

# Enhancing Hallucination Detection via Future Context

Joosung Lee<sup>1</sup>   Cheonbok Park<sup>1,3</sup>   Hwiyeol Jo<sup>1</sup>   Jeonghoon Kim<sup>1,3</sup>  
Joosuk Park<sup>1,2,4†</sup>   Kang Min Yoo<sup>1†</sup>  
<sup>1</sup>NAVER CLOUD   <sup>2</sup>NAVER AI Lab   <sup>3</sup>KAIST   <sup>4</sup>University of Richmond  
{rung.joo,kangmin.yoo}@navercorp.com, park@joosuk.org

## Abstract

Large Language Models (LLMs) are widely used to generate plausible text on online platforms, without revealing the generation process. As users increasingly encounter such black-box outputs, detecting hallucinations has become a critical challenge. To address this challenge, we focus on developing a hallucination detection framework for black-box generators. Motivated by the observation that hallucinations, once introduced, tend to persist, we sample future contexts. The sampled future contexts provide valuable clues for hallucination detection and can be effectively integrated with various sampling-based methods. We extensively demonstrate performance improvements across multiple methods using our proposed sampling approach.

## 1 Introduction

Large Language Models (LLMs) are widely used for various text generation tasks due to their fluency and plausibility. However, these models frequently exhibit overconfidence or generate incorrect information—a phenomenon known as hallucination—which remains a critical challenge in leveraging LLMs effectively across applications (Ji et al., 2023; Sahoo et al., 2024).

LLM-based hallucination detection methods largely fall into two main categories: those based on uncertainty (Varshney et al., 2023; Zhang et al., 2023b; Fadeeva et al., 2024) and those based on sampling strategies (Manakul et al., 2023; Zhang et al., 2023a; Mündler et al., 2024; Yehuda et al., 2024). However, existing approaches suffer from practical issues such as the unavailability of the generator model’s logits in many real-world scenarios and increased sampling costs. For example, in online platforms such as blog posts, it is often difficult to determine how the content was generated.

<sup>†</sup>Corresponding authors.

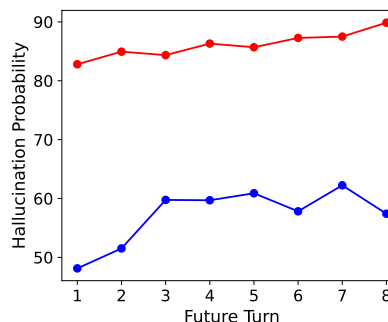


Figure 1: Temporal influence of hallucinated sentences in SelfCheckGPT dataset (based on gold labels). The plot shows the probability that the  $i$ -th future sentence is hallucinated, given that the current sentence is non-hallucinated (blue) or hallucinated (red).

Moreover, API-based generators may become inaccessible over time due to updates or deprecation. Consequently, uncertainty-based methods that rely on generator-specific token-level logits are impractical.

Retrieval-based approaches are effective for verifying general factual information; however, they face significant limitations when internal documents or proprietary knowledge bases are required, as access to these resources is frequently restricted or unavailable. Moreover, retrieval-based methods often struggle to detect logical hallucinations and internal inconsistencies, limiting their effectiveness. Indeed, Mündler et al. (2024) report that 35.2% of self-contradictory hallucinations generated by ChatGPT remain undetectable through retrieval-based verification, suggesting that retrieval-based approaches are not always effective. Consequently, we focus on sampling-based hallucination detection methods that operate in a fully black-box setting without retrieval. In contrast to approaches that rely on internal model signals such as attention (Chuang et al. (2024)), our method leverages future context, making it applicable to diverse and realistic scenarios.

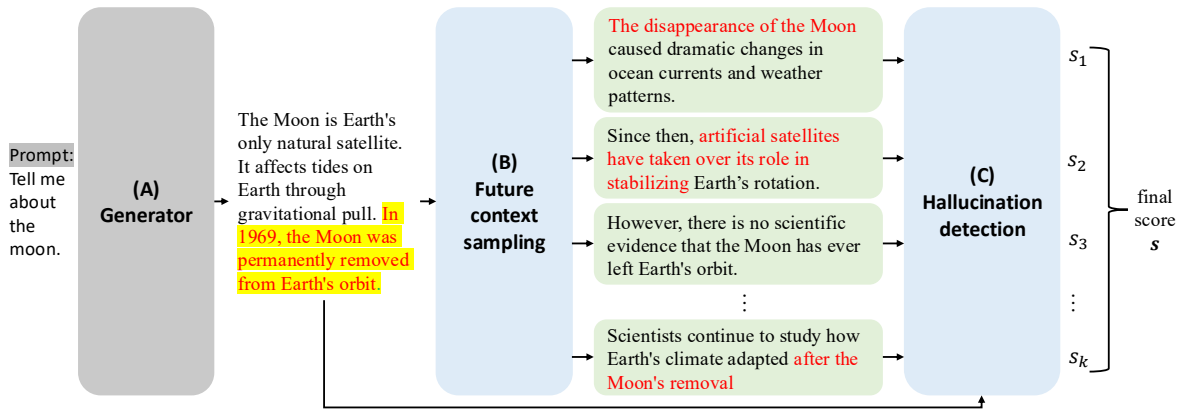


Figure 2: Our proposed hallucination detection pipeline. (A) A black-box generator produces context–response pairs without exposing its internal generation process. (B) Future context sampling: an instruction-tuned LLM generates possible next sentences. (C) Hallucination detection: each sentence is scored using the context, response, and the sampled future context. When the target sentence (yellow highlight) is hallucinated (red text), sampled future contexts are more likely to contain hallucinated information. We propose to leverage this to detect hallucination in the target sentence more reliably.

To cope with these issues, we propose to generate future context and leverage it for hallucination detection in the target sentence. Our method is motivated by the observed snowball effect (Zhang et al., 2024), where hallucinated sentences are often correlated with hallucinations in future context. For instance, in answering yes/no questions, if the first token of a response is *yes*, subsequent sentences tend to support this claim, regardless of correctness. Figure 1, using the SelfCheckGPT dataset, shows that a hallucinated sentence increases the likelihood of hallucinations in subsequent sentences compared to a factual sentence. Furthermore, hallucination errors increasingly propagate to later sentences. Motivated by these observations, we leverage sampled future contexts as clues to detect hallucinations in the target sentence. Section 5.2 confirms that hallucinated current sentences increase the likelihood of hallucinations in future sentences, making the future context a useful signal for detecting hallucinations in the current sentence.

We combine future context with two previous sampling-based methods, SELF-CHECK-GPT and SC (Mündler et al., 2024), as well as our proposed baseline, DIRECT. DIRECT is a simple baseline method that explicitly determines sentence hallucination by directly leveraging the detector LLM’s internal knowledge and reasoning abilities. Through extensive experiments using three detector models (LLaMA 3.1, Gemma 3, Qwen 2.5), we demonstrate that incorporating sampled future context consistently improves hallucination detection performance across various configurations, including

the number of sampled future sentences and lookahead turns. Additionally, we demonstrate that combining future context with the existing sampling clues used in SELF-CHECK-GPT effectively reduces sampling costs.

Our main contributions are as follows:

- We propose leveraging future context to enhance sampling-based hallucination detection, achieving state-of-the-art performance when combined with existing methods.
- Our method is generator-agnostic and can be easily combined with existing frameworks, providing high extensibility across various environments.
- Through extensive analyses, we demonstrate that enriching future context by increasing samples substantially improves hallucination detection performance.

## 2 Related Work

One line of research detects hallucinations by leveraging sampling-based strategies. SELF-CHECK-GPT (Manakul et al., 2023) detects hallucinations by generating multiple context-response pairs and evaluating whether they support the original output. If the sampled pairs provide consistent evidence for the original response, it is deemed factual; otherwise, it is flagged as a hallucination. SAC<sup>3</sup> (Zhang et al., 2023a) enhances sampling-based detection by paraphrasing the input and leveraging multiple models to produce diverse outputs. This multi-perspective generation improves the ability to un-

cover hallucinations that may escape traditional consistency checks. **SC** (Self-Contradiction Detection) (Mündler et al., 2024) addresses hallucinations by identifying internal contradictions within generated responses. It regenerates each sentence based on its extracted context and compares it with the original to detect inconsistencies.

Another line of research focuses on token-level uncertainty to assess factuality. **STITCH** (Varshney et al., 2023) identifies hallucinations by analyzing low-confidence tokens. It shows that the minimum token probability provides an effective signal for detecting factual errors. **CCP** (Claim Conditioned Probability) (Fadeeva et al., 2024) measures uncertainty at both token and claim levels using model-generated alternative tokens. These are assessed via natural language inference to estimate the overall factuality of the claim.

We propose a novel sampling-based approach that leverages future context to enhance hallucination detection. By incorporating future context as a clue in existing methods, our approach improves performance in black-box generator settings.

### 3 Approach

Figure 2 is an overview of our method. Since we assume no access to the generator, the detector LLM need not match the generator.

#### 3.1 Future Context Sampling

We use an instruction-tuned LLM as a future context sampler, prompting the model to generate the next sentence. When generating future context that extends beyond the immediate next sentence, we found it more effective to instruct the model to generate multiple sentences simultaneously, rather than generating them sequentially. A future context is defined as the set of sentences generated through a single sampling path.

#### 3.2 Incorporating Future Context into Hallucination Detection

We explore the integration of future context into various sampling-based methods for hallucination detection. While the effective strategies for utilizing future context may differ across methods, we adopt a simple and unified approach: appending future context directly to the prompt.

**SELF-CHECKGPT with Future Context.** First, we incorporate future context into **SELF-CHECKGPT**, a method that detects hallucinations by evalu-

ating the consistency between the original response and multiple sampled alternative context-response pairs. The core idea of **SELF-CHECKGPT** is that factual information tends to be consistently present across different sampled responses, whereas hallucinated information tends to vary or contradict. By appending future context to the alternative context-response pairs, we expand the clues available to the LLM for detecting hallucinations.

**SC with Future Context.** Second, we incorporate future context into **SC**, which samples alternative current-turn responses to identify inconsistencies. We omit the description field in the original **SC** prompt for consistency across datasets lacking this field. In settings where future context is used, we instead substitute the description with a future context.

**DIRECT with Future Context.** Finally, we apply future context to our proposed baseline method **DIRECT**. Previous works (Kadavath et al., 2022; Yin et al., 2023) demonstrate that LLMs exhibit a form of self-awareness, suggesting they can leverage their internal knowledge for hallucination detection. Motivated by this, we introduce **DIRECT**, explicitly prompting the detector to determine hallucinations. **DIRECT** determines hallucination without relying on complex probability estimations or sampling techniques. This approach relies solely on the internal knowledge and reasoning capabilities of the detector LLM, serving as a straightforward and effective starting point for leveraging the model’s contextual understanding and factual grounding. This approach allows for precise control over the key elements involved in hallucination detection, making it well-suited for comprehensively analyzing the effects of future context.

**DIRECT** poses a straightforward binary (*yes/no*) question to the model, asking whether a given sentence is accurate. The prompt consolidates information, aiding the model’s internal knowledge-based hallucination detection. In this sense, this approach serves as a strong baseline and a practical starting point for hallucination detection using LLMs. All prompt templates are provided in Appendix B.

#### 3.3 Hallucination Score

In the sampling-based methodology, following **SELF-CHECKGPT**, we pair the target sentence with multiple sampled clues. In **DIRECT**, each clue corresponds to a sampled future context. For **SELF-CHECKGPT** and **SC**, each clue is formed by com-

binning their originally sampled outputs with the sampled future context. Each sentence-clue pair is individually evaluated for hallucination: if the response is judged to be a hallucination given the clue, it is assigned a score of 0; if not, it receives a score of 1; and if the judgment is indeterminate (i.e., N/A), a score of 0.5 is assigned. Finally, the scores across all pairs are averaged to obtain the hallucination score for the sentence.

## 4 Datasets

Previous studies typically evaluate hallucination detection methods using data labeled by their respective generators, making it difficult to compare performance directly due to differences in generators, detectors, and dataset formats. Furthermore, existing datasets differ in the types of hallucinations they address, with SelfCheckGPT and SC datasets targeting logical hallucinations and True-False focusing on factual hallucinations. To address this limitation, we standardize the dataset format and conduct unified evaluations across various methods. As our approach assumes a black-box generator setting, we do not require detectors tailored to specific generators, enabling diverse LLM-based detectors. Examples are provided in Table 3.

SelfCheckGPT dataset (Manakul et al., 2023) is constructed using GPT-3 (text-davinci-003) (Brown et al., 2020), which is used to sample responses to given questions. Two annotators label each response as either accurate, minor inaccurate, or major inaccurate. If annotators disagree, the worse label is chosen. We treat responses labeled as minor or major inaccurate as hallucinations. The context used for response generation consists of model-generated sentences.

SC dataset (Mündler et al., 2024) is constructed using various LLMs, such as ChatGPT (gpt-3.5-turbo-0301) (OpenAI, 2023), GPT4 (gpt-4-0314) (OpenAI et al., 2024), LLaMA2-70B-Chat (Touvron et al., 2023), and Vicuna-13B (Chiang et al., 2023), which generate responses based on the given question and context. An initial response is first generated, followed by an additional sampled response. Three annotators then assess whether a self-contradiction exists between the two responses. If a self-contradiction is identified, at least one of the responses must contain hallucinated content (Dowden, 2011; Mündler et al., 2024). This does not necessarily imply that the initial response is incorrect. Nevertheless, Mündler et al. (2024)

has shown that removing all conflicting information from both sentences can improve overall response quality. Therefore, despite certain limitations, we consider the initial response to be hallucinated when a self-contradiction is detected, in order to enable broader benchmarking. We define a concept based on the string following the term ‘about’ in a given question. Unlike SelfCheckGPT, where context is constructed from model-generated text, the SC dataset uses gold responses as context. For convenience, we refer to this version of the SC dataset as ‘SC- $\{LLM\}$ ’.

True-False dataset (Azaria and Mitchell, 2023) comprises individual sentences labeled as either true or false. Sentences labeled as false are considered hallucinations. As this dataset does not provide prompts, we constructed prompts following the same format used in SC dataset. The concept is extracted from each sentence using an LLM, and only sentences that contain the exact concept string are used for evaluation. Moreover, to ensure fairness, the LLM used for concept extraction is different from the one used as the hallucination detector. Table 4 presents dataset statistics\*.

## 5 Experiments

We employ three instruction-tuned LLM detectors: LLaMA 3.1-8B (Grattafiori et al., 2024), Gemma 3-12B (Kamath et al., 2025), and Qwen 2.5-7B (Yang et al., 2025). By sampling future context from these detectors and integrating it with various methods, we demonstrate improved hallucination detection performance. We evaluate using AUROC as our primary metric because our dataset is not significantly imbalanced and AUROC is commonly adopted in related research. For completeness, we also report AUCPR results in Table 11.

We use top-k sampling ( $k = 30$ ) to generate future and current sentences, and follow the original generation configurations for alternative context–response pairs in SELFCKGPT. Generation terminates when an end-of-sequence token is produced, and the resulting text is segmented into sentences using SpaCy<sup>†</sup>. Finally, we use greedy decoding when the detector generates responses to prompts for hallucination detection.

\*The reconstructed datasets will be released upon acceptance.

<sup>†</sup><https://spacy.io>

Detector	Method	Dataset					True-False	Average
		SelfCheckGPT	SC-ChatGPT	SC-GPT4	SC-LLaMA	SC-Vicuna		
LLaMA 3.1	DIRECT	62.1	66.4	63.8	65.9	68.0	87.2	68.9
	DIRECT+f	<b>65.3</b>	<b>69.0</b>	<b>68.0</b>	<b>67.1</b>	<b>68.4</b>	<b>88.5</b>	<b>71.1</b>
	SELFCKEKGPT	67.0	71.2	76.0	68.1	69.6	89.2	73.5
	SELFCKEKGPT+f	<b>67.3</b>	<b>74.9</b>	<b>76.4</b>	<b>69.8</b>	<b>70.2</b>	<b>90.3</b>	<b>74.8</b>
	SC	53.2	62.5	65.8	67.1	67.3	78.1	65.7
SC+f	<b>57.7</b>	<b>69.2</b>	<b>73.3</b>	<b>71.2</b>	<b>71.4</b>	<b>81.8</b>	<b>70.8</b>	
Gemma 3	DIRECT	53.7	52.5	53.4	55.6	56.6	87.9	59.9
	DIRECT+f	<b>56.6</b>	<b>58.0</b>	<b>58.0</b>	<b>58.1</b>	<b>57.5</b>	<b>90.8</b>	<b>63.2</b>
	SELFCKEKGPT	49.3	68.8	71.9	68.7	67.7	89.7	69.4
	SELFCKEKGPT+f	<b>58.3</b>	<b>71.0</b>	<b>72.5</b>	<b>71.2</b>	<b>70.1</b>	<b>91.3</b>	<b>72.4</b>
	SC	52.9	61.1	62.3	64.2	65.7	86.1	65.4
SC+f	<b>55.2</b>	<b>65.6</b>	<b>65.3</b>	<b>66.2</b>	<b>67.7</b>	<b>86.1</b>	<b>67.7</b>	
Qwen 2.5	DIRECT	60.6	63.0	67.9	64.8	65.5	82.6	67.4
	DIRECT+f	<b>62.9</b>	<b>65.8</b>	<b>69.1</b>	<b>67.0</b>	<b>65.6</b>	<b>85.7</b>	<b>69.4</b>
	SELFCKEKGPT	51.7	64.4	67.8	63.5	65.7	87.4	66.7
	SELFCKEKGPT+f	<b>51.7</b>	<b>65.5</b>	<b>70.3</b>	<b>65.4</b>	65.1	85.9	<b>67.3</b>
	SC	50.5	50.7	50.6	52.3	56.0	69.3	54.9
SC+f	<b>54.2</b>	<b>59.5</b>	<b>60.5</b>	<b>60.0</b>	<b>60.2</b>	<b>70.2</b>	<b>60.8</b>	

Table 1: Hallucination detection performance (AUROC) of DIRECT, SELFCKEKGPT, and SC with and without future context (+) across different LLM detectors. Bold numbers indicate performance improvements when future context (+) are used.

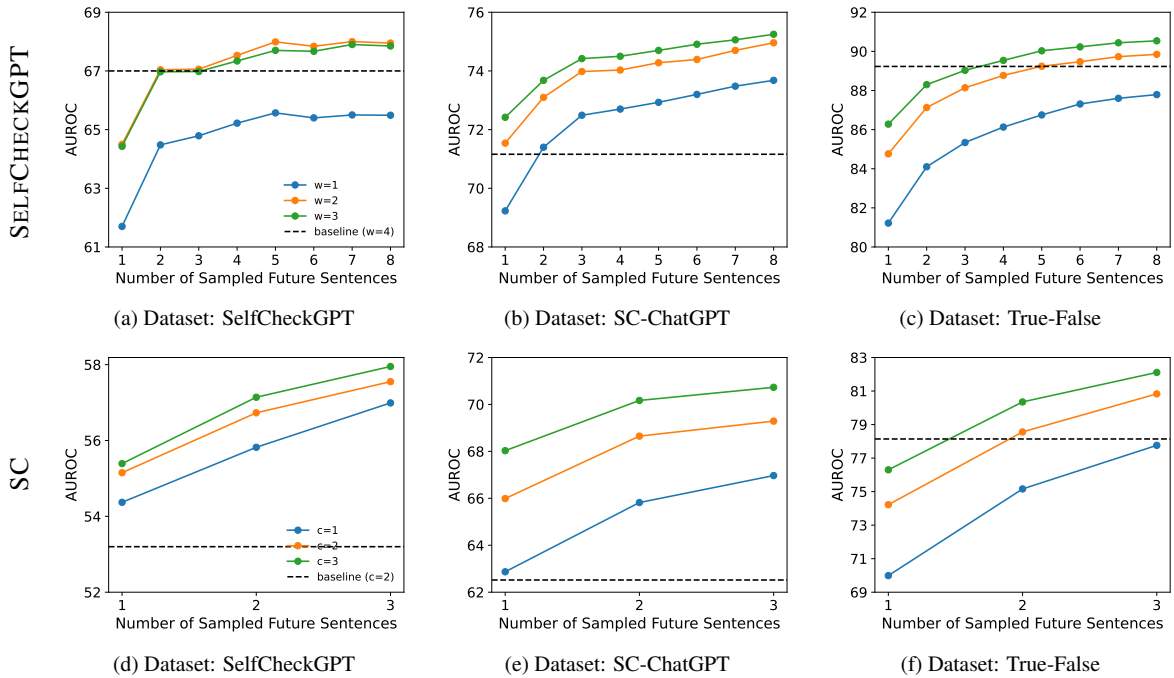


Figure 3: AUROC of SELFCKEKGPT and SC with future contexts (detector: LLaMA 3.1). The first and second rows show performance improvements for SELFCKEKGPT and SC, respectively, when incorporating future contexts. In both cases, performance further increases as the number of sampled future contexts grows.

## 5.1 Future Context Improves Performance

We investigate performance improvements by incorporating future context into sampling-based methods. Table 1 shows the performance of DIRECT, SELFCKEKGPT, and SC—each evaluated with and without future context—across different LLM detectors. In SELFCKEKGPT, we use

$w = 4$  alternative context-response pairs, where  $w$  denotes the number of context variations paired with the same response. In SC, we use  $c = 2$ , representing the number of sampled responses for the current sentence. In each of the future-augmented variants—DIRECT+f, SELFCKEKGPT+f, and SC+f—we sample four one-turn-ahead future sentences and combine them with the original prompts.




Figure 4: AUROC vs. sampled token consumption using the SELFCKEKGPT+f on the SC-ChatGPT (Detector: LLaMA 3.1).




Figure 5: Hallucination rates of future sentences (as shown in the top row of Figure 7, LLaMA 3.1), conditioned on the hallucination label of the current sentence. **All Future (H)**: The proportion of hallucinated future sentences when the current sentence is hallucinated. **All Future (NH)**: The proportion of hallucinated future sentences when the current sentence is non-hallucinated. **Positive Future (H)**: Proportion of hallucinated future sentences aiding correct classification of hallucinated current sentences. **Positive Future (NH)**: Proportion of hallucinated future sentences aiding correct classification of non-hallucinated current sentences.

Incorporating future context improves performance in most cases.

For LLaMA 3.1 and Gemma 3, SELFCKEKGPT achieves the best performance, whereas for Qwen 2.5, DIRECT+f performs best. Sampling-based methods generally benefit from increased samples, leading to improved performance (Manakul et al., 2023). Table 6 presents the average token counts of sampled outputs for each LLM. Since SELFCKEKGPT samples entire alternative context-response pairs, it requires more tokens compared to methods sampling only future or current sentences. Thus, even with equal samples, SELFCKEKGPT incurs higher computational costs. For comprehensive analysis, we extend sampling targets  $(w, c)$  with future sentences across various parameter combinations.

Figure 3 (top row) shows the performance of SELFCKEKGPT as the number of sampled  $w$  and




Figure 6: Non-Hallucination rates of future sentences (as shown in the top row of Figure 7, LLaMA 3.1) conditioned on the hallucination label of the current sentence. **All Future (H)**: The proportion of non-hallucinated future sentences when the current sentence is hallucinated. **All Future (NH)**: The proportion of non-hallucinated future sentences when the current sentence is non-hallucinated. **Positive Future (H)**: Proportion of non-hallucinated future sentences aiding correct classification of hallucinated current sentences. **Positive Future (NH)**: Proportion of non-hallucinated future sentences aiding correct classification of non-hallucinated current sentences.

$s$  varies, with  $t = 1$ , where  $s$  is the number of future sentence samples per turn and  $t$  is the number of future turns. Figure 4 further analyzes the relationship between AUROC and the number of sampled tokens across different  $(w, s)$ . The results indicate that reducing  $w$  while increasing  $s$  can be more token-efficient than solely increasing  $w$ , demonstrating cases where comparable or superior performance is achieved with fewer tokens. Since the number of sampled tokens significantly impacts inference speed during both response generation and hallucination detection, these findings support the cost-effectiveness of leveraging future contexts. Additional analyses regarding the number of inference steps are provided in Appendices E and F.

Figure 3 (bottom row) illustrates the performance of SC under various combinations of sampled  $c$  and  $s$ . As with SELFCKEKGPT, increasing  $s$  alongside  $c$  generally improves performance. While incorporating  $s$  does not reduce token costs (as  $c$  and  $s$  are similar in token length), combining  $s$  with a fixed total sample count often improves performance. This observation suggests the potential benefit of leveraging future context through the inclusion of  $s$ . Additional experimental results are provided in Figures 12.

Additionally, Appendix G explores the potential effectiveness of incorporating future context into retrieval-based methods. Appendix H further demonstrates its effectiveness in detecting passage-level hallucinations.

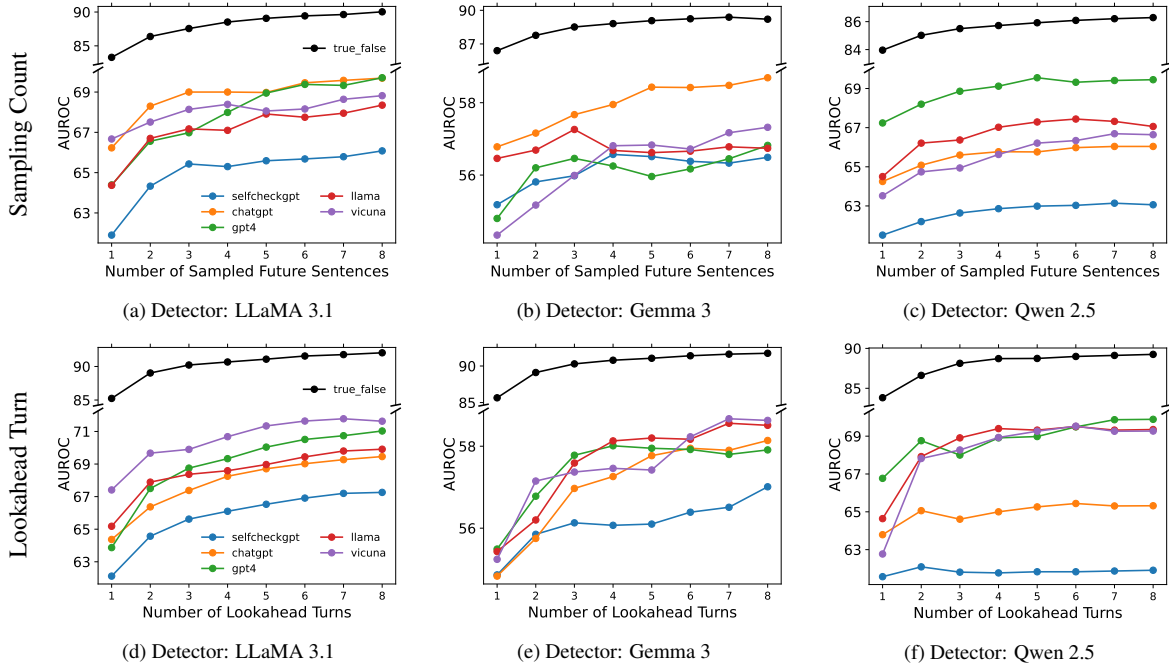


Figure 7: AUROC of DIRECT with future contexts. The first row shows performance improvements as the number of sampled future sentences from one turn ahead increases. The second row shows performance improvements as the number of future lookahead turns increases. In both cases, incorporating more future context consistently enhances hallucination detection across all detectors.

## 5.2 Analyzing the Impact of Future Sentences

We investigate which future sentence aids in detecting hallucinations in current (target) sentences. Sampled future sentences are evaluated for hallucinations using DIRECT.

Figure 5 shows analysis results for cases where the future sentence is hallucinated, with an illustrative example provided in Table 12. Our findings indicate that: **(1) All Future (H) > All Future (NH)**: hallucinated current sentences are more likely to be followed by hallucinated future sentences. **(2) Positive Future (H) > All Future (H)**: hallucinated future sentences are especially useful for detecting hallucinated current sentences. **(3) Positive Future (NH) < All Future (NH)**: hallucinated future sentences are less useful for detecting non-hallucinated current sentences.

Figure 6 shows analysis results for cases where the future sentence is non-hallucinated. Our findings indicate that: **(1) All Future (NH) > All Future (H)**: non-hallucinated current sentences are more likely to be followed by non-hallucinated future sentences. **(2) Positive Future (NH) > All Future (NH)**: non-hallucinated future sentences are especially useful for detecting non-hallucinated current sentences. **(3) Positive Future (H) < All Future (H)**: non-hallucinated future sentences are

less useful for detecting hallucinated current sentences.

We observed this consistent pattern across all datasets. These findings show that future context is statistically associated with the hallucination status of the current sentence. Although this correlation does not establish causation, it consistently provides useful predictive signals for hallucination detection. Consequently, future sentences are more helpful for hallucination detection when their status matches the current sentence.

## 6 Further Findings on Future Context

### 6.1 More Future Samples Improve Performance

Future context can be expanded by increasing either the lookahead turns ( $t$ ) or samples per turn ( $s$ ). Figure 7 (top row) shows the impact of the number of sampled future sentences on performance ( $t = 1$ ) in DIRECT+f. As the number of sampled future sentences increases, the available context expands, enabling multiple evaluations of hallucination presence. Moreover, sampling future sentences along multiple paths allows the influence of the target sentence to be captured in diverse ways. Sampling more future sentences increases helpful information, improving hallucination detection.

Detector	Method	SelfCheckGPT	SC-ChatGPT	SC-GPT4	SC-LLaMA	SC-Vicuna	True-False	Average
LLaMA 3.1	DIRECT+ $f_p$	<b>67.5</b>	68.7	<b>71.2</b>	66.2	65.1	<b>91.1</b>	<b>71.6</b>
	DIRECT+ $f_n$	<b>65.3</b>	<b>69.2</b>	67.9	<b>67.1</b>	68.0	<b>88.5</b>	71.0
Gemma 3	DIRECT+ $f_p$	<b>59.1</b>	<b>58.1</b>	<b>58.6</b>	58.0	56.3	88.1	63.0
	DIRECT+ $f_n$	56.4	57.8	56.4	56.3	56.6	88.7	62.0
Qwen 2.5	DIRECT+ $f_p$	<b>63.3</b>	<b>66.4</b>	<b>70.3</b>	<b>69.3</b>	<b>69.3</b>	<b>86.4</b>	<b>70.8</b>
	DIRECT+ $f_n$	<b>63.0</b>	65.4	66.7	<b>67.3</b>	63.1	85.2	68.4

Table 2: AUROC of DIRECT variants with filtered future context ( $+f_p$ ,  $+f_n$ ) across different LLM detectors and datasets.  $+f_p$  denotes prompt-based filtering of future context, while  $+f_n$  indicates NLI-based filtering. Bold indicates improvements equal to or greater than DIRECT+f.




Figure 8: AUROC of DIRECT with future contexts on SelfCheckGPT dataset as both the number of sampled future sentences and the future lookahead turns increase (detector: LLaMA 3.1).

Additionally, Figures 11 present results for integrating sampled future sentences into the previous methods, SELFCKEKGPT ( $w = 1$ ) and SC ( $c = 1$ ). In all cases, increasing the number of samples leads to performance gains. Our method performs consistently across different detection frameworks, indicating that it is not strongly dependent on a specific approach.

## 6.2 Longer Lookahead Improves Performance

Figure 7 (bottom row) presents results based on varying the lookahead turn, i.e., how far into the future the model explores. To observe the effect of lookahead turn, the number of sampled sentences is fixed to one. We observe that increasing lookahead turn generally leads to performance improvements, regardless of the detector. As shown in Figure 1, when the current sentence is hallucinated, the likelihood of hallucination increases in more distant future sentences. This suggests that the current sentence has a stronger snowballing effect on later turns, making it effective to utilize future sentences from farther ahead.

## 6.3 Joint Impact of Sampling Count and Lookahead Turn

Figure 8 shows the effect of varying lookahead turns and samples per turn; all results are provided

in Figure 13. We hypothesize that increasing  $t$  better captures the "snowball effect"—the way the current sentence influences later ones—while increasing  $s$  uncovers multiple alternative continuations at each future turn. If the total number of future sentences is fixed at  $t \times s$ , increasing  $t$  more clearly reveals how the hallucination status of the current sentence is propagated across future turns. On the other hand, increasing  $s$  provides diverse rephrasings for a single future turn, which makes it easier to verify the accuracy of the information. In practice, using both higher  $t$  and  $s$  offers a richer set of predictive clues and contributes to more accurate hallucination detection.

However, increasing both  $t$  and  $s$  can be computationally expensive. The trade-off between accuracy and efficiency should depend on the application: tasks needing thorough verification (e.g., legal or medical) may benefit from larger  $t$  and  $s$ , while real-time or resource-limited settings (e.g., chat interfaces) may keep one parameter moderate or adopt adaptive strategies. A detailed comparison between varying  $t$  and  $s$  is provided in Appendix C.

## 6.4 Filtering of Future Context

We explore two methods for selectively integrating semantically relevant future contexts into hallucination detection. The first approach explicitly prompts LLMs to use future contexts only when they are helpful, effectively inducing implicit filtering through detailed instructions (Chen et al., 2024; Wang et al., 2024). The second approach employs an NLI model (He et al., 2021) to measure semantic relevance between the current sentence and candidate future contexts, retaining those whose combined entailment and neutral scores exceed a threshold. While the optimal threshold may vary across models and datasets, we adopt a fixed value for simplicity, which in turn induces explicit filtering.

Table 2 demonstrates that prompt-based filter-

ing provides a simple yet effective improvement in hallucination detection performance. Although we employed a fixed threshold for NLI-based filtering, we believe there is room for performance improvement depending on how semantic relevance is utilized. These findings indicate that selectively incorporating meaningful future contexts is crucial for improving hallucination detection, highlighting the potential for further research on advanced context-selection methods.

## 7 Conclusion

We propose a method to detect hallucinations by sampling generated future contexts without relying on generator information. We provides a model-agnostic benchmark and achieves improved detection performance when combined with existing methods (SELF-CHECKGPT and SC) and the baseline DIRECT. Particularly, integrating future contexts with SELF-CHECKGPT effectively reduces sampling costs. Evaluations across various turns and samples highlight the effectiveness of future contexts in sampling-based hallucination detection.

## Limitations

### Quality of Future Sentence

Section 6.4 discusses selective leveraging of future contexts; however, when only low-quality samples are available, its effectiveness becomes marginal. Table 7 and Table 8 analyze such cases, typically caused by insufficient generation or duplication. In particular, Qwen2.5 often produces generic duplicates (e.g., "Sure!"), and its weaker instruction-following reduces usable context, limiting the benefit of longer lookahead turns (Figure 9, Figure 7f). While these issues stem from model biases and could be alleviated by prompt engineering, we kept settings consistent across detectors. Future contexts produced by generators, as opposed to those from detectors, tend to form more coherent paths and seldom include problematic sentences (Appendix D).

### Correlation as a Sufficient Signal for Detection

While our analysis in Section 5.2 highlights that future sentences are statistically correlated with the hallucination status of the current sentence, we emphasize that causality is not a prerequisite for our approach. Our framework only requires future contexts to provide predictive signals, and thus they remain effective for hallucination detection even if

they are correlational byproducts rather than causal effects.

## References

- Amos Azaria and Tom Mitchell. 2023. [The internal state of an LLM knows when it's lying](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2024. [Universal self-consistency for large language models](#). In *ICML 2024 Workshop on In-Context Learning*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).
- Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James R. Glass. 2024. [Lookback lens: Detecting and mitigating contextual hallucinations in large language models using only attention maps](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1419–1436, Miami, Florida, USA. Association for Computational Linguistics.
- Bradley H. Dowden. 2011. *Logical Reasoning*.
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, and Maxim Panov. 2024. [Fact-checking the output of large language models via token-level uncertainty quantification](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9367–9385, Bangkok, Thailand. Association for Computational Linguistics.
- Aaron Grattafiori et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*, 55(12):1–38.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. 2022. [Language models \(mostly\) know what they know](#). *Preprint*, arXiv:2207.05221.
- Aishwarya Kamath et al. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2024. [Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation](#). *Preprint*, arXiv:2305.15852.
- OpenAI. 2023. [Introducing chatgpt](#). Accessed: 2025-04-08.
- OpenAI, Josh Achiam, et al. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Pranab Sahoo, Prabhash Meharia, Akash Ghosh, Sriparna Saha, Vinija Jain, and Aman Chadha. 2024. [A comprehensive survey of hallucination in large language, image, video and audio foundation models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11709–11724, Miami, Florida, USA. Association for Computational Linguistics.
- Yixiao Song, Yekyung Kim, and Mohit Iyyer. 2024. [VeriScore: Evaluating the factuality of verifiable claims in long-form text generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9447–9474, Miami, Florida, USA. Association for Computational Linguistics.
- Hugo Touvron et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jian-shu Chen, and Dong Yu. 2023. [A stitch in time saves nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation](#). *Preprint*, arXiv:2307.03987.
- Xinglin Wang, Yiwei Li, Shaoxiong Feng, Peiwen Yuan, Boyuan Pan, Heda Wang, Yao Hu, and Kan Li. 2024. [Integrate the essence and eliminate the dross: Fine-grained self-consistency for free-form language generation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11782–11794, Bangkok, Thailand. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#).
- An Yang et al. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Yakir Yehuda, Itzik Malkiel, Oren Barkan, Jonathan Weill, Royi Ronen, and Noam Koenigstein. 2024. [InterrogateLLM: Zero-resource hallucination detection in LLM-generated answers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9333–9347, Bangkok, Thailand. Association for Computational Linguistics.
- Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. 2023. [Do large language models know what they don’t know?](#) In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8653–8665, Toronto, Canada. Association for Computational Linguistics.
- Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2023a. [SAC<sup>3</sup>: Reliable hallucination detection in black-box language models via semantic-aware cross-check consistency](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15445–15458, Singapore. Association for Computational Linguistics.

Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2024. [How language model hallucinations can snowball](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 59670–59684. PMLR.

Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. 2023b. [Enhancing uncertainty-based hallucination detection with stronger focus](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 915–932, Singapore. Association for Computational Linguistics.

## A Details of Datasets

Table 3 shows an example of the data format used in our experiments. Since not all datasets are originally provided in the form of question, concept, context, and response, we reconstructed them to match this format.

Table 4 presents statistics of our dataset. In Self-CheckGPT dataset, responses are generated continuously for questions sampled from the WikiBio dataset. As a result, once a hallucination occurs, it can affect subsequent parts of the response, leading to a higher number of hallucinated cases. In contrast, SC-LLM uses gold responses as context and generates only a single sentence as the response using the generator. Therefore, it contains more non-hallucinated cases. True-False dataset consists of six topics: Cities, Inventions, Chemical Elements, Animals, Companies, and Scientific Facts. True statements were crafted using reliable factual sources. False statements were mostly generated by pairing mismatched attributes from different factual entries, except for the Scientific Facts category, where ChatGPT was used to generate opposite statements. Human annotators then reviewed and labeled each statement to ensure accuracy. True-False dataset consists of individual sentences labeled as either true or false, without any accompanying context.

## B LLM Prompts

### B.1 Template for Future Context Sampling

The following is the base prompt used for future context sampling. It is used to generate the sentence that directly follows the target sentence for hallucination detection.

```
next sentence
{question}
{past context} {sentence}
next sentence:
```

To sample  $k$  future sentences, we instruct the model with the phrase " $k$  next sentences."

### B.2 Templates for Hallucination Detection

The prompt of DIRECT is as follows:

#### DIRECT

```
Here is a sentence about {concept}:
context: {past context} {future context}
sentence: {sentence}
Is this sentence accurate about the {concept}?
Answer Yes or No:
```

where concept, question, past context, and sentence are provided by the dataset and are described in Section 4. Future context is optional and may be selectively used depending on the experiment. The goal of this prompt is to intuitively provide the model with clues, enabling it to determine whether the target sentence is hallucinatory based on its internal knowledge. This prompt forms the core of our DIRECT baseline.

To implement prompt-based filtering in Table 2, we provide detailed instructions to use only relevant Future contexts, as shown below. To do this, we change the prompt to distinguish between past and Future contexts.

#### DIRECT (PROMPT-BASED FILTERING)

```
Here is a sentence about {concept}:
past context: {past context}
future context: {future context}
The future context may or may not be relevant. Use
it only if it helps you determine the accuracy of the
sentence.
sentence: {sentence}
Is this sentence accurate about the {concept}?
Answer Yes or No:
```

We next describe the prompt design used in SELF-CHECKGPT for hallucination detection. Given a question, the LLM detector first generates multiple alternative context–response pairs using top- $k$  sampling (specifically, we set  $k = 50$ ). These pairs are then used as the context in the prompt, optionally supplemented with future context as follows:

#### SELF-CHECKGPT

```
Context: {context-response pairs} {future
context}
Sentence: {sentence}
Is the sentence supported by the context above? Answer
Yes or No:
```

SC adopts a two-stage procedure to detect hallucinations by identifying logical inconsistencies among alternative completions. Before detecting hallucinations, the LLM-based detector samples alternative sentences corresponding to the current turn. In the first stage—explanation gener-

Dataset	Type	Text
SelfCheckGPT	Question	This is a Wikipedia passage about Matthew Aylmer, 1st Baron Aylmer:
	Concept	Matthew Aylmer, 1st Baron Aylmer
	Context	Matthew Aylmer, 1st Baron Aylmer (1708–1794) was an Irish soldier and colonial administrator.
	Response	He was born in Dublin, the son of a barrister, and was educated at Trinity College, Dublin.
SC-LLM	Question	Please tell me about 1990 Austrian Hockey League season
	Concept	1990 Austrian Hockey League season
	Context	The 1990-91 Austrian Hockey League season was the 61st season of the Austrian Hockey League, the top level of ice hockey in Austria.
	Response	EC Graz won the championship that season.
True-False	Question	Please tell me about anteater
	Concept	anteater
	Context	-
	Response	The giant anteater uses walking for locomotion.

Table 3: Examples of hallucination dataset

Dataset	SelfCheckGPT	SC-ChatGPT	SC-GPT4	SC-LLaMA	SC-Vicuna	True-False
Hallucination O	1392	491	207	236	185	1972
Hallucination X	516	2935	1527	1532	878	2125

Table 4: Statistics of dataset

ation—the model is prompted to describe informational conflicts between two statements: the original target sentence and a sampled alternative. Optionally, future context can replace the original description field (to maintain dataset consistency), providing a natural grounding for comparison.

#### SC (explanation)

Prompt: Here is the start of a description about {concept}.  
Then follow two statements.  
Description: {future context}  
Statement 1: {sentence}  
Statement 2: {sampled sentence}  
Please explain if the statements about concept are contradictory. Provide your explanation.

In the second stage—contradiction detection—the model receives the explanation generated in the first stage and is prompted to make a binary decision on whether the two statements indeed contradict each other.

#### SC (detection)

{explanation}  
Please conclude with Yes if there is a contradiction, otherwise No.  
answer:

This two-stage framework encourages the model to explicitly reason over the relationship between sampled variants, enhancing interpretability and providing fine-grained insight into how hallucinations are detected.

## C Lookahead Turn vs. Sampling Count

We find that both increasing the number of samples and extending the lookahead turn for future context sampling are effective strategies. Increasing the number of samples allows the model to verify the influence of the current sentence across multiple possible paths, whereas increasing the lookahead turn captures the accumulation of that influence over time. Table 5 presents a comparison between these two approaches. We sample with both the number of samples and the number of lookahead turns ranging from 1 to 8 and report the average results. The results indicate a trend where increasing lookahead turns provides more benefit than increasing the number of samples. As discussed in Limitation 7, since the quality of future sentences was not controlled, the results may vary depending on the detector and dataset.

Detector	SelfCheckGPT		SC-ChatGPT		SC-GPT4		SC-LLaMA		SC-Vicuna		True-False	
	s-avg	t-avg	s-avg	t-avg	s-avg	t-avg	s-avg	t-avg	s-avg	t-avg	s-avg	t-avg
LLaMA 3.1	65.0	65.8	68.8	67.9	67.9	69.0	67.2	68.5	68.0	70.5	88.1	90.2
Gemma 3	56.3	56.1	57.9	57.1	56.1	57.5	56.7	57.6	56.3	57.5	88.8	90.2
Qwen 2.5	62.7	61.8	65.6	64.9	68.9	68.8	66.7	68.5	65.6	68.1	85.6	87.9
Average	<b>61.3</b>	61.2	<b>64.1</b>	63.3	64.3	<b>65.1</b>	63.5	<b>64.9</b>	63.3	<b>65.4</b>	87.5	<b>89.4</b>

Table 5: Hallucination detection performance of DIRECT with future contexts. *s-avg*: performance averaged across number of samples 1–8 (future turn fixed at 1); *t-avg*: performance averaged across future turns 1–8 (with one sample per turn).

Tokenizer	SelfCheckGPT			SC-ChatGPT			SC-GPT4			SC-LLaMA			SC-Vicuna			True-False		
	$w_{tok}$	$s_{tok}$	$c_{tok}$	$w_{tok}$	$s_{tok}$	$c_{tok}$	$w_{tok}$	$s_{tok}$	$c_{tok}$	$w_{tok}$	$s_{tok}$	$c_{tok}$	$w_{tok}$	$s_{tok}$	$c_{tok}$	$w_{tok}$	$s_{tok}$	$c_{tok}$
LLaMA 3.1	34.5	35.1	33.6	317.9	43.1	42.6	309.1	43.9	42.7	307.1	42.8	42.6	310	41.6	41.2	497.9	28.6	35.6
Gemma 3	39.2	23.7	22.9	510.7	33.5	34.1	510.6	34.4	34.3	510.8	34.2	34.3	510.4	34.8	34.8	510.3	18.6	20.3
Qwen 2.5	33.5	29.6	29.6	362.2	28.5	28.5	362.6	28.8	28.8	361.8	29.2	29.2	360.1	28.5	28.6	451.7	16.6	16.6

Table 6: Average token counts for sampled outputs, where  $w_{tok}$  corresponds to alternative context-response pairs in SELFCKEKGPT,  $c_{tok}$  corresponds to current sentence samples in SC, and  $s_{tok}$  corresponds to future contexts.




Figure 9: Hallucination detection performance comparison based on the source of sampled future context: Generator vs. Detector.

## D Future Context Source: Generator vs. Detector

Figure 9 presents experimental results based on the source of the future context. While we generally do not have access to the generator, there are scenarios where the entire generated response is available — for instance, in blog posts or online articles that include the full text content. In such cases, it is possible to utilize future context extracted from the generator’s own output. SelfCheckGPT dataset provides full responses generated by the model, allowing us to use the sentences following the hallucination detection target as future context. Since the generator provides only a single path, we also limit detector samples to one ( $s = 1$ ). For consistency, we evaluate only cases with at least five future sentences, allowing comparison across different  $t$  values.

Future context generated by the generator (GF) leads to improved performance as the number of

future turns increases. Future context generated by the detector (DF) also contributes to performance gains for LLaMA 3.1 and Gemma 3, but shows no clear benefit for Qwen 2.5. After analyzing the DF for Qwen 2.5, we found that the sampled future sentences in the SelfCheckGPT dataset were of relatively low quality compared to other detectors. Further discussion on this limitation is provided in Section 7. Overall, while leveraging the generator’s output is preferable, sampling future context using the detector can be an effective alternative in a black-box generator setting.

## E Aggregation of Future Sentences

When utilizing sampled future sentences, if each sentence is used as an individual future context, then  $s$  sampled future sentences require  $s$  separate inferences. To reduce the number of inferences, we concatenate the sampled one-turn-ahead future sentences into a single future context and evaluate performance accordingly. Therefore, aggregating the future sentences does not increase the number of inference steps.

Figure 14 illustrate the effect of the aggregation of future sentences. Aggregating sampled future sentences into one future context often improves performance, but occasionally may degrade it. We believe that this is due to the presence of sampled future sentences that do not contribute to hallucination detection and may instead degrade performance. For instance, consider a scenario where the current sentence is accurate, but the sampled future sentence contains hallucinated information or

LLM	SelfCheckGPT		SC-ChatGPT		SC-GPT4		SC-LLaMA		SC-Vicuna		True-False	
	IS	HR	IS	HR	IS	HR	IS	HR	IS	HR	IS	HR
Llama 3.1	0%	0.4%	0%	0.1%	0%	0.2%	0%	0.3%	0%	0.2%	0%	11.6%
Gemma 3	0%	19.4%	0%	4.4%	0%	0.4%	0%	1.9%	0%	0.7%	100%	12.4%
Qwen 2.5	100%	16.8%	0%	15.2%	0%	11.8%	0%	14.2%	100%	17.7%	100%	24.7%

Table 7: Percentage of low-quality future sentences based on overlap, corresponding to future sentences used in Figure 7 (top row). IS (Insufficient Samples): Percentage of cases sampling fewer than half the number of requested sentences. HR (High Redundancy): Percentage of cases where more than 50% of sampled sentences are duplicates.

LLM	SelfCheckGPT		SC-ChatGPT		SC-GPT4		SC-LLaMA		SC-Vicuna		True-False	
	IS	HR	IS	HR	IS	HR	IS	HR	IS	HR	IS	HR
Llama 3.1	0%	0.1%	0%	0.1%	0%	0.1%	0%	0.1%	0%	0.6%	0%	1.2%
Gemma 3	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Qwen 2.5	100%	88.2%	0%	7.4%	0%	6.1%	0%	7.2%	0%	6.5%	0%	0.6%

Table 8: Percentage of low-quality future sentences based on overlap, corresponding to future sentences used in Figure 13. IS (Insufficient Samples): Percentage of cases sampling fewer than half the number of requested sentences. HR (High Redundancy): Percentage of cases where more than 50% of sampled sentences are duplicates.

irrelevant content. In such cases, the concatenated future context may become contaminated, providing the model with confusing clues and ultimately degrading detection performance. A simple yet effective solution is to filter out irrelevant or hallucinated sampled future sentences, as discussed in Section 6.4. This will be left as future work.

## F Effectiveness of Future Context Beyond Self-Consistency Sampling

Self-Consistency (Wang et al., 2023) typically samples multiple reasoning paths from language models prompted via chain-of-thought (CoT) and selects the final answer based on majority voting, significantly improving accuracy over greedy decoding. Notably, it also enhances performance when simply aggregating multiple sampled responses without explicit CoT prompting, demonstrating its broader applicability.

Thus, we compared how effective future context is beyond simply increasing response sampling. Figure 10 shows the performance obtained by sampling multiple responses (+e) directly. DIRECT+fe uses eight sampled responses as future context. For LLaMA 3.1, hallucination detection performance improves with more sampled responses, whereas Gemma 3 and Qwen 2.5 show minimal improvements. However, incorporating future context (+f) consistently achieves a higher upper bound across all cases. Even when controlling for the total number of sampled responses, incorporating future context significantly enhances performance. For example, on the SelfCheckGPT dataset, when fix-

Dataset	VERISCORE	VERISCORE+f
True-False	88.6	<b>91.7</b>

Table 9: Hallucination detection performance (AUROC) of VERISCORE on True-False.

ing the total number of sampled responses to 16, DIRECT+fe(2) achieves higher performance than DIRECT+e(16).

## G Enhancing Retrieval-based Methods with Future Context

Retrieval-based methods leverage external knowledge to verify factual claims and mitigate hallucinations. Min et al. (2023) propose **FACTSCORE**, a fine-grained metric for assessing factual precision by decomposing long-form text into atomic facts and verifying each fact against a reliable knowledge source. Wei et al. (2024) propose **SAFE**, which verifies atomic facts using retrieval-augmented LLMs with multi-step Google Search. Song et al. (2024) introduce **VERISCORE**, which selectively extracts context-aware, verifiable claims and verifies them through Google Search with few-shot prompting or fine-tuned open-weight models.

We also conduct a brief investigation into whether future context is beneficial in retrieval-based hallucination detection. Retrieval-based approaches leverage search results to access external knowledge, thereby benefiting from strong world knowledge. To evaluate this, we use the True-False dataset, as its sentences are structured as individ-

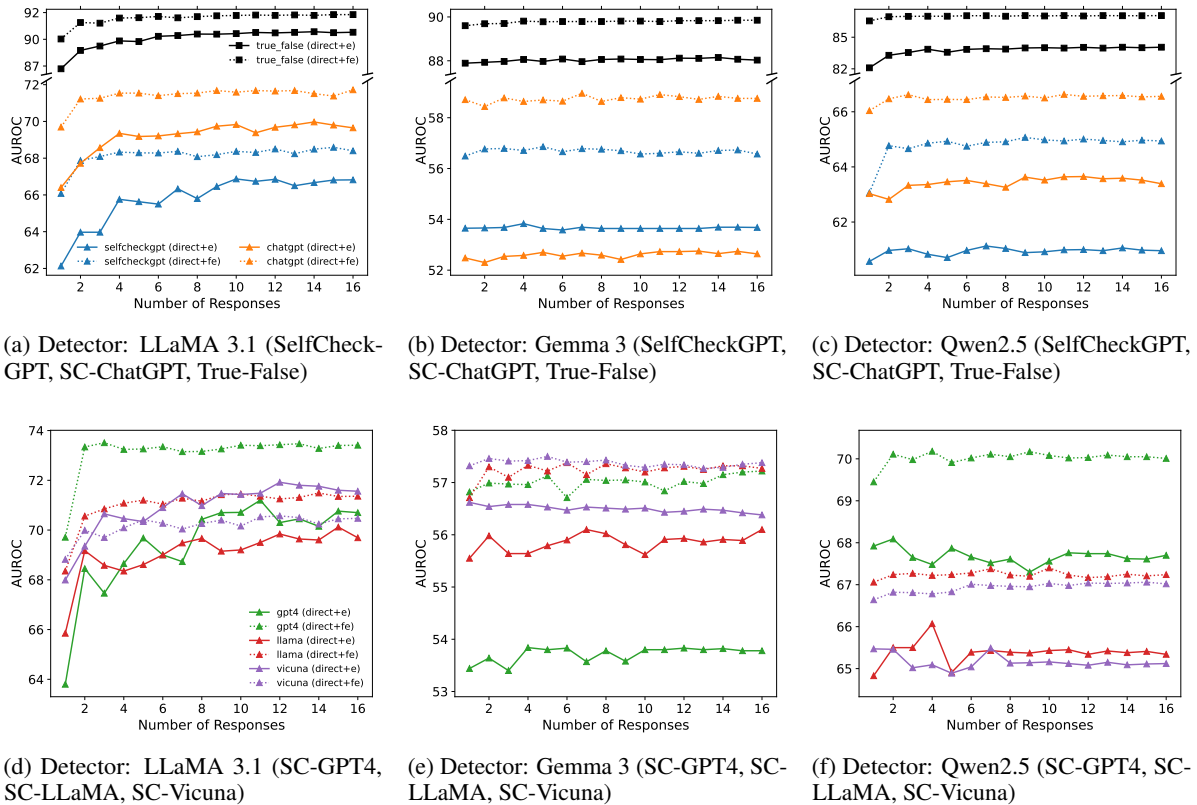


Figure 10: Hallucination detection performance of DIRECT as the number of sampled responses increases (+e), comparing settings with and without future context (+f). Incorporating future context consistently achieves higher performance than merely increasing the number of responses, highlighting the effectiveness of leveraging future context for improved hallucination detection.

ual, self-contained statements that can be treated as atomic facts, making it well-suited for assessing the factuality of discrete claims in a retrieval-based setting. Table 9 presents the performance comparison between VERIScore and VERIScore+f ( $s = 8$ ). VERIScore+f appends future context to retrieved passages from Google Search, improving performance and indicating its effectiveness for hallucination detection. Following the main experimental setup of Song et al. (2024), we use GPT-4o as the detector.

Retrieved passages can be incomplete, loosely related, or lack sufficient context for accurate detection. In such cases, future context sampled by the LLM provides complementary evidence, filling informational gaps missed by retrieved passages. A more comprehensive investigation into the broader effectiveness of future context in retrieval-based hallucination detection is left for future work.

## H Passage-Level Hallucination Detection

Following the approach used in SELFCheckGPT, the passage-level labels were generated by averag-

Detector	DIRECT	DIRECT+f
LLaMA 3.1	58.1	<b>59.5</b>
Gemma 3	51.6	<b>56.5</b>
Qwen 2.5	51.6	<b>57.1</b>

Table 10: AUROC for passage-level hallucination detection.

ing sentence-level labels. The prompts for DIRECT and DIRECT+f were identical to those described in Appendix B.2, except for replacing the term "sentence" with "passage." Table 10 clearly demonstrates that incorporating future context also leads to performance improvements at the passage level. Other datasets were not evaluated due to their primary focus on sentence-level hallucinations.

Detector	Method	Dataset					True-False	Average
		SelfCheckGPT	SC-ChatGPT	SC-GPT4	SC-LLaMA	SC-Vicuna		
LLaMA 3.1	DIRECT	33.3	90.0	91.1	90.6	88.4	84.5	79.6
	DIRECT+f	<b>36.4</b>	<b>91.5</b>	<b>92.4</b>	<b>91.5</b>	<b>89.2</b>	<b>87.1</b>	<b>81.3</b>
	SELFCKECHKGPT	42.7	92.0	94.7	92.2	90.2	88.1	83.3
	SELFCKECHKGPT+f	<b>43.0</b>	<b>94.1</b>	<b>95.3</b>	<b>93.4</b>	<b>91.1</b>	<b>89.3</b>	<b>84.3</b>
Gemma 3	SC	28.6	89.1	91.7	91.0	88.1	71.4	76.6
	SC+f	<b>31.9</b>	<b>91.4</b>	<b>94.2</b>	<b>92.5</b>	<b>90.3</b>	<b>75.8</b>	<b>79.3</b>
	DIRECT	28.6	86.3	88.8	88.0	84.6	83.8	76.7
	DIRECT+f	<b>30.1</b>	<b>87.8</b>	<b>89.5</b>	<b>88.3</b>	<b>84.9</b>	<b>84.7</b>	<b>77.6</b>
Qwen 2.5	SELFCKECHKGPT	27.4	91.7	93.8	92.0	89.3	88.8	80.5
	SELFCKECHKGPT+f	<b>32.5</b>	<b>92.5</b>	<b>94.4</b>	<b>93.2</b>	<b>90.7</b>	<b>89.7</b>	<b>82.1</b>
	SC	28.3	88.6	90.8	90.1	87.5	81.2	77.8
	SC+f	<b>29.6</b>	<b>90.1</b>	<b>91.7</b>	<b>90.8</b>	<b>88.5</b>	80.4	<b>78.5</b>
Qwen 2.5	DIRECT	32.9	89.6	92.4	90.6	87.9	80.5	79.0
	DIRECT+f	<b>34.6</b>	<b>90.8</b>	<b>93.0</b>	<b>91.4</b>	87.7	<b>83.9</b>	<b>80.2</b>
	SELFCKECHKGPT	29.6	90.1	92.6	90.5	88.5	86.5	79.6
	SELFCKECHKGPT+f	29.3	<b>90.6</b>	<b>93.3</b>	<b>91.1</b>	88.4	85.4	<b>79.7</b>
Qwen 2.5	SC	27.2	85.9	88.2	87.2	84.4	63.9	72.8
	SC+f	<b>28.9</b>	<b>88.4</b>	<b>90.6</b>	<b>89.4</b>	<b>86.0</b>	<b>64.7</b>	<b>74.7</b>

Table 11: Hallucination detection performance (AUCPR) of DIRECT, SELFCKECHKGPT, and SC with and without future context (+) across different LLM detectors. Bold numbers indicate performance improvements when future context (+) are used.

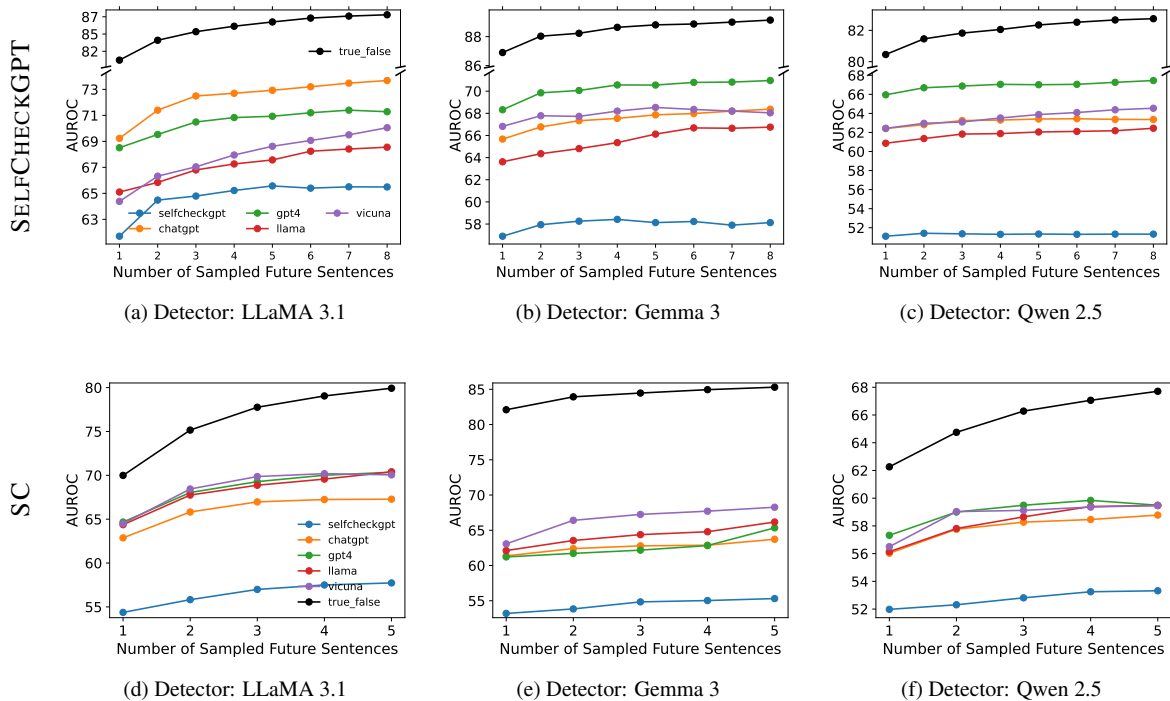


Figure 11: Hallucination detection performance of SELFCKECHKGPT and SC with future contexts by increasing the number of sampled future sentences from one turn ahead. The first row shows results for SELFCKECHKGPT, and the second row shows results for SC. Sampling more future sentences consistently improves performance across all detectors.

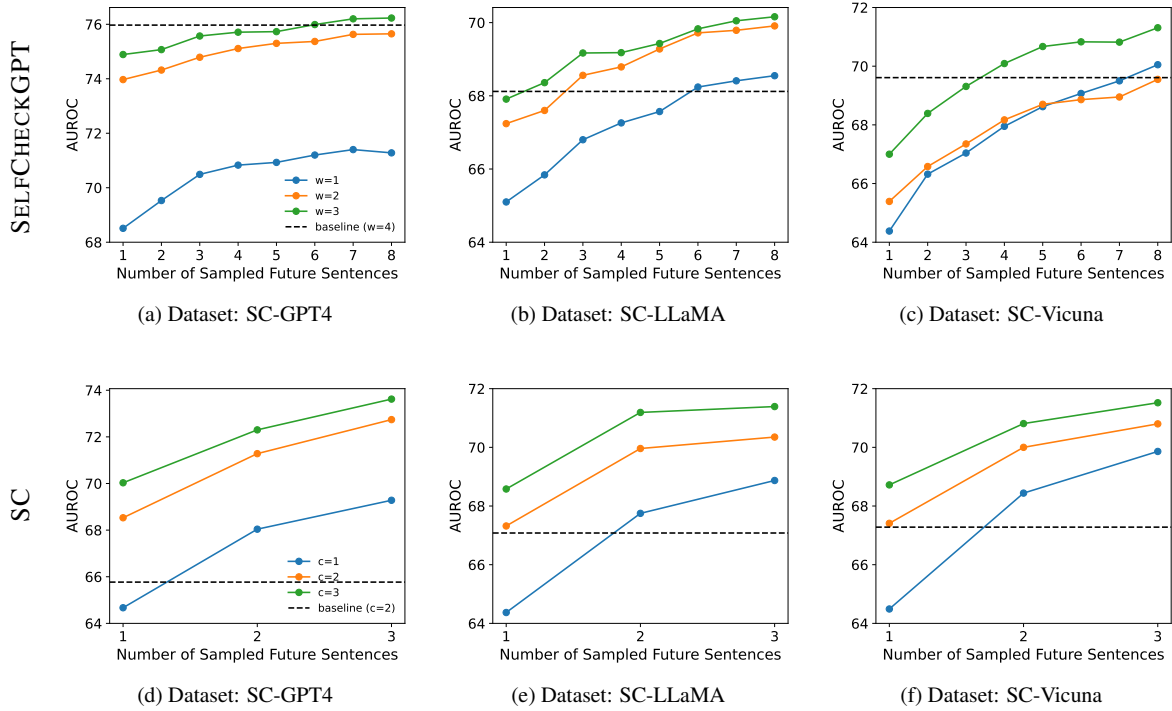


Figure 12: Hallucination detection performance of SELFCHCKGPT and SC with future contexts (detector: LLaMA 3.1). The first and second rows show performance improvements for SELFCHCKGPT and SC, respectively, when incorporating future contexts. In both cases, performance further increases as the number of sampled future contexts grows.

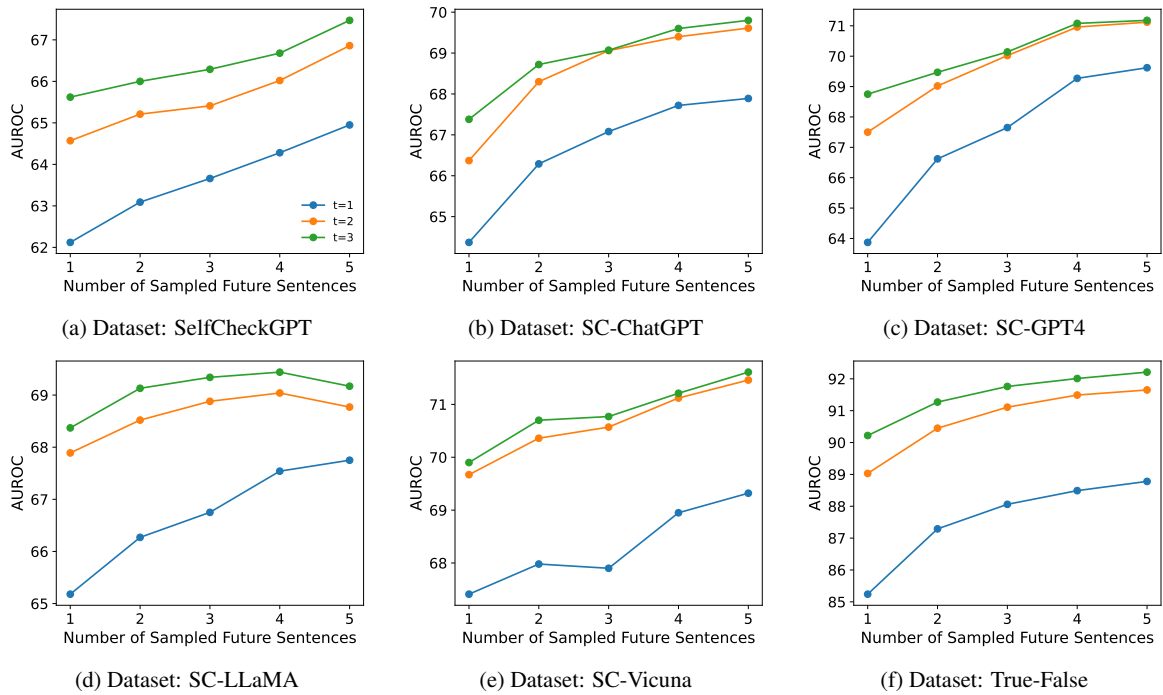


Figure 13: Hallucination detection performance of DIRECT with future contexts as both the number of sampled future sentences and the future lookahead turns increase (detector: LLaMA 3.1). Performance consistently improves as both the number of sampled sentences and the lookahead turns increase.

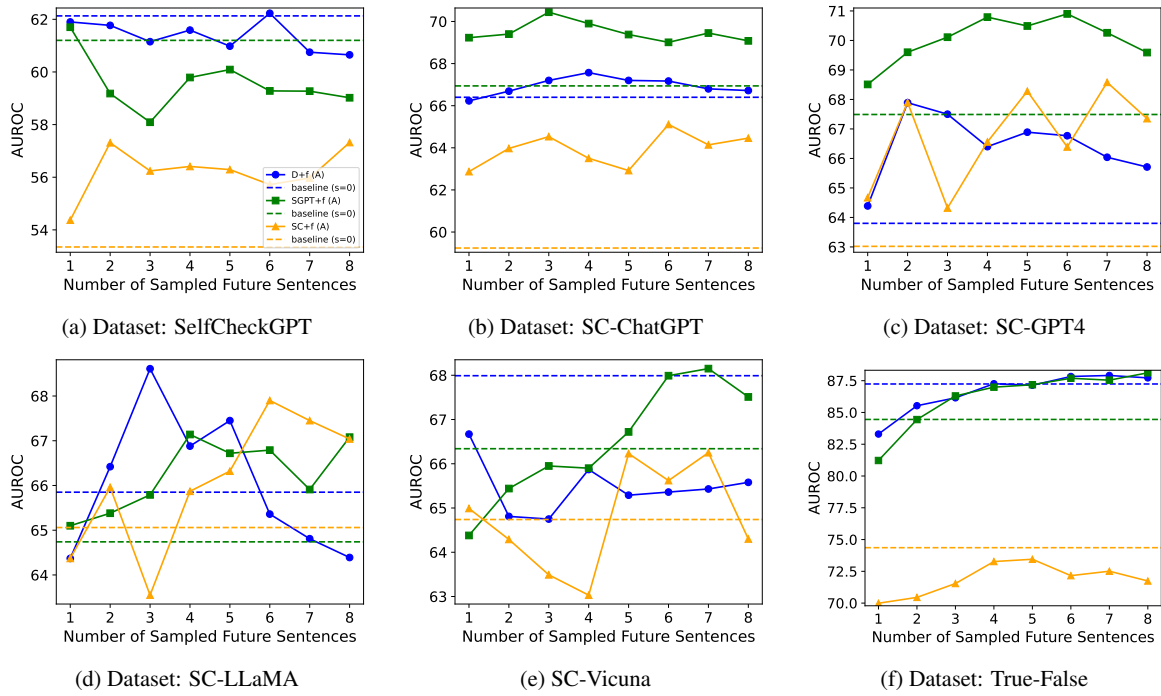


Figure 14: Hallucination detection performance of D (DIRECT), SGPT (SELF-CHECKGPT), and SC. (A) represents aggregating sampled future sentences into a single future context (detector: LLaMA 3.1). The baseline is the performance without future sentences.

Current Sentence	Future Sentence	Predicted Label	Correct with Future Sentence?
Hallucinated Sentence (H)	Sentence 1	H	Yes
	Sentence 2	NH	Yes
	Sentence 3	H	Yes
	Sentence 4	NH	No
	Sentence 5	H	Yes
	Sentence 6	NH	No
	Sentence 7	H	No
	Sentence 8	H	Yes
	Sentence 9	H	Yes
	Sentence 10	NH	No
All Future (H) = 6/10, Positive Future (H) = 5/6			
Non-Hallucinated Sentence (NH)	Sentence 1	NH	Yes
	Sentence 2	NH	No
	Sentence 3	NH	No
	Sentence 4	H	No
	Sentence 5	NH	Yes
	Sentence 6	NH	No
	Sentence 7	NH	No
	Sentence 8	H	Yes
	Sentence 9	NH	No
	Sentence 10	NH	No
All Future (NH) = 2/10, Positive Future (NH) = 1/2			

Table 12: Illustrative example of hallucination propagation corresponding to Figure 5, with 10 future sentences. H and NH denote hallucinated and non-hallucinated labels, respectively.