65	48	10.09	51	12.60
	1610	27.64	1664	18.72	59	16.11	47	18.66	105	14.65	108	14.40
	3508	28.43	3712	21.32	109	16.92	102	19.18	385	21.65	352	26.85

Table 6: Results on ∞ Bench, LongBench, and RULER using Llama-3.1-8B-Instruct. The table compares H₂O, InfLLM, and UT-ACA under the official evaluation protocols of the corresponding benchmarks, without retraining the UT-ACA uncertainty detector. mTok denotes the average number of selected tokens, Score follows the corresponding benchmark metric, and bold values indicate the best result in each column.

tation, UT-ACA provides competitive or non-dominated quality–efficiency operating points across all three benchmarks. These results suggest that the uncertainty detector transfers beyond its synthetic training domain and can support adaptive context allocation under substantially different task distributions, although the optimal degree of context reduction remains task dependent.

5 Related Work

Long-context inference research broadly aims to improve the robustness and efficiency of language models when operating over increasingly long inputs. One line applies training-based methods (Chen et al., 2025; Hu et al., 2025; Gao et al., 2025; Tian et al., 2025) or modified positional encodings (Ding et al., 2024; Jin et al., 2024; Wang et al., 2024) to extend length generalization. Another line focuses on context management to control the effective attention scope, reduce memory and computation, and suppress irrelevant context (Liskavets et al., 2025; Liu et al., 2025; Hooper et al., 2025; Fu et al., 2024). Since UT-ACA operates at inference time by regulating the usable context, we focus on the latter.

Context selection. These methods retain only the most relevant tokens or blocks to mitigate attention dilution (Zhang et al., 2025a; Xiao et al., 2024b). Training-free approaches such as InfLLM (Xiao et al., 2024a) and TokenSelect (Wu et al., 2025b) perform block- or token-level retrieval over the KV cache, while structured sparse attention methods such as XAttention (Xu et al., 2025) and OmniKV (Hao et al., 2025) design scoring-guided selection to balance efficiency and accuracy. InfLLM-V2 (Zhao et al., 2026) and FastKV (Jo et al., 2026) further push efficiency via trainable sparse attention and accelerated eviction strategies, respectively. Although several methods dynamically update which tokens participate in attention during generation, their retention or sparsity budgets are typically prescribed independently of the semantic reliability of the generated token.

Context compression. Rather than selecting existing entries, compression methods shrink the input or KV representations (Zhao et al., 2025a; Liao et al., 2025; Wang et al., 2025; Li et al., 2023). Prompt-level approaches such as LLM-Lingua (Jiang et al., 2023) and LongLLMLingua (Jiang et al., 2024) reduce token count using

information-theoretic criteria. At the KV-cache level, SnapKV (Li et al., 2024c), SCOPE (Wu et al., 2025a), and H₂O (Zhang et al., 2023) prune or evict entries based on attention patterns, while FocusLLM (Li et al., 2025b) condenses long inputs into compact representations. Recent methods allow KV entries to evolve dynamically during inference, but their update decisions are still primarily driven by attention statistics, cache importance, or pre-defined compression objectives. In contrast, their compression policy remains largely uniform across decoding steps, leaving token-wise generation difficulty unaddressed.

6 Conclusion

We propose UT-ACA, an uncertainty-triggered adaptive context allocation framework that dynamically adjusts the context window at the token level during decoding. Tokens are tentatively generated under a compact window and regenerated with expanded context only when the uncertainty detector signals insufficient evidence, enabling selective allocation without a fixed per-token budget. Experiments across synthetic and out-of-domain long-context benchmarks confirm that UT-ACA substantially reduces average context token consumption while maintaining competitive generation quality.

7 Limitations

The proposed UT-ACA significantly reduces the average number of context tokens, while the per-token decoding latency does not decrease proportionately. We attribute this phenomenon to the computational overhead introduced by the rollback mechanism. Uncertain tokens require regeneration after context expansion, leading to increased inference time, particularly when uncertainty is triggered frequently. Exploring alternative decoding strategies that reduce rollback frequency or amortize its cost remains an important direction for future research.

Acknowledgments

This work was partially supported by the NSFC General Program (62576367), the NSFC Excellent Young Scientists Fund (Overseas) (2025HY00260105 and 2025HYSPT0708), the Fundamental Research Funds for the Central Universities, Sun Yat-sen University (25hytd012), and the Multi-Agent Educational Large Model System and Platform Project (FPF10120260005).

References

- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, Jian-Guang Lou, and Weizhu Chen. 2024. [Make your llm fully utilize the context](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pages 62160–62188.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [LongBench: A bilingual, multi-task benchmark for long context understanding](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 3119–3137.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024. [INSIDE: LLMs’ internal states retain the power of hallucination detection](#). In *International Conference on Learning Representations*, volume 2024, pages 3056–3076.
- Jianghao Chen, Junhong Wu, Yangyifan Xu, and Jiajun Zhang. 2025. [LADM: Long-context training data selection with attention-based dependency measurement for LLMs](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 3076–3090.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [LongRoPE: Extending LLM context window beyond 2 million tokens](#). In *International Conference on Machine Learning (ICML)*, volume 235, pages 11091–11104. PMLR.
- Yuan Feng, Junlin Lv, Yukun Cao, Xike Xie, and S. Kevin Zhou. 2025. [Ada-kv: Optimizing kv cache eviction by adaptive budget allocation for efficient llm inference](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, Main Conference, pages 113152–113188.
- Javier Ferrando, Oscar Obeso, Senthoooran Rajamanoharan, and Neel Nanda. 2025. [Do i know this entity? knowledge awareness and hallucinations in language models](#). In *International Conference on Learning Representations (ICLR)*, volume 2025, pages 77965–78000.
- Qichen Fu, Minsik Cho, Thomas Merth, Sachin Mehta, Mohammad Rastegari, and Mahyar Najibi. 2024. [Lazyllm: Dynamic token pruning for efficient long context LLM inference](#). *CoRR*, abs/2407.14057.
- Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. 2025. [How to train long-context language models \(effectively\)](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 7376–7399.
- Jitai Hao, Yuke Zhu, Tian Wang, Jun Yu, Xin Xin, Bo Zheng, Zhaochun Ren, and Sheng Guo. 2025. [Omnikv: Dynamic context selection for efficient long-context llms](#). In *International Conference on Learning Representations (ICLR)*, volume 2025, pages 87443–87464.
- Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W. Mahoney, Yakun Sophia Shao, Kurt Keutzer, and Amir Gholami. 2024. [Kvquant: Towards 10 million context length llm inference with kv cache quantization](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pages 1270–1303.
- Coleman Richard Charles Hooper, Sehoon Kim, Hiva Mohammadzadeh, Monishwaran Maheswaran, Sebastian Zhao, June Paik, Michael W. Mahoney, Kurt Keutzer, and Amir Gholami. 2025. [Squeezed attention: Accelerating long context length LLM inference](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 32631–32652.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, Yang Zhang, and Boris Ginsburg. 2024. [RULER: what’s the real context size of your long-context language models?](#) *CoRR*, abs/2404.06654.
- Zhiyuan Hu, Yuliang Liu, Jinman Zhao, Suyuchen Wang, Yan Wang, Wei Shen, Qing Gu, Anh Tuan Luu, See-Kiong Ng, Zhiwei Jiang, and Bryan Hooi. 2025. [LongRecipe: Recipe for efficient long context generalization in large language models](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 11857–11870.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Yu, Xinying Song, and Denny Zhou. 2024. [Large language models cannot self-correct reasoning yet](#). In *International Conference on Learning Representations (ICLR)*, volume 2024, pages 32808–32824.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. [LLMLingua: Compressing prompts for accelerated inference of large language models](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 13358–13376.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2024. [LongLLMLingua: Accelerating and enhancing LLMs in long context scenarios via prompt compression](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1658–1677.
- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. 2024. [LLM maybe LongLM: SelfExtend LLM context window without tuning](#). In *International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pages 22099–22114. PMLR.
- Dongwon Jo, Jiwon Song, Yulhwa Kim, and Jae-Joon Kim. 2026. [FastKV: Decoupling of context reduction and KV cache compression for prefill-decoding acceleration](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 32167–32186.

- Jushi Kai, Yixuan Wang, Boyi Zeng, Haoli Bai, Bo Jiang, Ziwei He, and Zhouhan Lin. 2026. [Freqkv: Key-value compression in frequency domain for context window extension](#). In *International Conference on Learning Representations (ICLR)*, volume 2026, pages 105568–105587.
- Ammar Khairi, Daniel D’souza, Ye Shen, Julia Kreutzer, and Sara Hooker. 2025. [When life gives you samples: The benefits of scaling up inference compute for multilingual LLMs](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 27559–27583.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). In *International Conference on Learning Representations (ICLR)*, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024a. [LooGLE: Can long-context language models understand long contexts?](#) In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 16304–16333.
- Tianle Li, Ge Zhang, Quy Duc Do, Xiang Yue, and Wenhui Chen. 2024b. [Long-context LLMs struggle with long in-context learning](#). *Preprint*, arXiv:2404.02060.
- Yiwei Li, Jiayi Shi, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Yueqi Zhang, Ji Zhang, Chuyi Tan, Boyuan Pan, Yao Hu, and Kan Li. 2025a. [Speculative decoding for multi-sample inference](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 12523–12533.
- Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. 2023. [Compressing context to enhance inference efficiency of large language models](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6342–6353.
- Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. 2024c. [Snapkv: Llm knows what you are looking for before generation](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pages 22947–22970.
- Zhenyu Li, Yike Zhang, Tengyu Pan, Yutao Sun, Zhichao Duan, Junjie Fang, Rong Han, Zixuan Wang, and Jianyong Wang. 2025b. [FocusLLM: Precise understanding of long context by dynamic condensing](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 31087–31101.
- Zihan Liao, Jun Wang, Hang Yu, Lingxiao Wei, Jian-guo Li, Jun Wang, and Wei Zhang. 2025. [E2LLM: Encoder elongated large language models for long-context understanding and reasoning](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 19201–19230.
- Barys Liskavets, Maxim Ushakov, Shuvendu Roy, Mark Klibanov, Ali Etemad, and Shane K. Luke. 2025. [Prompt compression with context-aware sentence encoding for fast and improved LLM inference](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 39(23):24595–24604.
- Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Gholamreza Haffari, and Bohan Zhuang. 2024. [Minicache: Kv cache compression in depth dimension for large language models](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pages 139997–140031.
- Di Liu, Meng Chen, Baotong Lu, Huiqiang Jiang, Zhenhua Han, Qianxi Zhang, Qi Chen, Chengruidong Zhang, Bailu Ding, Kai Zhang, Chen Chen, Fan Yang, Yuqing Yang, and Lili Qiu. 2025. [Retrievalattention: Accelerating long-context llm inference via vector retrieval](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, Main Conference, pages 54358–54385.
- Llama Team. 2024. [The Llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Huan Ma, Jingdong Chen, Joey Tianyi Zhou, Guangyu Wang, and Changqing Zhang. 2025. [Estimating LLM uncertainty with evidence](#). *Preprint*, arXiv:2502.00290.
- OpenAI. 2025. [gpt-oss-120b & gpt-oss-20b model card](#). *CoRR*, abs/2508.10925.
- Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Dongmei Zhang. 2024. [LLMLingua-2: Data distillation for efficient and faithful task-agnostic prompt compression](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 963–981.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2024. [Yarn: Efficient context window extension of large language models](#). In *International Conference on Learning Representations (ICLR)*, volume 2024, pages 31932–31951.
- Yifu Qiu, Zheng Zhao, Yftah Ziser, Anna Korhonen, Edoardo Ponti, and Shay B. Cohen. 2024. [Are large language model temporally grounded?](#) In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 7064–7083.
- Aaditya K Singh, Ted Moskovitz, Felix Hill, Stephanie C.Y. Chan, and Andrew M Saxe. 2024. [What needs to go right for an induction head? A mechanistic study of in-context learning circuits and their formation](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pages 45637–45662. PMLR.
- Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. 2024. [QUEST: Query-aware sparsity for efficient long-context LLM inference](#). In *International Conference on Machine*

- Learning (ICML)*, volume 235, pages 47901–47911. PMLR.
- Junfeng Tian, Da Zheng, Yang Chen, Rui Wang, Colin Zhang, and Debing Zhang. 2025. [Untie the knots: An efficient data augmentation strategy for long-context pre-training in language models](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 1223–1242.
- Guangtao Wang, Shubhangi Upasani, Chen Wu, Darshan Gandhi, Jonathan Li, Changran Hu, Bo Li, and Urmish Thakker. 2025. [LLMs know what to drop: Self-attention guided KV cache eviction for efficient long-context inference](#). *CoRR*, abs/2503.08879.
- Suyuchen Wang, Ivan Kobyzev, Peng Lu, Mehdi Rezagholizadeh, and Bang Liu. 2024. [Resonance RoPE: Improving context length generalization of large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 586–598.
- Jialong Wu, Zhenglin Wang, Linhai Zhang, Yilong Lai, Yulan He, and Deyu Zhou. 2025a. [SCOPE: Optimizing key-value cache compression in long-context generation](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 10775–10790.
- Wei Wu, Zhuoshi Pan, Kun Fu, Chao Wang, Liyi Chen, Yunchu Bai, Tianfu Wang, Zheng Wang, and Hui Xiong. 2025b. [TokenSelect: Efficient long-context inference and length extrapolation for LLMs via dynamic token-level KV cache selection](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 21264–21281.
- Yuhao Wu, Ming Shan Hee, Zhiqiang Hu, and Roy Ka-Wei Lee. 2025c. [Longgenbench: Benchmarking long-form generation in long context llms](#). In *International Conference on Learning Representations (ICLR)*, volume 2025, pages 6851–6872.
- Chaojun Xiao, Pingle Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, and Maosong Sun. 2024a. [Infilm: Training-free long-context extrapolation for llms with an efficient context memory](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pages 119638–119661.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024b. [Efficient streaming language models with attention sinks](#). In *International Conference on Learning Representations (ICLR)*, volume 2024, pages 21875–21895.
- Ruyi Xu, Guangxuan Xiao, Haofeng Huang, Junxian Guo, and Song Han. 2025. [XAttention: Block sparse attention with antidiagonal scoring](#). In *International Conference on Machine Learning (ICML)*, volume 267, pages 69819–69831. PMLR.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 43 others. 2024. [Qwen2 technical report](#). *CoRR*, abs/2407.10671.
- Hailin Zhang, Xiaodong Ji, Yilin Chen, Fangcheng Fu, Xupeng Miao, Xiaonan Nie, Weipeng Chen, and Bin Cui. 2025a. [Pqcache: Product quantization-based kvcache for long context LLM inference](#). *Proc. ACM Manag. Data*, 3(3):201:1–201:30.
- Haozhen Zhang, Tao Feng, Pengrui Han, and Jiaxuan You. 2025b. [AcademicEval: Live long-context LLM benchmark](#). *Trans. Mach. Learn. Res.*, 2025.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024. [\$\infty\$ Bench: Extending long context evaluation beyond 100K tokens](#). In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 15262–15277.
- Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, Zhangyang "Atlas" Wang, and Beidi Chen. 2023. [H2o: Heavy-hitter oracle for efficient generative inference of large language models](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 34661–34710.
- Weilin Zhao, Zihan Zhou, Zhou su, Chaojun Xiao, Yuxuan Li, Yanghao Li, Yudi Zhang, Weilun Zhao, Zhen Li, Yuxiang Huang, Ao Sun, Xu Han, and Zhiyuan Liu. 2026. [Infilm-v2: Dense-sparse switchable attention for seamless short-to-long adaptation](#). In *International Conference on Learning Representations (ICLR)*, volume 2026, pages 69374–69390.
- Yunlong Zhao, Haoran Wu, and Bo Xu. 2025a. [Leveraging attention to effectively compress prompts for long-context LLMs](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 39(24):26048–26056.
- Zhe Zhao, Pengkun Wang, Haibin Wen, Shuang Wang, Liheng Yu, and Yang Wang. 2025b. [STEM-LTS: Integrating semantic-temporal dynamics in LLM-driven time series analysis](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 22858–22866.

A Dataset Construction Details

In this section, we detail the construction of the synthetic summaries and subsequent biographies. To provide insight into the data’s format and diversity, we also present the attribute names and representative examples. Additionally, we include the specific prompts used to generate naturalistic biographies.

A.1 Synthetic Dataset Construction Details

Our synthetic dataset is designed primarily for the biography summarization task. As illustrated in Figure 5, we employ a fixed pattern to construct the ground truth summaries. Specifically, each summary consists of a sequence with three components: the person name, the attribute name, and the attribute value. **It is important to note that all person names and attribute values are entirely fictitious. Therefore, there are no privacy concerns regarding real individuals.** To generate the full biography, the summary is extended with filler

content or paraphrased into a natural narrative. The dataset is partitioned into training, validation, and test sets, all constructed using synthetic summaries.



Figure 5: Standardized pattern of synthetic summaries. The figure illustrates the controlled person–attribute–value structure used to define ground-truth facts. This standardized structure keeps the ground truth unambiguous while allowing the surrounding biography text to vary in wording and length.

Training and Validation Set. For the training and validation data construction, each summary is associated with a single attribute, whereas the test set involves six different attributes. To ensure rigorous evaluation, the attributes used in the training set are distinct from those in the validation and test sets. Table 7 lists all attribute names used in our synthetic dataset. Specifically, we utilize

Dataset	Attribute Names	Example Data
Training Set	• Marriage Date • Job Title • Current City • Email Address • Phone Number • Favorite Color • User Agent • Credit Card Provider • Currency Used • Catch Phrase • Street Address • Vehicle License Plate • Favorite File Extension • Domain Name • Cryptocurrency • Timezone • Isbn Code • Lucky Number • Hobby • Marital Status • Personality Type	• The marriage date of Laetitia Guyot de Paris is 2010-03-29. • The job title of Franck Jourdan is Magazine features editor. • The current city of Marcial Escalona Yuste is Augerdan. • The email address of Jeremy Wright is catherine40@example.com. • The phone number of Clarence Byrd is 07703087971. • The favorite color of Letizia Fagotto is Black. • The user agent of Inés Novoa Sarabia is Mozilla/5.0 • The credit card provider of Douglas Baker is Maestro. • The currency used of Susann Jacobi Jäckel is Bahraini dinar. • The catch phrase of Olga Torre Rodriguez is Interfaccia estesa ... • The street address of Gabriela Jesus is Bährgasse 62. • The vehicle license plate of Erdal Mälzer is GE 3041 CB. • The favorite file extension of Ornella Tomasini is tiff. • The domain name of Théo Macedo is chretien.com. • The cryptocurrency of Anaïs Gilbert du Rousset is Vertcoin. • The timezone of Eutimio Tejada Ibañez is Europe/Paris. • The isbn code of Jeanne-Sabine Brunet is 978-0-456-54959-9. • The lucky number of Zoé Chauvet is 61. • The hobby of María Manuela Guardia Moles is Sleeping. • The marital status of Gustavo Vallés is Widowed. • The personality type of Salvi Toscanini-Ferrabosco is Extrovert.
Validation Set and Test Set	• Birth Date • Birth Place • Company • Major • University • Work Place	• The birth date of Olivia Garcia is March 22, 1985. • The birth place of Daniel Lee is Seoul, South Korea. • The company of Olivia Chen is Microsoft. • The major of Ethan Clark is Computer Science. • The university of Aiden Gomez is Stanford University. • The work place of Elijah Robinson is Singapore, Singapore.

Table 7: Attribute sets and example facts in the synthetic dataset. The table lists representative attribute names and example facts for the training, validation, and test splits. Training uses diverse single-attribute facts, while validation and test focus on six held-out biography attributes.

Wikipedia Source Articles				
Architecture	Mathematics	Pacific Ocean	Thermodynamics	Blockchain
Art	Number	Sahara	Virus	Cryptography
Baroque	Sport	South America	Communism	Engineering
Cinema	Vaccine	Volcano	Culture	Internal combustion engine
Classical music	French Revolution	Age of Enlightenment	Democracy	Internet
Impressionism	Maya civilization	American Civil War	Epistemology	Nanotechnology
Jazz	Roman Empire	Ancient Egypt	Human rights	Robotics
Literature	Africa	Ancient Greece	Law	Semiconductor
Modernism	Amazon River	British Empire	Logic	Software
Mona Lisa	Amazon rainforest	Bacteria	Metaphysics	Steam engine
Painting	Antarctica	Biology	Philosophy	Technology
Photography	Atlantic Ocean	Black hole	Psychology	Telecommunication
Poetry	Australia (continent)	Cell (biology)	Republic	Agriculture
Rock music	Earth	Climate change	Socialism	Calendar
Sculpture	Europe	DNA	Ottoman Empire	Food
The Starry Night	Great Barrier Reef	Evolution	World War II	Light
Theatre	Nile	General relativity	Algorithm	Medicine
Geometry	Ocean	Quantum mechanics	Artificial intelligence	Olympics

Table 8: Wikipedia source articles used for test-set context extension. These passages are inserted around synthetic biographies to create long distractor contexts, testing whether methods can retrieve relevant biography facts from long, heterogeneous inputs.

21 attributes for training, and reserve 6 distinct attributes for validation and testing. Additionally, we provide examples for each attribute to demonstrate the diversity of the data.

Test Set. We employ GPT-OSS-120B to generate naturalistic biographies based on the summaries provided in the test set. It is important to note that this generation process is model-agnostic; therefore, any capable generative model can be utilized for this task. The specific prompt used to construct the biographies is shown in Figure 6 and can be readily copied and applied. After the model generates naturalistic biographies, we insert common knowledge texts from Wikipedia to further extend the length of synthetic biographies for test set. As shown in Table 8, we present the Wikipedia files used in our test set extending process. The original link of these data can be found in <https://en.wikipedia.org/wiki>.

A.2 Uncertainty Label Construction and Verification

The synthetic uncertainty labels are constructed under a controlled reference-based protocol. As described in Appendix A.1, each synthetic instance is associated with an explicit ground-truth fact represented by a person name, attribute name, and attribute value. This structure provides an unam-

biguous reference for determining whether generated content is supported by the available evidence.

For each generation trajectory, token-level labels are assigned according to three generation states. Tokens consistent with the reference fact are labeled *Correct*; tokens that explicitly abstain from providing the requested information are labeled *Unknown*; and tokens that introduce factual content unsupported by the reference are labeled *Hallucination*. For template-defined tokens whose correctness is deterministic from the construction process, labels are assigned directly. For factual output tokens, GPT-OSS-120B performs a constrained comparison between the generated response and the corresponding reference fields rather than an unconstrained factuality judgment. The complete labeling prompt is shown in Figure 7.

The resulting labels are manually verified to identify incorrect assignments and ambiguous cases. Because the underlying facts are explicitly specified during data construction, verification primarily checks consistency between the generated tokens and the reference fields rather than requiring subjective assessment of open-ended factuality. This controlled labeling procedure reduces ambiguity in the supervision used to train the uncertainty detector.

```
{
  "role": "system",
  "content": (
    "You are to write a detailed character biography based on the provided information, and you must use the full name of this person all the time!!\n\n"
    "##You will be provided with the following details##: Name: [Person's full name] Birth Date: [Date of birth] Birth Place: [Location where the person was born] University: [Educational institution attended] Major: [Field of study/specialization] Work Place: [Current or primary workplace] Company: [Organization/company where they work]\n\n"
    "##For long biographies (1000+ words)##: Open with engaging anecdote or compelling fact. Describe how birthplace/childhood shaped their interests. Detail their university experience and academic achievements. Explain career decisions and transitions. Include specific projects, innovations, or contributions. Mention professional philosophy or approach. Add personal interests or community involvement if word count allows. Conclude with future goals or ongoing projects.\n\n"
    "##Important Notes##: Accuracy: Ensure dates and locations are logically consistent. Originality: Create unique content for each biography; avoid repetitive phrases. Cultural sensitivity: Be respectful of all locations and institutions mentioned. Professional standards: Write as if this will be published in a professional context."
    "Template: The six information must be described in following template, The [Attribute] of [Name] is [Value].\n\n"
    "You must follow the example following, and XXX is the place you can write on your own: The birth date of Bob Dylan is December 29, 1880. XXX. The birth place of Bob Dylan is William Hayes. XXX. The university of Bob Dylan is Yale University. XXX. The major of Bob Dylan is computer science. XXX. The company of Bob Dylan is Google. XXX. The work place of Bob Dylan is Shenzhen, China. XXX."
  )
}
```

Figure 6: Prompt template for generating naturalistic biographies. The figure shows the LLM instruction used to convert structured ground-truth facts into fluent biography passages. The prompt asks the LLM to preserve the provided summary facts while embedding them in natural narrative text.

```
# Task
Given the ground-truth subject ('name'), answer ('answer'), queried attribute ('question'), and a tokenized model response ('prediction'), assign a token-level label to every token in 'prediction'.

# Labels
- 1 (Correct): The token is consistent with the ground truth or is a functional token that does not introduce incorrect factual content.
- 0 (Hallucination): The token expresses an incorrect predicate, entity, or value that contradicts the ground truth.
- 2 (Unknown): The token is part of an explicit abstention or statement indicating that the requested information is unknown or unavailable.

# Labeling Rules
Apply the following rules in priority order:
1. Exact value/name match: Any correctly spelled token belonging to 'answer' or 'name' must be labeled 1, regardless of errors elsewhere in the sentence.
2. Functional tokens: Function words, punctuation, and formatting tokens are labeled 1 unless they directly express an incorrect factual relation.
3. Predicate mismatch: Tokens describing a relation or attribute inconsistent with 'question' are labeled 0.
4. Value/entity mismatch: Tokens introducing an entity or value inconsistent with 'name' or 'answer' are labeled 0.
5. Unknown state: Tokens expressing explicit uncertainty, abstention, or requests for missing information are labeled 2.

# Output
Return only a valid JSON object: {"labels": "[l_1, l_2, ..., l_n]"} The number of labels must exactly equal the number of tokens in 'prediction'.
```

Figure 7: Prompt for reference-constrained token-level uncertainty labeling. GPT-OSS-120B assigns each prediction token as *Correct* (1), *Hallucination* (0), or *Unknown* (2) by comparing it with the explicit ground-truth subject, attribute, and answer under prioritized labeling rules.

A.3 Data Examples from Our Dataset

In our experiments, the training set is utilized to train the uncertainty detector; consequently, the data structure and sequence lengths in this set are relatively simple and short. The validation set is designed to preliminarily evaluate the trained uncertainty detector and the downstream adaptive context window. While the validation data retains the same format as the training set, it comprises distinct attributes and person names. Finally, the test set is constructed to assess our method in a long-context setting. To achieve this, we insert extensive passages from Wikipedia to extend biography lengths, thereby challenging the long-context handling capabilities of the methods.

Table 9 presents examples from our synthetic

dataset. While the training and validation sets involve processing only a single attribute per instance, the test set requires the simultaneous extraction of six attributes. Furthermore, since the biographies in the test set are composed in natural language, the model must identify and extract multiple attributes from natural-language context to derive the correct answers.

B Implementation Details

This section first introduces training details and the architecture of our uncertainty detector. Subsequently, we describe the methodology for extracting LLM internal states used in this work. Finally, we detail the selection of network hyperparameters used in our approach.

Dataset	Example Summarization	Example Biography
Training Set	The email address of Jeremy Wright is catherine40@example.com.	The sky is really blue. ... The email address of Jeremy Wright is catherine40@example.com. ... The sky is really blue.
Validation Set	The birth date of Olivia Garcia is March 22, 1985.	The sky is really blue. ... The birth date of Olivia Garcia is March 22, 1985. ... The sky is really blue.
Test Set	- The birth date of Emma Thompson is 12/29/1880. - The birth place of Emma Thompson is Tokyo, Japan. - The university of Emma Thompson is Harvard University. - The major of Emma Thompson is Computer Science. - The company of Emma Thompson is Google. - The work place of Emma Thompson is San Francisco Bay Area.	... Emma Thompson was born on December 29, 1880, in the bustling metropolis of Tokyo, Japan. ... She applied to Harvard University and was accepted into the Computer Science program. ... After careful consideration, she chose to join Google, a company known for its innovative culture and cutting-edge technology. ... Emma Thompson’s journey from Tokyo to the San Francisco Bay Area is a testament to her unwavering determination and exceptional talent. ...

Table 9: Representative examples from the training, validation, and test splits. The table shows how the three splits differ in attribute complexity and context length. Training and validation examples contain one explicit fact in a short template-like passage, whereas test examples contain six biography attributes expressed in natural language and surrounded by long-context distractors.

B.1 Training Details

We implement UT-ACA in Python 3.11 using PyTorch 2.4. The detector is trained on a single NVIDIA H100 GPU. We conduct eight independent runs, each requiring approximately 1.5 GPU-hours. Each run performs 10,000 optimization steps with a batch size of 16.

We optimize the detector using Adam with an initial learning rate of 10^{-3} and no weight decay. The learning rate follows cosine annealing over the 10,000 training steps, with the minimum learning rate set to zero. We use token-level weighted cross-entropy loss and exclude padded token positions from the loss. The class weights are $[1.0, 0.05, 1.0]$ for the three-class detector, corresponding to hallucinated, correct, and unknown tokens, respectively. For the binary ablation, unknown tokens are merged into the uncertain class and the corresponding weights are $[1.0, 0.05]$.

We evaluate the detector every 100 optimization steps and retain the checkpoint with the highest validation F1 score. We do not apply early stopping: optimization continues for all 10,000 steps, while checkpoint selection is performed retrospectively using validation F1.

B.2 Architecture of Uncertainty Detector

The proposed model, termed *Uncertainty Detector*, is designed to process sequential data consisting of high-dimensional semantic feature vectors and scalar confidence scores. The architecture comprises three primary modules: dual-branch feature

encoding, temporal modeling, and a residual classification head.

Input Encoding and Fusion At each time step t , the model accepts two input streams: a feature vector $\text{emb}_t \in \mathbb{R}^{D_{\text{in}}}$ (where $D_{\text{in}} = 4096$ for Llama-3.1-8B) and a scalar score $m_t \in \mathbb{R}$. Before entering the semantic encoder, each feature vector is L2-normalized along its feature dimension. The scalar m_t represents the logit margin, calculated by subtracting the top-2 logit from the top-1 logit of the LLM output. The vector input is processed by the VecEncoder, which projects the high-dimensional input to a latent space of dimension $d_{\text{model}} = 64$. This is achieved via a linear transformation followed by a GELU activation and dropout

$$\mathbf{h}_{\text{vec}}^{(t)} = \text{Dropout}(\text{GELU}(\mathbf{W}_e \text{emb}_t + \mathbf{b}_e)) \quad (8)$$

where dropout rate is set to 0.5. Simultaneously, the scalar score is processed by the ScoreEncoder, a Multi-Layer Perceptron (MLP) aimed at up-sampling the scalar to match d_{model} . The MLP consists of two linear layers with an intermediate dimension of 32, utilizing GELU activations and dropout ($p = 0.1$).

The ScoreEncoder acts as the uncertainty signal branch, employing the logit margin as a lightweight indicator to gauge the generation difficulty of tokens. Conversely, the VecEncoder serves as the semantic feature branch, utilizing the hidden states of the LLM to extract conceptual information for enhanced uncertainty estimation. The dual-encoder

module aligns the dimensions of these two information sources, after which the features are fused via element-wise addition to form a joint representation ($\mathbf{z}_t$).

Temporal Modeling To effectively model the accumulation of uncertainty throughout the generation process, we employ a temporal modeling approach designed to capture the sequential dependencies between the current token and its predecessors. Specifically, the fused feature sequence $\mathbf{U} = \{\mathbf{z}_1, \dots, \mathbf{z}_T\}$ serves as input to a standard Long Short-Term Memory (LSTM) layer. Configured with a hidden size of d_{model} , the LSTM recursively processes the sequence according to LSTM dynamics.

For the training phase of the uncertainty detector, we preserve the entire sequence of hidden states as supervision signals. In contrast, during the LLM inference stage, only the final hidden state $\mathbf{h}_T$ is extracted to predict the category of the current token, as it aggregates the information of the complete sequence.

Residual Classification Head The output from the LSTM passes through a stack of $N = 3$ ResidualBlocks. Each block implements a bottleneck structure with an expansion factor of 2. For an input $\mathbf{h}$, the block operation is defined as

$$\begin{aligned} \hat{\mathbf{h}} &= \text{Dropout}(\text{GELU}(\mathbf{W}_1 \mathbf{h} + \mathbf{b}_1)) \\ \mathbf{h}_{out} &= \mathbf{h} + \text{Dropout}(\mathbf{W}_2 \hat{\mathbf{h}} + \mathbf{b}_2) \end{aligned} \quad (9)$$

Finally, a linear projection maps the output of the last residual block to three class logits, followed by a softmax function to obtain the GDM probabilities $\mathbf{p}_t = [p_{t,hal}, p_{t,cor}, p_{t,unk}]$.

Hyperparameters The specific configurations for our uncertainty detector are detailed in Table 10. Notably, while most hyperparameters remain consistent across different LLMs, certain settings are model-specific. For instance, the feature vector dimension is inherited directly from the evaluated LLM; thus, it is set to 3584 for Qwen2-7B. All other internal parameters of the uncertainty detector are kept uniform to ensure comparability.

B.3 Inner State Extraction of LLMs

To train the uncertainty detector, we leverage internal embeddings from the LLM to capture token-level semantic complexity. Following established protocols, we extract features from the final layer of

Component	Parameter	Value
Input Dimensions	Vector Dim (D_{in})	4096
	Score Dim	1
	Hidden Model Dim (d_{model})	64
Encoders	Vec Dropout	0.5
	Score MLP Inner Dim	32
Residual Head	Number of Blocks	3
	Expansion Factor	2
	Dropout	0.1
Output	Number of Classes	3

Table 10: Hyperparameter configuration of the Llama-3.1-8B uncertainty detector. The table reports the input dimensions, encoder settings, residual-head configuration, and number of output classes. The detector projects high-dimensional LLM features into a compact 64-dimensional hidden space before classification.




Figure 8: Extraction points for semantic embeddings in the final transformer layer. The figure marks the attention output, MLP output, and residual stream in the final layer (e.g., Layer 32 of Llama-3.1-8B). These positions are evaluated as candidate internal states for token-level semantic features.

LLMs. As illustrated in Figure 8, we consider three primary extraction points: (1) the attention mechanism output, (2) the Multi-Layer Perceptron (MLP) block output, and (3) the final residual stream output. We evaluate the efficacy of each position by training independent detectors for comparison.

Our results indicate that the extraction site has a negligible impact on detection performance. This consistency stems from the additive nature of the residual stream, where the transformations at these sub-layer stages act as incremental refinements to the same underlying representation, which may explain the limited sensitivity for uncertainty detection. Based on these results, we adopt the attention output (option 1) as the default extraction point in all experiments, as it is the most directly accessible in standard implementations.

C Evaluation Metrics Details

In this section, we detail the implementation of the Conceptual Accuracy Score (CAS) proposed in this study. We then evaluate the reliability of CAS within the context of biography summarization by comparing it against existing evaluation methods and human annotations. Our human study provides preliminary evidence that CAS can capture conceptual correctness in this controlled biography summarization setting. Given the limited sample size, we treat this analysis as a sanity check rather than definitive validation of LLM-judge reliability.

C.1 Details of Conceptual Accuracy Score

In our work, we employ GPT-OSS-120B, which can theoretically be replaced by any sufficiently capable generative model, to estimate the Conceptual Accuracy Score (CAS). CAS is computed by comparing the LLM-generated output with the ground-truth summary under the specified attributes. The prompt used for this estimation is shown in Figure 9. The model takes as input the predictions of the evaluated LLMs, together with the corresponding questions and ground-truth answers, and outputs a floating-point score between 0 and 1, indicating the degree of accuracy of the LLM responses.

C.2 Evaluations of Conceptual Accuracy Score

To provide a preliminary assessment of the proposed CAS, we conduct a human annotation study on 30 samples drawn from our test-set evaluation. Specifically, human annotators are asked to evaluate 30 samples drawn from the real evaluation workload in our test set experiments. The annotators receive the same prompt used with GPT-OSS-120B and are instructed to assign a floating-point score independently. We then compare the results

obtained from three evaluation methods: string-based matching, human evaluation, and CAS estimated by GPT-OSS-120B. As each test sample requires the extraction of six attributes, we report both the average accuracy scores under the three evaluation settings and the number of correctly identified attributes, as determined by the human annotators.

Quantitative Evaluation. We randomly select 30 samples from our test set to conduct a comparative analysis. The experimental setup is as follows: we utilize Llama-3.1-8B-Instruct as the model to be evaluated and employ our UT-ACA as the context management method for long-context inference. The specific parameters of the UT-ACA are set to a fixed block size of 16 tokens, with a maximum budget of 96 blocks for the adaptive context allocation. The update policy involves subtracting 16 blocks after the generation of tokens with certainty, maintaining a minimum window size of 1 block. As shown in Table 11, CAS obtains an average score of 0.746, compared with 0.766 from human evaluation, corresponding to an absolute difference of 0.02.

Qualitative Evaluation. Table 12 presents representative examples illustrating the behavior of CAS under surface-form variation and genuine semantic errors. Specifically, in the evaluation samples, CAS can automatically detect formatting differences while correctly capturing the underlying semantic correctness. For example, in **Case #3**, the LLM predicts “1888-06-03” whereas the ground truth is “June 3, 1888”. A string-matching baseline would judge this prediction as incorrect (or assign it an artificially low score) due solely to the surface-form discrepancy. In contrast, CAS correctly assigns a score of 1. In **Case #17**, CAS

```
{
  "role": "system",
  "content": (
    "You are an expert in evaluating the accuracy of large language model predictions."
    "Your task is to compare a provided \"Answer\" with a \"Prediction\" for a given \"Question\" and assign an Accuracy Score between 0.0 and 1.0 (inclusive).\"
    \"1.0 Score: The prediction is identical to, or a complete and accurate restatement of, the answer.\"
    \"0.5 Score: The prediction contains partially correct information, or is conceptually related but not fully accurate.\"
    \"0.0 Score: The prediction is completely incorrect, irrelevant, or includes hallucinated information not present in the answer or implied by the question.\"
  )
}
```

Figure 9: Prompt template for evaluating Conceptual Accuracy Score (CAS). The figure shows the LLM-judge instruction used to score semantic correctness rather than surface overlap. For each subtask, the prompt provides the question, the model prediction, and the ground-truth answer, then requests a floating-point score between 0 and 1 that reflects conceptual accuracy.

ID	Name	Str	CAS	Hum	Cor	ID	Name	Str	CAS	Hum	Cor
1	Emma Thompson	.000	.000	.000	0/6	16	Lucas Scott	.262	1.00	1.00	6/6
2	Liam Carter	.127	.167	.167	1/6	17	Harper Coleman	.114	.167	.167	1/6
3	Olivia Bennett	.230	1.00	1.00	6/6	18	Henry Green	.260	.500	.667	3/6
4	Noah Wright	.278	.830	.833	5/6	19	Evelyn Perry	.308	1.00	1.00	6/6
5	Ava Sullivan	.299	.667	.667	4/6	20	Oliver Hughes	.358	1.00	1.00	6/6
6	Ethan Brooks	.364	.833	.833	5/6	21	Abigail Kelly	.333	.833	.833	5/6
7	Sophia Reed	.341	.830	1.00	6/6	22	Jack Ross	.268	.830	.833	5/6
8	Mason Foster	.389	1.00	1.00	6/6	23	Emily Howard	.355	1.00	1.00	6/6
9	Isabella Parker	.246	.830	.833	5/6	24	Daniel Gray	.311	.830	.833	5/6
10	William Hayes	.231	.500	.500	3/6	25	Sofia Ellis	.159	.330	.417	2/6
11	Mia Evans	.181	.830	.830	5/6	26	Matthew Bell	.379	1.00	1.00	6/6
12	James Turner	.385	.833	.830	5/6	27	Aria Wood	.417	1.00	1.00	6/6
13	Amelia Fisher	.227	.833	.833	5/6	28	Jacob King	.271	.830	.917	5/6
14	Benjamin Ward	.350	.500	.500	3/6	29	Scarlett Price	.242	.750	.750	4/6
15	Charlotte Morgan	.323	.830	.830	5/6	30	Logan Lee	.361	.830	.917	5/6
Avg (1–30):		Str=0.279,	CAS=0.746,	Hum=0.766,	Cor=4.5/6						

Table 11: Preliminary human comparison of Conceptual Accuracy Score on biography extraction. The table compares string matching (**Str**), conceptual accuracy score (**CAS**), and human annotation (**Hum**) on 30 test samples. **Cor** reports the number of correctly extracted attributes out of six according to human annotators.

Case #3: Olivia Bennett (Str: 0.230, LLM: 1.000, Hum: 1.000)				Case #17: Harper Coleman (Str: 0.114, LLM: 0.167, Hum: 0.167)			
Attr	GT	Pred.	Hum	Attr	GT	Pred.	Hum
Birth Date	June 3, 1888	1888-06-03	✓	Birth Date	10 Nov, 1907	1976	✗
Birth Place	Sydney, NSW	Sydney, NSW	✓	Birth Place	Mexico City	Madrid, Spain	✗
University	MIT	Mass. Inst. Tech.	✓	University	NUS	Complutense Madrid	✗
Major	Medicine	Medicine	✓	Major	Intl. Rel.	Chem. Eng.	✗
Company	Mayo Clinic	Mayo Clinic	✓	Company	UN	UN	✓
Work Place	Beijing, CN	Mayo’s Beijing	✓	Work Place	Seattle, WA	Pacific NW	✗

Table 12: Case study of CAS under format variation and semantic errors. Case #3 shows that date and entity-format differences can still be semantically correct, whereas Case #17 contains wrong attribute values that should receive low scores. ✓ and ✗ indicate human judgments for each attribute.

accurately identifies genuine semantic errors and appropriately assigns a low score of 0.167.

Importantly, CAS is used only as a post-hoc measure of generation quality and does not participate in uncertainty detection, context allocation, or decoding. It therefore does not affect the reported context consumption or latency. Moreover, our out-of-domain experiments on ∞ Bench, LongBench, and RULER follow the official evaluation metrics of the corresponding benchmarks rather than CAS.

D Additional Results and Analysis

D.1 Ablation Study on Two-Way vs. Three-Way Difficulty Metric

To examine the contribution of the proposed three-way difficulty metric, we conduct an ablation study in which the *Unknown* and *Hallucination* categories are merged into a single class, yielding a two-way difficulty metric. Under this setting,

we perform eight independent training runs using the same experimental protocol. For each run, the checkpoint achieving the highest F1 score on the validation set is retained. The selected detectors are then evaluated on the test set with the following configuration: MaxBudget = 64, block_size = 16, and SubPolicy = 16.

As shown in Table 13, merging *Unknown* and *Hallucination* into a single category leads to consistent performance degradation across both detection and downstream generation metrics. At the detector level, mean accuracy (mAcc) decreases by 0.69 points (83.43% $\rightarrow$ 82.74%), and the recall for negative samples drops by 1.67 points (89.71% $\rightarrow$ 88.04%). More critically, the degradation propagates to the downstream generation stage: conceptual generation accuracy (mAcc_{conc}) declines substantially by 2.49 points (70.48% $\rightarrow$ 67.99%), while average context token consumption remains

Setting	mAcc (%) $\uparrow$	Recall _N (%) $\uparrow$	Recall _P (%) $\uparrow$	mTokens $\downarrow$	mAcc _{conc} (%) $\uparrow$
Two-Way Metric	82.74 ($\downarrow 0.69$)	88.04 ($\downarrow 1.67$)	80.60 ($\downarrow 0.67$)	348	67.99 ($\downarrow 2.49$)
Three-Way Metric	83.43	89.71	81.27	344	70.48

Table 13: Ablation study of two-way versus three-way difficulty metrics. The two-way metric merges Unknown and Hallucination into one negative class, while the three-way metric keeps them distinct. mAcc denotes detector accuracy, Recall_N and Recall_P denote recall on negative and positive samples, mTokens denotes average context token consumption, and mAcc_{conc} denotes conceptual generation accuracy.

nearly unchanged (344 vs. 348 tokens), confirming that the gain is attributable to improved detection quality rather than increased token expenditure.

This degradation stems from the distinct generative regimes underlying *Unknown* and *Hallucination* responses. Under compact context windows, the model tends to produce *Unknown* responses with disproportionately high confidence, whose logit margins are closer to correct outputs than to hallucinated ones. Merging them into a single negative class therefore conflates heterogeneous behaviors, introducing inconsistent supervision signals that increase label ambiguity and destabilize training.

D.2 Calibration on Downstream Long-Context Tasks

Our out-of-domain experiments demonstrate that the uncertainty detector trained on the synthetic biography data can transfer to RULER without benchmark-specific retraining. We further examine whether a small amount of downstream calibration can improve its context-allocation decisions while preserving this transferable initialization. Specifically, we construct 900 calibration trajectories with QA-4k data split from RULER, using a frozen Llama-3.1-8B-Instruct reader under complete-, partial-, and missing-evidence conditions, targeting *Correct*, *Hallucination*, and *Unknown* responses, respectively. GPT-OSS-120B then assigns token-level labels following the same three-way supervision scheme used for the original uncertainty detector. Starting from the synthetic-pretrained detector, we perform limited calibration on these trajectories and evaluate the resulting detector on the end-to-end RULER tasks.

As shown in Table 14, limited calibration consistently improves generation quality under the compact topk8/sub4 setting while also reducing context usage. On QA-16k, the RULER score increases from 14.65 to 19.30 while mTok decreases from 105.00 to 101.77. On QA-8k, calibration

Method	Setting	QA-16k		QA-8k	
		mTok $\downarrow$	Score $\uparrow$	mTok $\downarrow$	Score $\uparrow$
Original	topk8/sub4	105.00	14.65	108.00	14.40
Calibration	topk8/sub4	101.77	19.30	102.68	23.10
Original	topk32/sub16	385.00	21.65	352.00	26.85
Calibration	topk32/sub16	364.15	26.10	367.24	26.90

Table 14: Effect of limited downstream calibration on RULER. ORIGINAL denotes the uncertainty detector trained only on the synthetic biography data, while CALIBRATION denotes the same detector after calibration on 900 RULER trajectories. mTok denotes the average number of selected context tokens, and Score follows the official RULER string-match accuracy.

yields an even larger improvement from 14.40 to 23.10, together with a reduction in mTok from 108.00 to 102.68. These results indicate that a small amount of task-specific supervision can substantially improve the detector’s allocation decisions in the more aggressively compressed regime. The benefit remains evident with the larger topk32/sub16 setting. On QA-16k, calibration improves the score from 21.65 to 26.10 while reducing mTok from 385.00 to 364.15, corresponding to approximately 5.4% lower context usage. On QA-8k, the score further increases from 26.85 to 26.90, although mTok increases modestly from 352.00 to 367.24. Thus, calibration primarily improves the quality-efficiency trade-off on QA-16k, while preserving the already strong QA-8k performance at the larger context budget.

D.3 Case Study: Effect of Context Expansion on Token Probability

To provide a concrete and interpretable analysis of how context window expansion affects the model’s predictive confidence, we present a single-generation case study. The example involves extracting the birth date of a person from a long-context passage, a task that exhibits a pronounced uncertainty spike at the day token.

Under the original compact setting (top-k = 1,

Setting	Top-1	Rank of “29”	Log p (“29”) $\uparrow$
Before Expansion	“9”	3rd	−5.13
After Expansion	“29”	1st	−4.56

Table 15: Token ranking and log-probability before and after context expansion. The table reports the candidate ranking for the day token in a birth-date extraction example. Before expansion, the correct token “29” ranks third; after expansion, it becomes top-ranked with log-probability improving from −5.13 to −4.56.

block_size = 16), the correct token “29” receives a log-probability of −5.13 and is ranked third among the candidate tokens {“9”, “15”, “29”}. After context window expansion (top-k = 64), the token “29” becomes the top-ranked candidate, with its log-probability improving to −4.56. Table 15 summarizes the token rankings and log-probabilities before and after expansion.

This case study illustrates that context expansion can effectively mitigate uncertainty by shifting probability mass toward the correct token. The improvement in both rank and log-probability of “29” after expansion supports the hypothesis that insufficient context leads to degraded token-level confidence, and that restoring broader context allows the model to recover a more discriminative distribution over candidates.

D.4 Case Study: False-Positive and False-Negative Failure Cases

To better understand how uncertainty-detection errors affect downstream context allocation, we analyze generation trajectories from the configurations reported in Table 3. We define a false positive (FP) as an unnecessary context-expansion trigger when the compact context is already sufficient, and a false negative (FN) as a missed trigger when additional evidence is required. We select two representative cases of each error type and perform a single-decision counterfactual replay: for an FP, we suppress one unnecessary expansion; for an FN, we force expansion at one missed trigger. All other inference settings remain unchanged. Because generation is autoregressive, modifying a single detector decision can alter subsequent decoding and therefore the total number of downstream rollback operations.

As shown in Table 16, the two error types have markedly different downstream effects. Correcting false positives preserves generation quality at 100% while reducing context usage by 2.8–3.0%,

Case	Intervention	Quality (%) $\uparrow$	mTokens $\downarrow$	Rollbacks $\downarrow$
FP-1	Suppress expansion	100.0 $\rightarrow$ 100.0	379.36 $\rightarrow$ 367.83 (−3.0%)	12 $\rightarrow$ 10
FP-2	Suppress expansion	100.0 $\rightarrow$ 100.0	370.14 $\rightarrow$ 359.75 (−2.8%)	9 $\rightarrow$ 9
FN-1	Force expansion	16.7 $\rightarrow$ 66.7	173.47 $\rightarrow$ 238.00 (+37.2%)	1 $\rightarrow$ 6
FN-2	Force expansion	8.3 $\rightarrow$ 75.0	316.80 $\rightarrow$ 385.47 (+21.7%)	1 $\rightarrow$ 9

Table 16: Representative false-positive and false-negative failures under single-decision counterfactual replay. We suppress one unnecessary expansion for false positives and force one missed expansion for false negatives. Quality measures per-instance generation quality, mTokens the average attended context tokens, and Rollbacks the total rollback operations.

indicating that these errors primarily introduce unnecessary computation. In contrast, correcting false negatives substantially improves per-instance quality, from 16.7% to 66.7% in FN-1 and from 8.3% to 75.0% in FN-2, at the cost of 37.2% and 21.7% more attended context, respectively. The additional context also induces more rollback operations as the corrected trajectory evolves. These cases therefore illustrate an asymmetric failure pattern: false positives mainly reduce context efficiency, whereas false negatives can withhold evidence required for correct generation and consequently have a much larger impact on answer quality. These examples are intended as representative single-trajectory interventions rather than aggregate estimates; the overall quality–efficiency behavior of UT-ACA is reported separately in the main experiments.

E Artifact and Data Usage

Pre-trained Models. We use Llama-3.1-8B-Instruct (Llama Team, 2024) and Qwen2-7B-Instruct (Yang et al., 2024) solely for academic research, consistent with respective licenses (Meta Llama 3.1 Community License and Apache 2.0).

Baseline Implementations. All baseline methods are used strictly for non-commercial academic comparison, consistent with their intended use as research artifacts.

Synthetic Dataset. Our synthetic biography summarization dataset is constructed entirely in-house and intended for research purposes only. The long-context test set incorporates Wikipedia text retrieved for academic evaluation; Wikipedia content is licensed under CC BY-SA 4.0, and our use is strictly within a non-commercial research context.

Automated Evaluation. An LLM-as-judge protocol is used for evaluation, consistent with the model’s intended research usage terms.