

# Small Agent Can Also Rock! Empowering Small Language Models as Hallucination Detector

Xiaoxue Cheng<sup>1\*</sup>, Junyi Li<sup>2\*</sup>, Wayne Xin Zhao<sup>1†</sup>, Hongzhi Zhang<sup>4</sup>,  
Fuzheng Zhang<sup>4</sup>, Di Zhang<sup>4</sup>, Kun Gai<sup>4</sup> and Ji-Rong Wen<sup>1,3</sup>

<sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China

<sup>2</sup>Department of Computer Science, National University of Singapore

<sup>3</sup>School of Information, Renmin University of China <sup>4</sup>Kuaishou

chengxiaoxue@ruc.edu.cn junyi\_cs@nus.edu.sg batmanfly@gmail.com

## Abstract

Hallucination detection is a challenging task for large language models (LLMs), and existing studies heavily rely on powerful closed-source LLMs such as GPT-4. In this paper, we propose an autonomous LLM-based agent framework, called **HaluAgent**, which enables relatively smaller LLMs<sup>1</sup> (e.g., Baichuan2-Chat 7B) to actively select suitable tools for detecting multiple hallucination types such as text, code, and mathematical expression. In HaluAgent, we integrate the LLM, multi-functional toolbox, and design a fine-grained three-stage detection framework along with memory mechanism. To facilitate the effectiveness of HaluAgent, we leverage existing Chinese and English datasets to synthesize detection trajectories for fine-tuning, which endows HaluAgent with the capability for bilingual hallucination detection. Extensive experiments demonstrate that only using 2K samples for tuning LLMs, HaluAgent can perform hallucination detection on various types of tasks and datasets, achieving performance comparable to or even higher than GPT-4 without tool enhancements on both in-domain and out-of-domain datasets. We release our dataset and code at <https://github.com/RUCAIBox/HaluAgent>.

## 1 Introduction

Recently, large language models (LLMs) (Zhao et al., 2023) have demonstrated exceptional capabilities across a variety of tasks within the field of natural language processing. However, *hallucination* (Ji et al., 2023; Rawte et al., 2023; Zhang et al., 2023; Huang et al., 2023; Ye et al., 2023) in text generated by LLMs remains an underlying concern, impeding the application of LLMs in real-world scenarios (Kaddour et al., 2023). Xu et al. (2024) have indicated that, despite the existence of some

effective hallucination mitigation strategies, the occurrence of hallucinations in LLMs is inevitable. Therefore, reliable and effective hallucination detection methods are necessary and urgent.

Existing hallucination detection methods can be roughly categorized into two primary approaches. One line of work relies on the internal knowledge of LLMs to directly identify hallucinations via prompts (Li et al., 2023; Lei et al., 2023) or evaluating the semantic consistency among multiple responses to the same question generated by LLMs (Manakul et al., 2023). However, these methods are usually constrained not only by LLMs’ internal knowledge but also by their abilities to utilize knowledge. Another line of work extends the detection ability of LLMs by employing external tools (e.g., search engine) to obtain supporting evidence for hallucination detection (Chern et al., 2023; Wei et al., 2024; Min et al., 2023). However, these approaches mostly depend on closed-source powerful LLMs such as GPT-4. Moreover, their detection process is usually pre-determined by human, making it difficult for the model to autonomously and effectively execute the hallucination detection.

To address these issues, in this paper, we propose the **HaluAgent**, an autonomous LLM-based agent framework for hallucination detection. This agent is based on smaller open-source models (i.e., Baichuan2-Chat 7B and 13B (Yang et al., 2023a)) and capable of *bilingual* hallucination detection in Chinese and English. The motivations are twofold: (1) designing autonomous detection agents that can actively make decisions and judgements, without human assistance; (2) enabling relatively smaller models to effectively perform complex detection, without reliance on close-sourced LLM APIs. To achieve this, we make three major technical contributions. First, we extend the LLM’s capability to detect a broader range of hallucination forms such as text, code, math expression, or their combination by curating a multi-functional toolbox,

\*Equal contributions

†Corresponding author

<sup>1</sup> “Smaller” in this work is relative to larger language models (e.g., over 100B).

Methods	Base Model	Task Agnostic	Tool Usage	Fine Grained	Extensibility
SelfCheckGPT	ChatGPT	✗	✗	✗	✗
SAFE	GPT-4	✗	✓	✓	✗
FacTool	GPT-4	✗	✓	✓	✗
HaluAgent	Baichuan2-Chat	✓	✓	✓	✓

Table 1: Comparison of different methods. **Task Agnostic** means whether the method is designed to specific tasks; **Fine Grained** describes whether providing detailed hallucination sentences; **Extensibility** means whether the method can extend to more tasks and tools.

in contrast to previous work only focused on textual hallucination or limited tools (Manakul et al., 2023; Min et al., 2023). Second, we design a fine-grained three-stage detection framework along with memory mechanism, including sentence segmentation, tool selection and verification, and reflection. Third, we leverage existing hallucination datasets to synthesize detection trajectories for fine-tuning the LLM, where we first employ GPT-4 to execute the above three stages until obtaining detection results consistent with the ground-truth label and then synthesize the instruction data. We compare HaluAgent and previous work in Table 1.

To verify the effectiveness, we evaluate HaluAgent on both in-domain and out-of-domain datasets at response- and sentence-level granularities. After fine-tuning with only 2K trajectories, the detection performance of HaluAgent has been improved significantly (*e.g.*, overall accuracy increases from 46.44% to 79.70% in four in-domain datasets and from 49.50% to 78.43% in two out-of-domain datasets), reaching a level comparable to or even higher than GPT-4 without tool enhancement. In the sentence-level detection experiments, HaluAgent also achieved substantial improvements, particularly with F1 scores on the math and science datasets increasing from 19.51% to 68.80% and from 17.54% to 94.16%, respectively.

## 2 Approach

In this section, we introduce **HaluAgent**, our proposed autonomous agent for detecting hallucinations across various text types. The core of our HaluAgent framework is a well-instructed LLM, which can autonomously leverage tools to detect a broader range of hallucination types. First, we define several tasks for hallucination detection and then design a toolbox with supporting tools to extend the LLM’s capability. To enable step-by-step detection, we design a three-stage detection frame-

work equipped with memory mechanism. Finally, we synthesize high-quality detection trajectory data to fine-tune open-source LLMs. We present the overall architecture of HaluAgent in Figure 1.

### 2.1 Task Definition

Hallucination refers to seemingly plausible yet factually unsupported content (Huang et al., 2023). Unlike previous work that mostly focused on detecting text-based hallucinations (Manakul et al., 2023; Yehuda et al., 2024), we consider a broader range of hallucination forms such as text, code, and mathematical expression. Hence, in this work, we conduct hallucination detection in five tasks as:

- **Knowledge-based QA** involves generating answers to the input question (Lan et al., 2021), and we aim to identify misinformation from the answer text such as incorrect historical dates, misattributed quotes, or false scientific facts.

- **Conditional Text Generation** focuses on tasks with specific requirements given in the input instructions (Dathathri et al., 2020) such as generating text with particular length, format, or translating a paragraph into special language. We aim to detect hallucinations that deviate from the given instructions or include irrelevant information.

- **Semantic Consistency** is not specific to any particular type of text or task. We define hallucination in semantic consistency as those irrelevant or self-contradictory content in the responses.

- **Math Problem Solving** is related to generating a series of mathematical expressions to solve the math problem (Hendrycks et al., 2021), while the expressions may contain computational errors, *e.g.*, arithmetic or calculation mistakes.

- **Code Generation** aims to generate code snippets for the input query (Chen et al., 2021). We define hallucinations as code snippets that are syntactically incorrect, fail to execute as intended, or contain logical flaws and missing dependencies.

Note that in real-world scenarios the responses from LLMs may contain a mixture of these hallucination types, and we aim to develop a general and versatile hallucination detection agent that can deal with a broader range of hallucinations.

### 2.2 Toolbox for Hallucination Detection

Since detecting texts with a mix of hallucination types is challenging, we design a comprehensive toolbox to enhance the hallucination detection capability of LLMs following previous work (Chern

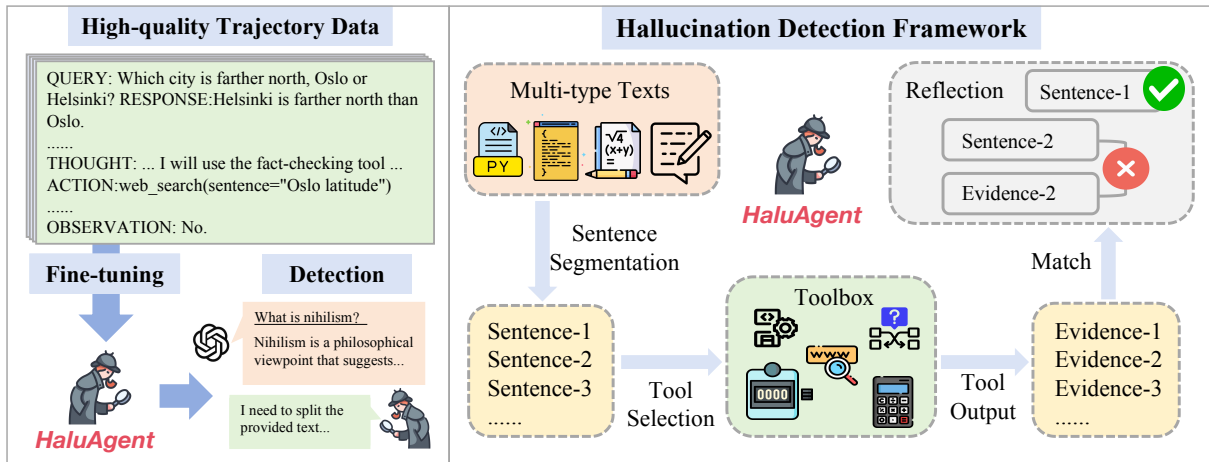


Figure 1: The overview of our proposed HaluAgent. The left part shows the process of fine-tuning open-source models and detecting hallucinations. The right part illustrates the hallucination detection pipeline of HaluAgent.

et al., 2023; Gou et al., 2024). Based on the hallucination types discussed in Section 2.1, we incorporate five types of external tools, *i.e.*, search engine, calculator, code interpreter, condition verifier, and semantic checker, and two internal system tools:

- **Search Engine** is utilized to retrieve supporting evidence from the web for identifying factually incorrect content, defined as *web\_search*. We use Google Programmable Search Engine API<sup>2</sup> to implement this tool. Considering the relevance between retrieved documents and the output text, we only use the top-5 documents as the most relevant retrieval results for hallucination detection.

- **Calculator** is used to verify inaccurate mathematical calculations in the model responses, targeting math-related hallucinations, defined as *calculator*. We implement calculator via the scientific computing library SymPy (Meurer et al., 2017).

- **Code Interpreter** can be used to validate code snippets by executing the snippets in a programming environment, defined as *code\_interpreter*, ensuring that the code is syntactically correct and functions as intended. We implement the code interpreter tool following CRITIC (Gou et al., 2024).

- **Condition Verifier** aims to detect hallucinations for conditional text generation by assessing whether the text is consistent with the given condition. For example, we utilize word counter (*word\_counter*) to compute the number of words in a length-constrained generation scenario.

- **Semantic Checker** mainly addresses hallucination types such as irrelevant responses and self-contradictions, which is defined as *match*. We leverage GPT-4 to examine the consistency and

relevance of semantics, primarily handling the semantic matching scenario.

- **System Tools** are developed to support the basic manipulation operations for hallucination detection, including sentence segmentation (*split\_text*) and returning detection results (*get\_answer*).

All tools are defined as functions in a unified way and we present the whole toolbox in Table 5 in Appendix A.2. It is worth noting that the toolbox is not limited to the existing tools and can be easily extended, *e.g.*, adding more task-specific tools.

### 2.3 HaluAgent Framework

Inspired by prior work on complex reasoning (Jiang et al., 2024), we consider the hallucination detection process as an agent task. Specifically, HaluAgent includes three stages: sentence segmentation, tool selection and verification, and reflection, along with memory mechanism. Below, we outline the workflow and components of our HaluAgent framework.

**Sentence Segmentation.** Since the responses of LLMs are usually the combination of facts, opinions, and various types of texts, we first segment the responses into several independent detection units. To be specific, we utilize the system tool, *split\_text*, to perform sentence segmentation. This tool requires the agent to split the input text into a set of sentences and complete any incomplete semantic information within the sentences, *e.g.*, pronouns and omitted content. Sentence segmentation reduces the complexity of hallucination detection tasks and allows HaluAgent to tailor its detection strategies for individual sentences, ensuring more accurate and fine-grained results by minimizing

<sup>2</sup><https://developers.google.com/custom-search>

interference from unrelated content.

**Tool Selection and Verification.** Next, HaluAgent verifies each sentence separately by selecting the appropriate tool from the toolbox based on the type of sentence content (e.g., *web\_search* for fact statements and *calculator* for mathematical expressions). HaluAgent compares each sentence with the execution results of the selected tool to identify any inconsistencies or inaccuracies, thereby detecting hallucinations. To prevent the agent from forgetting intermediate detection results and support the subsequent reflection, HaluAgent stores the detection result of each sentence as a triple, i.e., (*sentence*, *hallucination label*, *supporting evidence*). If the sentence is identified as containing hallucinations, it will be labeled “1”, otherwise “0”. These useful information will be stored based on a memory mechanism, allowing HaluAgent to refer to historical results and maintain a consistent understanding of the text veracity throughout the detection process.

**Reflection.** After individually examining all the sentences, HaluAgent can obtain preliminary detection results. However, due to the limited capacity of each tool and potential errors in their outputs, these detection results might not be fully accurate. To address this, HaluAgent performs the final reflection to double-check whether the previous detection results are correct from *local* and *global* perspectives. At the local level, HaluAgent will match each sentence with the corresponding evidence to ensure the local correctness of hallucination detection. However, different sentences may influence each other. Therefore, at the global level, HaluAgent will determine whether the current sentence is incorrect based on the context of other sentences. For instance, if the calculation result in a preceding sentence is incorrect, any subsequent steps based on this result should be considered incorrect, even if these steps are correct when checked in isolation. Any detection mistakes will be corrected and the detection results (i.e., hallucination label and supporting evidence) stored in memory are updated accordingly. After reflection, HaluAgent invokes a system tool, *get\_answer*, to output the final detection result. If any hallucinations are detected, HaluAgent outputs the specific sentences and supporting evidence from the detection tools.

Throughout the above process, HaluAgent first segments the input text into a set of semantically

Datasets	#Train	#Filtered Trajectory	#Test
WebQA	900	675	100
Ape210K	500	334	100
HumanEval	100	100	63
WordCnt	100	100	100
HaluEval-QA	900	808	100
All	2500	2017	463

Table 2: Statistics of synthetic detection trajectories.

complete sentences, then selects tools to check each sentence individually, and finally reflects on the detection results to further correct mistakes. To support this process, we use memory mechanism to store useful information such as historical detection trajectories and current detection results.

## 2.4 Bilingual Agent Tuning

Previous studies mostly depended on closed-source LLMs (e.g., GPT-4) to detect hallucinations (Chern et al., 2023; Wei et al., 2024). To empower smaller language models as effective hallucination detectors, we aim to perform supervised fine-tuning on smaller LLMs (e.g., Baichuan2-Chat 7B). Given the powerful agent capability, we leverage GPT-4 to synthesize high-quality hallucination detection trajectory data in Chinese and English following the detection framework in Section 2.3.

### 2.4.1 Trajectory Generation

**Data Source.** To curate high-quality trajectory data covering diverse hallucination types, we select and construct five datasets, i.e., HaluEval (Li et al., 2023) and WebQA (Li et al., 2016) for knowledge-based QA, Ape210K (Zhao et al., 2020) for math word problem, HumanEval (Chen et al., 2021) for code generation, and WordCnt for conditional text generation. Among them, HaluEval and HumanEval are English datasets, and WebQA, Ape210K, and WordCnt are Chinese datasets, which enable bilingual detection capabilities of HaluAgent. In our experiments, we use GPT-4 to synthesize WordCnt which is targeted at generating text with a specified length. Besides, we obtain the ground-truth hallucination labels for WebQA and Ape210K by ChatGPT and human annotators. We provide details of each dataset in Appendix A.3.

**Trajectory Format.** Based on our datasets, we employ GPT-4 to execute the detection process in Section 2.3 and generate corresponding trajectory data following the ReAct format (Yao et al., 2023). We begin by feeding the detection instruction and

the text to be detected as input for GPT-4. At each turn, the agent receives an *observation*, makes its plans and thoughts as *thought*, and invokes corresponding tools through *action*. The results from these tools are formulated as a new observation for the next turn. By iterating the above process, we can obtain a complete detection trajectory comprising the input instruction, detected text, intermediate steps (*i.e.*, observations, thoughts, actions), and the final detection result. To ensure the accuracy of the trajectories, we remove those samples that include wrong tool invocation, formatting errors, and inconsistency between the detection result and the ground-truth hallucination label. Finally, we produce 2,017 high-quality trajectories for supervised fine-tuning. We present an example in Figure 3 and the statistics of trajectory data in Table 2.

### 2.4.2 Trajectory Tuning

Based on the above formatted bilingual trajectory data, we perform supervised fine-tuning on Baichuan2-Chat 7B and 13B (Yang et al., 2023a), which are much smaller than the backbone models in previous studies (Chern et al., 2023; Wei et al., 2024). Formally, the hallucination detection trajectory for each sample can be represented as  $\langle o_0, t_1, a_1, o_1, t_2, a_2, \dots, o_{n-1}, t_n, a_n, o_n \rangle$ , where  $o_i$ ,  $t_i$ , and  $a_i$  denote the observation, thought, and action at the  $i$ -th turn, respectively. Specifically,  $o_0$  denotes the initial observation consisting of the input instruction and detected text, and  $o_n$  denotes the final detection result. At each turn, based on the historical trajectory  $c_i = \langle o_0, t_1, a_1, \dots, o_{i-1} \rangle$ , the agent aims to generate thought  $t_i$  and action  $a_i$ . Therefore, during the trajectory fine-tuning process, we only compute the cross-entropy loss for  $t_i$  and  $a_i$  while masking  $o_i$ :

$$\mathcal{L} = -\log \sum_{i=1}^n \Pr(t_i, a_i | c_i). \quad (1)$$

## 2.5 Comparison to Previous Work

To clarify the differences between HaluAgent and other hallucination detection methods, we aim to address the following two questions:

- **What are the benefits of designing an autonomous hallucination detection agent?** Hallucination detection is fundamental to related research. Most previous work either relied on the internal knowledge of LLMs (might be limited) (Cohen et al., 2023) or performed coarse-grained detection process (Manakul et al., 2023). Designing an

agent for hallucination detection offers a flexible alternative that can effectively extend the detection capability of LLMs through tool utilization. Existing agent-based detection methods do not consider tool utilization or only employ limited tools such as search engine for hallucination detection in long texts (Lei et al., 2023; Wei et al., 2024). In contrast, HaluAgent develops a comprehensive and extensible toolbox and performs fine-grained reasoning process for hallucination detection, providing a more adaptable and robust solution.

- **Can smaller language models perform well in challenging hallucination detection?** Existing methods (Chern et al., 2023; Manakul et al., 2023; Dhuliawala et al., 2023) heavily depended on powerful closed-source LLMs such as GPT-4, which leads to high computational costs and poses unavoidable limitations for the practical deployment of these technologies. Moreover, relying on closed-source models makes the detection results difficult to reproduce. Our work demonstrates that by incorporating the agent capabilities and tool integration, smaller language models can also effectively handle challenging hallucination detection tasks. This way can provide a more viable and economical choice, significantly reducing the need for closed-source LLMs.

## 3 Experiments

### 3.1 Experimental Setup

**In-domain/Out-of-domain Tasks.** We evaluate HaluAgent on both in-domain and out-of-domain datasets. As described in Section 2.4.1, we select HaluEval-QA, WebQA, Ape210K, HumanEval, and WordCnt as in-domain datasets. For out-of-domain datasets, we use a Chinese dataset, HaluQA (Cheng et al., 2023), and an English dataset, HaluEval 2.0 (Li et al., 2024), which cover diverse hallucination detection scenarios including knowledge, math, science texts. All datasets are associated with ground-truth response-level hallucination labels. We present the details in Appendix B.1.

**Sentence-level Tasks.** HaluAgent performs fine-grained hallucination detection by segmenting sentences and provides sentence-level detection results. We use FacTool (Chern et al., 2023) with annotated claim-level hallucination labels to evaluate the fine-grained detection capability. We consider each claim as a sentence and concatenate all claims as the input text. FacTool contains five sub-datasets:

Types	Datasets	GPT-4		Baichuan2-Chat		HaluAgent	
		prompt	pipeline	7B	13B	7B	13B
In-domain Datasets	WebQA	82.00/35.71	<u>91.00/57.14</u>	51.00/14.04	54.00/ <b>61.67</b>	80.00/54.55	<b>82.83/51.43</b>
	Ape210K	72.33/74.21	<u>76.63/75.10</u>	49.00/7.27	51.33/58.29	72.00/72.55	<b>73.40/73.68</b>
	HumanEval	71.43/79.07	<u>93.44/94.12</u>	34.92/49.38	47.62/19.51	<b>93.44/94.12</b>	<b>93.44/94.12</b>
	WordCnt	56.00/66.15	<u>100.00/100.00</u>	43.00/16.00	46.00/59.70	<b>100.00/100.00</b>	<b>100.00/100.00</b>
	HaluEval-QA	62.00/42.42	<u>77.53/75.61</u>	53.19/46.34	60.00/67.74	67.00/67.33	<b>71.00/72.38</b>
	Overall	69.76/71.66	<u>85.10/83.12</u>	46.44/32.20	52.70/55.58	79.70/79.50	<b>81.86/80.69</b>
Out-of-domain Datasets	HalluQA	61.00/74.84	<u>85.11/89.23</u>	33.00/12.99	56.00/67.16	67.48/76.09	<b>78.16/83.75</b>
	HaluEval 2.0	63.00/74.13	<u>85.71/87.36</u>	54.00/69.33	43.00/46.73	75.00/76.19	<b>79.00/78.79</b>
	Overall	62.00/74.50	<u>85.25/88.76</u>	43.50/50.22	49.50/58.10	69.97/76.12	<b>78.43/82.45</b>

Table 3: Evaluation results at Accuracy and F1 score on in-domain and out-of-domain datasets. **Bold** denotes the best methods among open-source models; underline denotes the best methods among closed-source models.

Chinese-QA, KB-QA, math problems, code generation, and scientific literature review. Since the code generation sub-dataset is collected from HumanEval and overlaps with our training data, we conduct sentence-level detection experiments on the other four sub-datasets.

**Baselines.** Unlike previous work relied on powerful closed-source models like GPT-4, HaluAgent is built upon relatively smaller language model by performing fine-tuning on hallucination detection trajectory data. Hence, we compare HaluAgent with two kinds of baselines: (1) Closed-source models: *GPT-4 prompt* and *GPT-4 pipeline* employ GPT-4 with simple task description prompts and our proposed detection pipeline, respectively; (2) Open-source models: *Baichuan2-Chat (7B and 13B)*, which is the backbone model of HaluAgent. For other detection methods without comparison in this work, they are either implemented based on ChatGPT/GPT-4 (an unfair comparison) or focus solely on specific tasks, making it difficult to adapt to other tasks, as shown in Table 1.

**Metrics.** Hallucination detection is essentially a binary classification task. Consequently, we adopt *Accuracy* and *F1 score* as metrics for response-level detection evaluation. Considering the imbalanced hallucination data distribution at the sentence level, we adopt *Accuracy*, *Precision*, *Recall*, and *F1 score* for sentence-level detection.

### 3.2 Response-level Detection

We present the evaluation results on in-domain and out-of-domain datasets in Table 3.

First, by comparing GPT-4 and Baichuan2-Chat, we can observe a large performance gap between

closed-source and open-source models when detecting hallucinations solely based on their internal knowledge. In Table 3, GPT-4 achieves detection accuracies of 82.00% and 72.33% on WebQA and Ape210K respectively, whereas Baichuan2-Chat only achieves 51.00% and 49.00%. This underscores the considerable difference in hallucination detection capabilities between these models.

Second, implementing a fine-grained detection framework and incorporating extensive tools can enhance the hallucination detection performance of LLMs across various tasks. Compared to GPT-4 prompt, GPT-4 pipeline consistently yields better detection results across all datasets. Especially for tasks that require tools to detect the hallucinations such as WordCnt, our tool-assistant framework boosts the detection accuracy of GPT-4 from 56.00% to 100.00%. This demonstrates that guiding the model to utilize appropriate tools tailored to the specific text is an effective strategy to enhance hallucination detection capabilities.

Finally, based on trajectory fine-tuning, smaller open-source models can be effective autonomous agents for hallucination detection, narrowing the gap with closed-source models. As can be observed from Table 3, HaluAgent consistently improves detection performance across in-domain datasets compared to Baichuan2-Chat. For instance, on WebQA and HaluEval-QA datasets, HaluAgent 7B and 13B improve the detection accuracy from 51.00%, 60.00% to 80.00%, 71.00%, respectively. Furthermore, substantial performance improvements are observed on out-of-domain datasets, with accuracy increasing from 33.00% to 67.48% (7B) for HalluQA and from 43.00% to 79.00% (13B) for HaluEval 2.0. With tool utilization and agent capabilities,

Models	Chinese-QA	KB-QA	Math	Science
GPT4-prompt	54.93/59.40/33.17/42.57	75.97/39.29/50.00/44.00	55.27/49.25/23.74/32.04	59.14/64.71/81.82/72.26
GPT4-pipeline	79.76/64.71/50.00/56.41	84.12/85.31/93.21/89.09	91.61/79.66/78.33/78.99	95.72/96.77/98.04/97.40
Baichuan2-Chat 7B	65.15/15.02/29.17/19.83	49.77/38.78/19.59/26.03	68.07/17.91/21.43/19.51	24.19/9.80/83.33/17.54
Baichuan2-Chat 13B	56.82/21.70/22.67/22.17	65.41/30.77/30.00/30.38	70.97/13.85/20.93/16.67	18.82/1.31/100.00/2.58
HaluAgent-7B	73.57/32.96/55.97/41.49	83.33/41.82/82.14/55.42	85.39/81.13/59.72/68.80	90.27/94.77/93.55/94.16
HaluAgent-13B	75.47/33.80/56.25/42.23	81.97/62.07/85.71/72.00	87.50/64.52/72.73/68.38	92.39/96.05/94.81/95.42

Table 4: Evaluation results of sentence-level detection on the four subsets of FacTool.

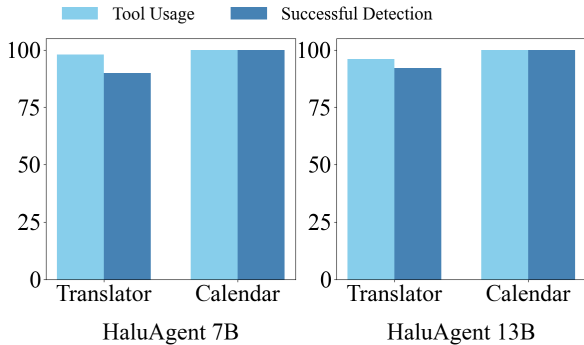


Figure 2: The usage rate of new tools and the proportion of successful detection.

HaluAgent achieves performance comparable to or even higher than GPT-4 prompt across all tasks, showing strong generalization capability.

### 3.3 Sentence-level Detection

In our framework, HaluAgent is capable of detecting hallucinations for each individual sentence. Thus, we evaluate the accuracy of HaluAgent in identifying hallucinations at the sentence level.

The sentence-level detection results are shown in Table 4. As we can see, HaluAgent achieves much higher F1 score compared to Baichuan2-Chat, especially for those tasks where the detection results can be precisely judged via tools. For example, HaluAgent 13B achieves an accuracy of 87.50% on the math dataset, and F1 score of 95.42% on the scientific literature review dataset, approaching the performance of the GPT-4 pipeline. We attribute this improvement to the fine-grained design of the HaluAgent framework and its capability to store and reflect on detection results for each sentence.

### 3.4 Further Analysis

**Extensibility Study.** To verify the extensibility of our HaluAgent framework, we introduce new tools to the fine-tuned models and test their abilities to use these tools for hallucination detection. Specifically, we incorporate a translator and a calendar tool into the HaluAgent toolbox for handling texts

related to translation and date calculations. We provide instructions and two examples as in-context demonstrations to guide the model to use these new tools. For evaluation, we employ ChatGPT to generate 100 samples involving translation and date calculation. Instruction and dataset details can be found in Appendix B.3.

We measure the proportion of correct tool usage and successful task completion rate by HaluAgent in Figure 2. As we can see, HaluAgent 7B and 13B achieve a usage rate of over 95% for the translator tool, and the usage rate for the calendar tool is 100%. Due to the models’ inherent multilingual capabilities, they sometimes directly assess the accuracy of the translation result instead of invoking the translator tool. With the assistance of the appropriate tools, HaluAgent achieves a success rate of 100% in hallucination detection for date calculation and over 90% success rate for translation task. These experimental results indicate that with instructions and demonstrations, HaluAgent can effectively use new tools to complete hallucination detection tasks without additional fine-tuning.

**Case Study.** To qualitatively demonstrate the effectiveness of HaluAgent, we present a case study about average speed calculation, which involves both commonsense knowledge and mathematical calculations. We compare HaluAgent with two hallucination detection methods based on GPT-4 with simple prompts and search engine tool, as shown in Figure 3. As we can see, GPT-4 fails to detect the minor calculation errors (*i.e.*,  $3 \text{ hours} + 0.5 \text{ hours} + 1.5 \text{ hours} = 4 \text{ hours}$ ) when only provided with a simple prompt of task description. Meanwhile, GPT-4 with a search engine tool can only verify the formula for average speed in the response but cannot check the correctness of each calculation step. Consequently, both methods yield incorrect detection results. In contrast, HaluAgent autonomously plans and selects suitable tools for different parts of the text (*i.e.*, search engine for knowledge checks and calculator for calculation checks), allowing for

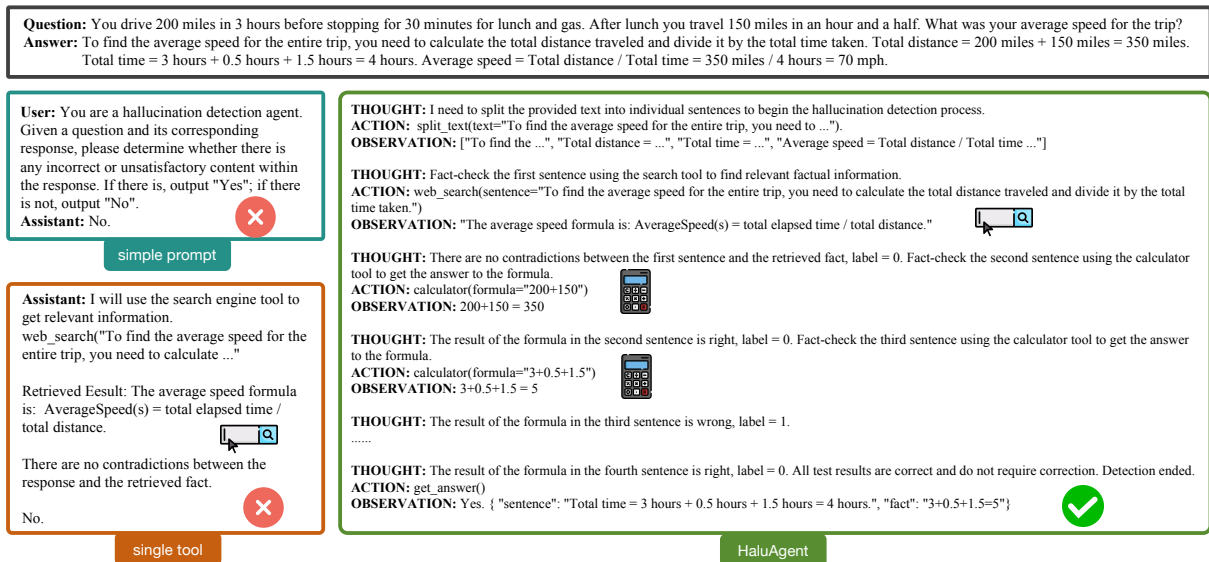


Figure 3: Case study between GPT-4 with a simple prompt and single tool, and the HaluAgent framework.

accurate verification of the content. This highlights HaluAgent’s capability to effectively verify complex texts with mixed types of hallucinations.

## 4 Related Work

**Hallucination Detection.** Hallucination Detection in LLMs is a pivotal concern due to its role in identifying inaccuracies or falsehoods within model responses. Most existing methods for hallucination detection are implemented based on powerful LLMs (Luo et al., 2024). One category of these methods relies on the internal knowledge and consistency of LLMs to detect hallucinations, by breaking down the detection process (Dhuliawala et al., 2023) or comparing multiple responses to the same query (Manakul et al., 2023). However, these approaches are inherently limited by the knowledge boundaries of LLMs. Another category of methods leverages external tools for hallucination detection (Chern et al., 2023; Wei et al., 2024). While tool utilization can compensate for the knowledge limitations of LLMs (Wang et al., 2024), existing methods often require manually designing specific tools tailored to particular text types. In contrast, HaluAgent is equipped with a versatile and expandable toolbox, enabling small language models to autonomously select appropriate tools for fine-grained detection.

**Agent Tuning.** Recently, the surprising planning and reasoning capabilities of LLMs have inspired research into their application as agents for specific tasks. Previous research (Yao et al., 2023;

Nakano et al., 2021; Singh et al., 2023; Yang et al., 2023b) has predominantly relied on prompting to use LLMs as agents, but this method demands advanced instruction-following capabilities that open-source LLMs typically do not match with API-based LLMs. To break this restriction, AgentTuning (Zeng et al., 2023) first fine-tunes open-source LLMs on agent interaction trajectories generated by powerful LLMs. Moreover, Agent-FLAN (Chen et al., 2024) goes a step further by decomposing and redesigning the training corpus. Besides general agents, KG-Agent (Jiang et al., 2024) fine-tunes LLaMA-7B model to achieve an agent specialized in reasoning over knowledge graphs. Similarly, HaluAgent is fine-tuned on trajectory data of hallucination detection tasks, thereby enhancing the detection capabilities of open-source LLMs.

## 5 Conclusion

In this work, we proposed an autonomous agent framework, HaluAgent, which is capable of bilingual hallucination detection in Chinese and English. In our approach, we first curated a multi-functional toolbox to extend the LLM’s capability to detect hallucinations. Next, we designed a fine-grained three-stage detection framework along with memory mechanism, including sentence segmentation, tool selection and verification, and reflection. Then, we leveraged existing datasets to synthesize detection trajectories by employing GPT-4 to execute the detection process following the HaluAgent framework. Finally, we fine-tuned Baichuan2-Chat 7B and 13B on the synthesized trajectories. The Halu-



Agent models achieved notable improvements on both in-domain and out-of-domain datasets, with performance comparable to or even higher than GPT-4 without tool enhancements. Due to its high flexibility and adaptability, HaluAgent can effectively serve as a hallucination detector when human users interact with LLMs in real-world scenarios. In future work, we will extend our method to deal with more types of tools and hallucinations.

## Limitations

Although HaluAgent significantly enhances the performance of small open-source models in hallucination detection tasks, our approach still has some limitations. First, we use only Baichuan2-Chat as the backbone LLM and do not compare it with other models of comparable parameter size, such as Llama2-7B (Touvron et al., 2023), Mistral-7B (Jiang et al., 2023), and Qwen-7B (Bai et al., 2023). Second, our work focuses on hallucination detection tasks and does not propose corresponding mitigation strategies. Third, we focus on detecting hallucinations that contain errors and contradictions, lacking consideration for hallucinations related to identity recognition and ethical issues. Finally, the training process of HaluAgent uses only correct trajectory data for supervised fine-tuning, without fully leveraging failed trajectory data or incorporating other types of data. In the future, we will further refine the HaluAgent framework to cover more hallucination types and fully leverage failed data to train the model. We also consider developing improved mitigation strategies based on HaluAgent.

## Acknowledgements

This work was partially supported by Beijing Natural Science Foundation under Grant No. 4222027 and L233008, National Natural Science Foundation of China under Grant No. 62222215. Xin Zhao is the corresponding author.

## References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi

Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen technical report](#). *CoRR*, abs/2309.16609.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebguss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.

Zehui Chen, Kuikun Liu, Qiuchen Wang, Wenwei Zhang, Jiangning Liu, Dahua Lin, Kai Chen, and Feng Zhao. 2024. [Agent-flan: Designing data and methods of effective agent tuning for large language models](#). *CoRR*, abs/2403.12881.

Qinyuan Cheng, Tianxiang Sun, Wenwei Zhang, Siyin Wang, Xiangyang Liu, Mozhi Zhang, Junliang He, Mianqiu Huang, Zhangyue Yin, Kai Chen, and Xipeng Qiu. 2023. [Evaluating hallucinations in chinese large language models](#). *CoRR*, abs/2310.03368.

I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, and Pengfei Liu. 2023. [Factool: Factuality detection in generative AI - A tool augmented framework for multi-task and multi-domain scenarios](#). *CoRR*, abs/2307.13528.

Roi Cohen, May Hamri, Mor Geva, and Amir Globerson. 2023. [LM vs LM: detecting factual errors via cross examination](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12621–12640. Association for Computational Linguistics.

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2023. [Chain-of-verification reduces](#)

- hallucination in large language models. *CoRR*, abs/2309.11495.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujia Yang, Nan Duan, and Weizhu Chen. 2024. [CRITIC: Large language models can self-correct with tool-interactive critiquing](#). In *The Twelfth International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *CoRR*, abs/2311.05232.
- 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.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Jinhao Jiang, Kun Zhou, Wayne Xin Zhao, Yang Song, Chen Zhu, Hengshu Zhu, and Ji-Rong Wen. 2024. [Kg-agent: An efficient autonomous agent framework for complex reasoning over knowledge graph](#). *CoRR*, abs/2402.11163.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. [Challenges and applications of large language models](#). *CoRR*, abs/2307.10169.
- Yunshi Lan, Gaole He, Jinhao Jiang, Jing Jiang, Wayne Xin Zhao, and Ji-Rong Wen. 2021. [A survey on complex knowledge base question answering: Methods, challenges and solutions](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4483–4491. ijcai.org.
- Deren Lei, Yaxi Li, Mengya Hu, Mingyu Wang, Vincent Yun, Emily Ching, and Eslam Kamal. 2023. [Chain of natural language inference for reducing large language model ungrounded hallucinations](#). *CoRR*, abs/2310.03951.
- Junyi Li, Jie Chen, Ruiyang Ren, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024. [The dawn after the dark: An empirical study on factuality hallucination in large language models](#). *CoRR*, abs/2401.03205.
- Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. [Halueval: A large-scale hallucination evaluation benchmark for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6449–6464.
- Peng Li, Wei Li, Zhengyan He, Xuguang Wang, Ying Cao, Jie Zhou, and Wei Xu. 2016. Dataset and neural recurrent sequence labeling model for open-domain factoid question answering. *arXiv preprint arXiv:1607.06275*.
- Junliang Luo, Tianyu Li, Di Wu, Michael Jenkin, Steve Liu, and Gregory Dudek. 2024. [Hallucination detection and hallucination mitigation: An investigation](#). *CoRR*, abs/2401.08358.
- Potsawee Manakul, Adian Liusie, and Mark J. F. 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, EMNLP 2023, Singapore, December 6-10, 2023*, pages 9004–9017. Association for Computational Linguistics.
- Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, Amit Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. 2017. [SymPy: symbolic computing in python](#). *PeerJ Computer Science*, 3:e103.
- Sewon Min, Kalpesh Krishna, Xinxu 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.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. [Webgpt: Browser-assisted question-answering with human feedback](#). *CoRR*, abs/2112.09332.
- Vipula Rawte, Amit P. Sheth, and Amitava Das. 2023. [A survey of hallucination in large foundation models](#). *CoRR*, abs/2309.05922.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori,

- Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code](#). *CoRR*, abs/2308.12950.
- Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. 2023. [Prog-prompt: Generating situated robot task plans using large language models](#). In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11523–11530. IEEE.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Yuhao Wang, Ruiyang Ren, Junyi Li, Wayne Xin Zhao, Jing Liu, and Ji-Rong Wen. 2024. [Rear: A relevance-aware retrieval-augmented framework for open-domain question answering](#). *arXiv preprint arXiv:2402.17497*.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). *CoRR*, abs/2403.18802.
- Ziwei Xu, Sanjay Jain, and Mohan S. Kankanhalli. 2024. [Hallucination is inevitable: An innate limitation of large language models](#). *CoRR*, abs/2401.11817.
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng Liu, Guangwei Ai, Guosheng Dong, Haizhou Zhao, Hang Xu, Haoze Sun, Hongda Zhang, Hui Liu, Jiaming Ji, Jian Xie, Juntao Dai, Kun Fang, Lei Su, Liang Song, Lifeng Liu, Liyun Ru, Luyao Ma, Mang Wang, Mickel Liu, MingAn Lin, Nuolan Nie, Peidong Guo, Ruiyang Sun, Tao Zhang, Tianpeng Li, Tianyu Li, Wei Cheng, Weipeng Chen, Xiangrong Zeng, Xiaochuan Wang, Xiaoxi Chen, Xin Men, Xin Yu, Xuehai Pan, Yanjun Shen, Yiding Wang, Yiyu Li, Youxin Jiang, Yuchen Gao, Yupeng Zhang, Zenan Zhou, and Zhiying Wu. 2023a. [Baichuan 2: Open large-scale language models](#). *CoRR*, abs/2309.10305.
- Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023b. [MM-REACT: prompting chatgpt for multimodal reasoning and action](#). *CoRR*, abs/2303.11381.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *International Conference on Learning Representations (ICLR)*.
- Hongbin Ye, Tong Liu, Aijia Zhang, Wei Hua, and Weiqiang Jia. 2023. [Cognitive mirage: A review of hallucinations in large language models](#). *CoRR*, abs/2309.06794.
- Yakir Yehuda, Itzik Malkiel, Oren Barkan, Jonathan Weill, Royi Ronen, and Noam Koenigstein. 2024. [In search of truth: An interrogation approach to hallucination detection](#). *CoRR*, abs/2403.02889.
- Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2023. [Agenttuning: Enabling generalized agent abilities for llms](#). *CoRR*, abs/2310.12823.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren’s song in the AI ocean: A survey on hallucination in large language models](#). *CoRR*, abs/2309.01219.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.
- Wei Zhao, Mingyue Shang, Yang Liu, Liang Wang, and Jingming Liu. 2020. [Ape210k: A large-scale and template-rich dataset of math word problems](#). *arXiv preprint arXiv:2009.11506*.

## A HaluAgent Framework

### A.1 Instructions

During the trajectory generation phase, we provide detailed instructions for the HaluAgent framework to prompt GPT-4 to generate the required detection trajectories. The English and Chinese prompts are shown in Figure 4 and Figure 5, respectively.

### A.2 Toolbox

HaluAgent includes both verification tools and system tools. The tool names and their usage instructions are summarized in Table 5.

### A.3 Data Source

We introduce the data sources for trajectory generation here, which include five datasets: HaluEval, WebQA, Ape210K, HumanEval, and WordCnt.

- **HaluEval** (Li et al., 2023) is a benchmark for evaluating hallucinations in LLMs across various tasks. We select 1,000 samples from the QA subset, with 900 samples used for trajectory generation and 100 samples for testing.

- **WebQA** (Li et al., 2016) is a QA dataset collected from the Baidu Zhidao platform. This dataset contains a large number of questions and corresponding answers. We sample 1,000 questions from the training set of WebQA and generate answers for each question using ChatGPT. Then, human annotators evaluate these answers by comparing them to the correct ones to determine if they contain hallucinations. We use 900 questions for generating detection trajectories and reserve 100 samples as the testset.

- **Ape210K** (Zhao et al., 2020) is a large-scale math word problem dataset, containing 210,000 Chinese primary school-level math problems. Each problem includes a gold answer and the equations needed to derive the answer. Solving Ape210K requires not only natural language understanding but also commonsense knowledge. We select 700 data samples from the training set of Ape210K for mathematical calculation. As with the WebQA processing method, we first use ChatGPT to generate answers for each question, and then human annotators label the responses. We reserve 100 data samples as the testset, and use the remaining 600

Category	Tool Name	Tool Usage Instruction
Verification Tools	web_search	<i>Input:</i> sentence: str $\rightarrow$ <i>Output:</i> fact Conduct a web search and return factual information.
	calculator	<i>Input:</i> formula: str $\rightarrow$ <i>Output:</i> result Perform calculations based on the input formula and return the result.
	code_interpreter	<i>Input:</i> code: str $\rightarrow$ <i>Output:</i> label Execute code and return a label indicating whether the execution was successful.
	word_count	<i>Input:</i> length: int, text: str $\rightarrow$ <i>Output:</i> count, label Count words in the text and provide a label indicating whether the requirement is met.
	match	<i>Input:</i> sentence: str, context: str $\rightarrow$ <i>Output:</i> label Match a sentence with the given context and return a label indicating whether semantic matching is successful.
System Tools	split_text	<i>Input:</i> text: str $\rightarrow$ <i>Output:</i> sentences Split text into individual sentences.
	get_answer	<i>Input:</i> $\rightarrow$ <i>Output:</i> result(, evidence) Return the detection answer with optional evidence.

Table 5: Instructions of the toolbox in HaluAgent.

samples for trajectory generation.

- **HumanEval** (Chen et al., 2021) dataset contains 164 programming problems, including function names, comments, specific implementations, and multiple unit tests. It is widely used to test the code generation capabilities of LLMs. We use CodeLlama-34b-Instruct-hf (Rozière et al., 2023) to generate code for the HumanEval dataset, using 100 samples for trajectory generation and the rest for testing.

- **WordCnt** is a newly constructed dataset, representing conditional text generation tasks in scenarios when users interact with LLMs. Specifically, we create WordCnt by prompting GPT-4 to generate a set of text generation instructions with specific length requirements and the corresponding responses for each instruction. WordCnt consists of 200 samples, with 100 samples used for trajectory generation and 100 samples for testing.

## B Experiment Setup

### B.1 Out-of-domain Datasets

We present the details of the out-of-domain datasets here.

- **HalluQA** is a Chinese Hallucination Question-Answering benchmark covering misleading questions like identity awareness and knowledge-based questions. Each question includes one correct answer and several answers with hallucinations. We conduct experiments on 206 knowledge-related samples from the dataset, randomly selecting one answer from the correct answer and hallucinated answers for evaluation.

- **HaluEval 2.0** is a hallucination evaluation benchmark that contains large-scale questions from five domains: biomedicine, finance, science, education, and open domain. We evenly sample 100 of HaluEval 2.0 for hallucination detection evaluation.

### B.2 Baselines

In the experiment section, we compare the hallucination detection performance of GPT-4, Baichuan2-Chat, and HaluAgent. The detailed baseline settings are explained below.

- **GPT-4 prompt** involves providing GPT-4 with a simple description of the hallucination detection task, enabling the model to determine whether there are hallucinations in the text. The model’s response is either “Yes” or “No”, indicating the presence or absence of hallucinations. The detailed prompt is shown in Figure 6.

- **GPT-4 pipeline** guides GPT-4 through the hallucination detection process following the HaluAgent framework, which includes steps such as sentence segmentation and tool invocation. In addition to providing a “Yes” or “No” answer, it identifies the location of the hallucinations and provides supporting evidence. The instructions are shown in Figure 4 and Figure 5.

- **Baichuan2-Chat (7B and 13B)** do not have the capability to follow the HaluAgent detection framework. Therefore, we evaluate these models using the same simple hallucination detection prompt as used for GPT-4.

- **HaluAgent (7B and 13B)** are models fine-tuned with trajectory data. We evaluate them using HaluAgent instructions in a zero-shot setting, similar to the evaluation of the GPT-4 pipeline.

### B.3 Implementation Details of Scalability Study

#### B.3.1 Constructed Dataset

We construct a new dataset for translation and data calculation tasks via ChatGPT by providing examples. Our goal is to create a dataset specifically for evaluating these new tools, so the questions in the dataset are straightforward. We present some examples below:

- *Translate the following Spanish into Chinese: ¡Hola! ¿Cómo estás?*
- *How many days are there from 2014-02-06 to 2014-05-21?*

#### B.3.2 Instructions of New Tools

To guide HaluAgent in using the new tools, we include descriptions and usage examples of these tools in the instructions. The detailed prompt is shown in Figure 7.

---

You are an agent tasked with detecting hallucinations in reply texts using a specific framework. Below is a detailed explanation of the detection framework:

Firstly, you need to determine whether to split the input reply text into a list of sentences using a sentence segmentation tool. If required, you should check each sentence individually; otherwise, the entire text should be checked as a whole. You can choose an appropriate fact-checking tool to obtain relevant information and knowledge for verification, and then use the matching tool to output the judgment results, or directly output the judgment results. If you do not use the match tool and directly output the judgment results, you need to output the label in your thought. If there is an error, output "label = 1"; if there is no error, output "label = 0". After the verification is completed, you need to reflect on all detection results and output the label in your thought, then call `get_answer()` to produce the final detection result.

#### **Sentence Segmentation Tool:**

**`split_text(text: str) -> sentence_list`**

This function splits the text into a list of sentences.

#### **Fact-Checking Tools:**

**`web_search(sentence: str) -> fact`**

This function uses a search engine to find information related to the sentence.

**`calculator(sentence: str, formula: str) -> result, label`**

This function uses a calculator to obtain the result of a formula and checks if the result matches the sentence. If they match, the label is 0; otherwise, it is 1. Valid operators include +, -, \*, /, and parentheses. For instance, a valid input could be `“(1 + 2) * 3”`. If the input is an equation, it needs to be converted to a formula without unknowns.

**`word_count(length: int, text: str) -> count, label`**

This function calculates the word count of a text and outputs the count. If the word count does not meet the specified length, the label is 1; otherwise, it is 0.

**`code_interpreter() -> label`**

This function checks whether the code can execute correctly. If it executes correctly, the output label is 0; otherwise, it is 1.

#### **Matching Tool:**

**`match(sentence: str, context:str) -> label`**

This function checks a sentence against its context, which might include content from questions and replies around the detected sentence. It looks for irrelevant or contradictory answers. If any are found, the label is 1; otherwise, it is 0. If you think the output of match is wrong, you can correct the label in thought. For example, if you think the "label = 0" output by match is wrong, you can correct the answer and output "label = 1" in thought.

Every time it's your turn to respond, you must strictly follow this format to present your thoughts and actions: "THOUGHT: Your thought process. ACTION: Tool call, *e.g.*, `match(sentence="...", context="...")`". After each tool invocation, I will provide the output as follows: "OBSERVATION: Tool output".

---

Figure 4: Instructions of HaluAgent framework in English.

---

你是一个通过特定的框架检测回复文本中的幻象的智能体。下面是检测框架的详细说明。

首先，你需要判断是否要将输入中的回复文本拆分为句子列表。你可以使用拆分句子的工具。如果需要拆分，需要对每个句子逐一进行核查；否则就对整个回复文本进行核查。你可以选择适当的事实核查工具来获取用于核查的相关信息和知识然后使用匹配工具输出判断结果或者直接输出判断结果。如果不使用 `match` 工具而直接输出判断结果，则需要思考中输出 `label`。存在错误输出 `"label = 1"`；不存在错误输出 `"label = 0"`。核查完毕后，你需要在思考中反思所有检测结果并输出 `label`，在行为中调用 `get_answer()` 输出最终的检测结果，如果存在幻象一并输出幻象内容和证据。

#### 分句工具：

**`split_text(text: str) -> sentence_list`**

输入是文本，该函数将文本分割成句子列表。

#### 事实核查工具：

**`web_search(sentence: str) -> fact`**

输入是一个句子，该函数使用搜索引擎来搜索相关信息。调用 `web_search` 后必须接着调用 `match` 工具来判断回复与检索到的信息是否匹配。

**`date(date1, date2) -> days`**

该函数接收两个日期作为输入，输出两个日期之间的天数差。

**`calculator(sentence: str, formula: str) -> result, label`**

输入是需要检查的公式，此函数使用计算器来获取计算结果并判断得到的结果是否与句子匹配。如果匹配 `label` 为 0，否则为 1。有效的运算符有 `+`、`-`、`*`、`/` 和 `(,)`。例如，合法的输入可以是 `"(1 + 2) * 3"`。如果输入为方程，需要将其转换为不含未知数的算式。

**`word_count(length: int, text: str) -> count, label`**

输入文本的指定字数和一段文本。该函数计算这段文本的字数并输出为 `count`。如果字数不符合要求，输出 `label` 为 1，否则为 0。

**`code_interpreter() -> label`**

该函数检查代码是否能够正确执行。如果能正确执行，输出标签为 0，否则为 1。

**`translate(text, target_language) -> translated_text`**

该函数输入一段文本和目标语言，输出翻译后的文本。

#### 匹配工具：

**`match(sentence: str, context:str) -> label`**

输入是一个句子以及相应的上下文。上下文可以是问题和回复中的检测句子之前的内容。该函数检查句子中是否存在答非所问或自相矛盾的情况。如果有，则输出标签为 1，否则为 0。如果你认为 `match` 的输出是错误的，可以在思考中修正 `label`，例如如果你认为 `match` 输出的 `"label = 0"` 是错误的，可以在思考中输出 `"label = 1"`。

每次轮到你回复时，你必须严格遵循以下格式给出你的思考和行为：“思考：你的思考过程。行为：工具调用。如 `match(sentence="...", context="...")`”，其中思考部分是你的规划内容，行为部分必须为一个工具调用指令。每次你调用工具后，我会以这种格式为你提供结果：“观察：工具的输出结果”。

---

---

You are a hallucination detection agent. Given a question and its corresponding response, please determine whether there is any incorrect or unsatisfactory content within the response. If there is, output "Yes"; if there is not, output "No".

---

你是一个检测回复文本中的幻象的智能体。给定一个问题和对应的回复，请你判断这段回复中是否有不正确或不符合要求的内容。如果有，请输出“是”；如果没有，输出“否”。

---

Figure 6: Simple description of the hallucination detection task in English and Chinese.

---

`translate(text, target_language) -> translated_text`

该函数输入一段文本和目标语言，输出翻译后的文本。

问题：将以下西班牙语翻译成中文：¿Dónde está la tienda más cercana? 回复：最近的商店在哪里？

思考：这个问题涉及将西班牙语翻译成中文，我需要调用 `translate` 工具进行检查。

行为：`translate(text="¿Dónde está la tienda más cercana?", target_language="zh-cn")`

观察：最近的商店在哪里？

思考：回复中的“最近的商店在哪里？”与翻译器得到的“最近的商店在哪里？”含义相匹配，`label = 0`。检测结束。调用 `get_answer()` 来输出最终的检测结果。

行为：`get_answer()`

观察：否。

---

`date(date1, date2) -> days`

该函数接收两个日期作为输入，输出两个日期之间的天数差。

问题：2024-01-01 到 2024-06-05 共有多少天？ 回复：2024-01-01 到 2024-06-05 共有 150 天

思考：这个问题涉及计算两个日期之间的天数，我需要调用 `date` 工具进行检查。

行为：`date(date1="2024-01-01", date2="2024-06-05")`

观察：156 天

思考：回复与计算得到的天数不匹配，`label = 1`。检测结束。调用 `get_answer()` 来输出最终的检测结果。

行为：`get_answer()`

观察：是。"sentence": "2024-01-01 到 2024-06-05 共有 150 天", "fact": "156 天"

---

Figure 7: Description and usage example of new tools.