

# Mitigating Entity-Level Hallucination in Large Language Models

Weihang Su  
swh22@mails.tsinghua.edu.cn  
Department of Computer Science and  
Technology, Tsinghua University,  
Beijing 100084, China

Yichen Tang\*  
Department of Computer Science and  
Technology, Tsinghua University,  
Beijing 100084, China

Qingyao Ai†  
aiqy@tsinghua.edu.cn  
Department of Computer Science and  
Technology, Tsinghua University,  
Beijing 100084, China

Changyue Wang  
Department of Computer Science and  
Technology, Tsinghua University,  
Beijing 100084, China

Zhijing Wu  
School of Computer Science and  
Technology, Beijing Institute of  
Technology, Beijing 100081, China

Yiqun Liu  
Department of Computer Science and  
Technology, Tsinghua University,  
Beijing 100084, China

## ABSTRACT

The emergence of Large Language Models (LLMs) has revolutionized how users access information, shifting from traditional search engines to direct question-and-answer interactions with LLMs. However, the widespread adoption of LLMs has revealed a significant challenge known as hallucination, wherein LLMs generate coherent yet factually inaccurate responses. This hallucination phenomenon has led to users' distrust in information retrieval systems based on LLMs. To tackle this challenge, this paper proposes Dynamic Retrieval Augmentation based on hallucination Detection (DRAD) as a novel method to detect and mitigate hallucinations in LLMs. DRAD improves upon traditional retrieval augmentation by dynamically adapting the retrieval process based on real-time hallucination detection. It features two main components: Real-time Hallucination Detection (RHD) for identifying potential hallucinations without external models, and Self-correction based on External Knowledge (SEK) for correcting these errors using external knowledge. Experiment results show that DRAD demonstrates superior performance in both detecting and mitigating hallucinations in LLMs. All of our code and data are open-sourced at <https://github.com/oneal2000/EntityHallucination>.

## KEYWORDS

Hallucination, Retrieval Augmented Generation, Large Language Model

### ACM Reference Format:

Weihang Su, Yichen Tang, Qingyao Ai, Changyue Wang, Zhijing Wu, and Yiqun Liu. 2024. Mitigating Entity-Level Hallucination in Large Language Models. In *Proceedings of Pre-print*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

\* Contributed equally

† Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*Pre-print, Under Review, V 1.0*

© 2024 Association for Computing Machinery.  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

In recent years, large language models (LLMs) have achieved remarkable success in a variety of natural language processing (NLP) tasks and have already become an indispensable component of many AI applications [2, 5, 32, 43, 50]. Due to the outstanding capabilities and wide applications of LLMs, they have revolutionized how users access information, shifting from traditional search engines to direct question-and-answer interactions with LLMs. However, despite their impressive performance, it is well recognized that existing LLMs may generate text that, while appearing coherent and plausible on the surface, is fundamentally inaccurate or lacks grounding in reality. This phenomenon is commonly referred to as LLM hallucination [13, 23, 28, 41, 51].

To address hallucination, the most popular method adopted in existing studies is Retrieval-Augmented Generation (RAG). In this method, relevant knowledge is retrieved from an external corpus and utilized as input for LLMs, which has been proven to be effective in many NLP tasks [1, 14, 16, 18, 38, 40]. Traditional RAG methods often adopt single-round retrieval that uses the original input of LLM as the query to retrieve relevant external information. Although this method is effective for simple tasks, it frequently fails in complex tasks like long-form generation and multi-hop question-answering. This is primarily because the LLM needs continuously changing information in these complex text generation tasks [15]. In contrast, multi-round RAG [1, 15, 29, 44] performs multiple retrievals during the generation process of LLMs. Depending on the timing of retrieval augmentation, various methods have been proposed in this area. For example, RETRO [1] and IC-RALM [29] trigger the retrieval module based on a pre-defined number of generated tokens. FLARE [15] triggers the retrieval module whenever a token's predictive probability is below a certain threshold. DRAGIN [40] defines an empirical method based on uncertainty and self-attention to determine when to trigger retrieval.

Despite these advancements, a significant oversight remains: none of the existing studies explicitly verify whether the retrieval is triggered at an optimal timing. Typically, the evaluation methods of existing dynamic RAG approaches only focus on the performance in downstream QA tasks, without directly evaluating whether the chosen moment to invoke retrieval is appropriate. In this paper, we argue that retrieval augmentation without identifying when and where hallucination happens in LLM is suboptimal in terms

of effectiveness and efficiency. On the one hand, the retriever cannot be perfect, therefore unnecessary retrieval augmentation may introduce irrelevant or noisy data to LLMs. On the other hand, calling the retrieval module during the LLM generation process will increase the inference time and computational cost. Such cost is unworthy if retrieval augmentation is conducted in places where LLM hallucination doesn't exist.

To address these concerns, we evaluate existing dynamic RAG approaches on the timing of retrieval to find out if it coincides with the occurrence of hallucination. We then propose a dynamic retrieval-augmented generation strategy based on hallucination detection. To be specific, we propose DRAD, Dynamic Retrieval augmentation based on hallucination Detection (**DRAD**, illustrated in Figure 1) that synchronizes retrieval augmentation with real-time hallucination detection of LLMs during the text generation process. DRAD consists of two components: Real-time Hallucination Detection (RHD) and Self-correction based on External Knowledge (SEK). RHD is a real-time hallucination detection method that doesn't rely on any external models or external knowledge. It detects hallucinations by analyzing the uncertainty of the output entities of an LLM, particularly those with low probability and high entropy, that are likely to be potential hallucinations. When RHD determines that the model is likely to generate unreliable text, SEK invokes the retrieval module to retrieve relevant external knowledge and helps the LLM to make corrections to its outputs so that potential hallucinations can be prevented.

We conducted experiments on existing benchmarks [8, 10, 11, 27, 35] to evaluate the effectiveness of our framework. Experimental results show that RHD can achieve state-of-the-art (SOTA) performance in hallucination detection, and DRAD significantly outperforms existing single-round and multiple-round retrieval augmentation methods.

To summarize, the contributions of this paper are as follows:

- We introduce a new retrieval-augmented framework, i.e., DRAD<sup>1</sup>, which detects hallucinations during LLM's inference process and triggers RAG only when hallucinations are detected.
- We propose a real-time hallucination detection method, i.e., RHD, that achieves SOTA performance on existing benchmarks.
- We evaluate DRAD on multiple complex QA benchmark datasets, and the experimental results demonstrate that our proposed DRAD framework significantly reduces hallucinations in large models across three diverse text generation benchmarks, outperforming previous methods.

## 2 RELATED WORKS

### 2.1 Hallucination Detection

Given the significance of hallucination detection and mitigation, considerable research has focused on developing efficient and effective methods for hallucination detection. Some methods require the model to generate multiple outputs for the same input. For instance, SelfCheckGPT [27] introduces three methods: SelfCheckGPT\_BERTScore, SelfCheckGPT\_QA, and SelfCheckGPT\_n-gram.

<sup>1</sup>Our code and data are open-sourced on this anonymous Github link: <https://github.com/oneal2000/EntityHallucination>.

These methods allow an LLM to generate multiple outputs based on the same input. Subsequently, the likelihood of hallucinations occurring in the LLM is measured based on the consistency among these outputs. On the other hand, some methods require the introduction of new models. For example, MIND [41] is an unsupervised hallucination detection method based on the internal states of LLMs. The MIND framework consists of two steps: automatically generating training data and training a hallucination detector based on the hidden states of the selected LLM. The input to the hallucination detector is the hidden states of the LLM, and the output is a 0-1 label, representing whether the LLM is experiencing hallucinations. Liu et al. [23] specifically fine-tuned several models such as BERT [6], RoBERTa [24], and XLNet [47] to detect hallucinations. The input to the Hallucination Detector is the output text of an LLM, and the output is also a 0-1 label indicating whether the LLM is experiencing hallucinations. Our proposed hallucination detection method does not require generating multiple responses for the same input or introducing external information or models. Consequently, our approach excels in efficiency compared to existing works.

### 2.2 Retrieval-augmented Generation

In recent studies, Retrieval-Augmented Generation (RAG) has been widely used to improve the performance of LLMs. One of the most direct methods is single-round RAG [1, 9, 12, 14, 16, 18, 34], which utilize the initial input of the LLM to retrieve external knowledge. The retrieved external knowledge is subsequently integrated into the model's input. Previous studies, like REPLUG [34] and UniWeb [22], have explored using Language Model-based feedback and adaptive search engines for retrieval augmentation in single-round retrieval, enhancing predictions and selectively incorporating external knowledge.

Single-round RAG can be quite effective for straightforward tasks or situations where the user's information needs are well-defined. However, for the tasks that involve generating extensive text such as long-form generation and multi-hop QA, searching for external knowledge based on the initial input cannot sufficiently address the information needs for LLM [15]. As a result, researchers have begun to explore multi-round retrieval augmentation. For example, RETRO [1] and IC-RALM [29] trigger retrieval every 4 to 32 tokens, while IRCot [44] triggers retrieval whenever a new sentence is generated. Looking at it from another perspective, FLARE [15] triggers retrieval when any token in the generated text has a probability lower than a certain threshold. However, indiscriminately considering the probability of every token is not the optimal solution for multi-round retrieval, as many tokens are function words (such as 'am', 'to', 'in', etc) lacking semantic meaning. To address the limitations of FLARE, DRAGIN [40] optimize the timing (when to retrieve) and the query formulation method (what to retrieve) of the dynamic RAG framework, achieving state-of-the-art performance on downstream QA tasks.

## 3 METHODOLOGY

This section will introduce the DRAD framework in detail. DRAD consists of two components: Real-time Hallucination Detection (RHD) and Self-correction based on External Knowledge (SEK). We will introduce RHD in section 3.1 and SEK in section 3.2.

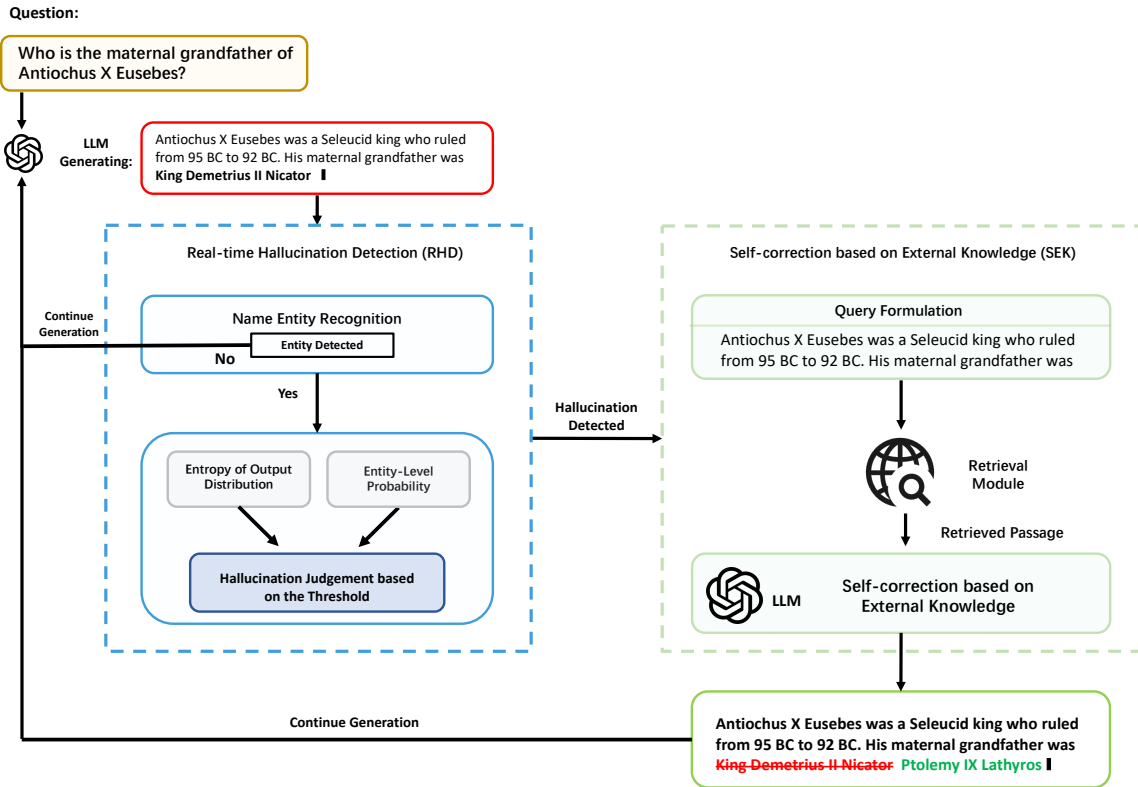


Figure 1: An illustration of our proposed DRAD Framework, which comprises two main components: RHD and SEK Module, highlighted in the diagram with blue and green frames, respectively.

### 3.1 Real-time Hallucination Detection

**3.1.1 Empirical Investigation.** Our approach referred to as Real-time Hallucination Detection (RHD, illustrated in Figure 2), is constructed based on the assumption that when an LLM is forced to provide answers or generate text beyond its existing knowledge, the uncertainty in its output significantly escalates [27]. For example, consider the input, "Bill Clinton was born and raised in \_\_\_\_". Since the LLM has seen extensive information about Bill Clinton during the pre-training phase, it can confidently assign a high probability to the token "Arkansas" to complete the sentence. Conversely, other locations such as "Alabama" or "Paris" would be considered low probability. However, in scenarios where the LLM is required to generate text about unfamiliar topics, such as "Alice's childhood neighbor now lives in \_\_\_\_," the model faces a challenge. It lacks specific knowledge to confidently choose a word to fill the gap, leading to a flat probability distribution across all tokens related to place names in its vocabulary. This uncertainty in choosing the right token can cause the model to select a random entity during text generation and cause hallucinations.

This insight leads us to establish a connection between uncertainty metrics and hallucination detection in text generation. When a model's internal parameters cannot reliably choose the correct response from multiple options, its outputs tend to include tokens with low probabilities and high entropy. Therefore, we propose


a method for real-time detection of hallucinations in LLM's outputs, utilizing an entity confidence metric derived from both the predictive probability and entropy of entities.

**3.1.2 Entity Probability.** The first method to evaluate the hallucination of an LLM involves examining its generation probability for each entity in its output. This process starts with the recognition of entities in the LLM output. Specifically, let  $O$  represent the output sequence generated by an LLM. We apply real-time entity recognition on  $O$  to identify entities. Assume an entity  $E$  is detected in  $O$ , consisting of  $n$  tokens. Next, for each token  $t_i$  in  $E$ , where  $1 \leq i \leq n$ , let  $p_i$  denote the generation probability of  $t_i$  as given by the output layer of the LLM. Finally, the overall probability of the entity  $E$  is computed as an aggregation of the probabilities of its constituent tokens. This can be expressed as:

$$P(E) = f(p_1, p_2, \dots, p_n), \quad (1)$$

where  $f$  is an aggregation function (e.g., product, sum, average, min, max) applied to the probabilities  $p_1, p_2, \dots, p_n$  of the tokens in  $E$ .

**3.1.3 Entity Entropy.** In addition to probability, *entropy* is widely used to evaluate the uncertainty inherent in the model's output distribution. Given an entity  $E$  comprised of  $n$  tokens, we assess the entropy for each token in the LLM output distribution. For each token  $t_i$  in the entity  $E$ , the entropy  $\mathcal{H}_i$  in the LLM output



**Figure 2: An illustration of our proposed Real-time Hallucination Detection (RHD) framework.**

distribution is defined as:

$$\mathcal{H}_i = - \sum_{\tilde{w} \in \mathcal{W}} p_i(\tilde{w}) \log p_i(\tilde{w}), \quad (2)$$

where  $p_i(\tilde{w})$  represents the likelihood of generating word  $\tilde{w}$ , and  $\mathcal{W}$  is the vocabulary of the LLM. To compute the overall entropy of the entity  $E$ , we aggregate the entropies of its constituent tokens. This can be done using a pooling function  $f$  that takes the entropies of individual tokens and combines them to yield the entropy of the entity:

$$\mathcal{H}(E) = f(\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_n). \quad (3)$$

This approach, using entropy as a measure, offers an additional dimension to evaluate the uncertainty in an LLM’s output, particularly in its ability to generate specific entities.

**3.1.4 Threshold for Hallucination Detection.** To explicitly determine whether an entity generation should be considered as a hallucination or not, we define two thresholds,  $\theta_1$  and  $\theta_2$ .

$$\text{Entity Hallucination} = \begin{cases} \text{yes} & \text{if } P(E) < \theta_1 \text{ or } \mathcal{H}(E) > \theta_2 \\ \text{no} & \text{otherwise} \end{cases} \quad (4)$$

Note that  $\theta$  determines the frequency of calling the retrieval module. The specific value of  $\theta$  can be adjusted based on the demands of the actual application, thereby offering flexibility in determining how often the retrieval module is consulted. We conducted a detailed experiment to explore the impact of the threshold on efficiency and effectiveness. The experimental results can be found in section 5.2.

## 3.2 Self-correction based on External Knowledge

The objective of the DRAD framework is to mitigate hallucinations in LLM using a Detect-Retrieve-Revise paradigm. The RHD module performs detection, which involves real-time identification of hallucinations. Following hallucination detection, DRAD retrieves relevant external knowledge based on the context surrounding the hallucination’s occurrence and then revises the hallucination. This Retrieve-Revise process is implemented through the SEK method (Self-correction based on External Knowledge). The SEK method comprises three main components: query formulation, relevant knowledge retrieval, and self-correction. This section provides a detailed description of the SEK method.

**3.2.1 Formulating Search Queries.** Given the output of an LLM  $O$  comprising tokens  $t_1, t_2, \dots, t_n$  and an entity  $E$  identified as a hallucination by the RHD module, where the token span of  $E$  is from  $t_k$  to  $t_{k+i}$ , the SEK query  $Q$  is formulated as:

$$Q = f(O, t_k, i) = \text{concat}(t_{k-m}, \dots, t_{k-1}, t_{k+i+1}, \dots, t_{k+m}). \quad (5)$$

In this equation,  $Q$  represents the formulated query,  $O$  denotes the output of the LLM consisting of a series of tokens,  $t_k$  to  $t_{k+i}$  define the hallucination entity span within the context, and  $\text{concat}$  is the concatenation function that assembles tokens surrounding the entity span to construct the query. The function  $f$  represents the process of query generation based on the given context and entity span. The parameter  $m$  is a flexible hyperparameter that determines the range of tokens included in the concatenation process.

**3.2.2 External Knowledge Retrieval.** Given an external knowledge corpus  $\mathcal{D}$ , composed of a set of documents  $\{d_1, d_2, \dots, d_n\}$ , and a retrieval function  $\mathcal{R}$ , the process of external knowledge retrieval for a given query  $q$  is mathematically formalized as follows:

$$\mathcal{S} = \{\mathcal{R}(q, d_i) | d_i \in \mathcal{D}\}. \quad (6)$$

In this equation,  $\mathcal{S}$  represents a score list where each score is computed by the retrieval function  $\mathcal{R}$  for the query-document pair  $(q, d_i)$ , assessing their relevance.

$$\mathcal{S}_{\text{sorted}} = \text{sort}_{\text{desc}}(\mathcal{S}), \quad (7)$$

where  $\mathcal{S}_{\text{sorted}}$  is the list of scores  $\mathcal{S}$  sorted in descending order based on relevance. Finally, the top  $k$  documents from  $\mathcal{S}_{\text{sorted}}$  are selected as the set of retrieved relevant documents:

$$\mathcal{D}_{\text{retrieved}} = \{d_i | (q, d_i, s_i) \in \mathcal{S}_{\text{sorted}} \wedge i \leq k\}. \quad (8)$$

In this equation,  $\mathcal{D}_{\text{retrieved}}$  denotes the subset of documents from  $\mathcal{D}$  that are most relevant to the query  $q$ , based on the top  $k$  scores in  $\mathcal{S}_{\text{sorted}}$ .

Note that in the field of information retrieval, various methods have been explored for the retrieval function  $R$ . These include vocabulary-based approaches such as TF-IDF [30], BM25 [31], and Query Likelihood [49], as well as neural network-based methods [3, 4, 7, 19–21, 25, 26, 36, 37, 39, 48].

**3.2.3 Self-Correction.** The RHD method detects the specific position in an LLM’s output where hallucination occurs. It is important to note that the text preceding this hallucination position is typically accurate and does not require modification. Thus, to revise the hallucination tokens, the first step of Self-correction is truncating the output of an LLM at the position of hallucination. This truncation process is mathematically represented as:

$$O' = \text{truncate}(O, t_k), \quad (9)$$

where  $O'$  denotes the truncated output,  $O$  is the original output of the LLM, and  $t_k$  denotes the first token of the hallucinated entity. Following this, the retrieved external knowledge set  $\mathcal{D}_{\text{retrieved}}$  (defined in Eq 8) is input into the LLM using a specific prompt template<sup>2</sup>. After the external knowledge is concatenated, the LLM regenerates the content at the position of the hallucination, utilizing the external knowledge. The regenerated content represents the LLM’s self-correction based on external knowledge, aiming to address the detected hallucinated entities in its output.

### 3.3 Discussion

In this subsection, we discuss the limitations of the RHD method. While the RHD method is effective at detecting hallucinations caused by gaps in the knowledge of LLMs, it is less successful at identifying hallucinations that are due to incorrect information acquired during the model’s pre-training phase. Hallucinations in LLMs primarily arise from two sources: first, the absence of relevant knowledge, which leads the LLM to generate the most probable but potentially incorrect token; second, the learning of incorrect knowledge during pre-training. The RHD method utilizes uncertainty for hallucination detection, making it proficient at recognizing the first type of hallucination, where LLMs often show uncertainty. However, it struggles with the second type, where LLMs might deliver incorrect responses with high confidence.

## 4 EXPERIMENTAL SETUP

### 4.1 Datasets

**4.1.1 Hallucination Detection.** We evaluate our proposed real-time hallucination detection method on WikiBio GPT-3 dataset [27]. This dataset comprises text generated by GPT-3 based on various topics, along with manually annotated labels indicating whether these passages contain hallucinations from a factual perspective.

**4.1.2 Retrieval-augmented Text Generation.** We evaluate the text generation performance of DRAD on the following datasets.

- **2WikiMultihopQA** [11]. We use the 2WikiMultihopQA dataset to assess the DRAD model’s ability to answer complex questions. 2WikiMultihopQA comprises complex questions that require two hops of information from Wikipedia articles. Answering these questions typically involves composing, comparing, or inferring multiple pieces of information to provide an accurate response.
- **StrategyQA** [8]. We utilize the StrategyQA dataset to assess the commonsense reasoning capabilities of DRAD. StrategyQA comprises a diverse range of crowdsourced yes/no questions.

<sup>2</sup>The prompt template is detailed in Section 4.2

- **NQ** [17]. We use the NQ dataset to assess the DRAD model’s capability to answer factual questions. NQ is a large-scale question-answering dataset derived from real queries of Google Search.

### 4.2 Settings for each Dataset

Our settings on various benchmarks are as follows:

- **2WikiMultihopQA** [11]. We follow the prompt template of Wang et al. [45] that instructs the selected LLM to generate the chain-of-thought reasoning process. Following the prompt template of [44] and [15], we add 8 example query-answer pairs to the input as prompts. For the evaluation metrics, we use pattern-matching techniques to extract the final answer from the output of the LLM. This extracted answer is subsequently compared to the reference answer. We employ various methods for this comparison, including the exact match (EM) metric at the answer level and token-level assessments of F1 score, precision, and recall.
- **StrategyQA** [8]. We follow the setting of Wei et al. [46] to instruct the LLM to generate the chain-of-thought reasoning process. Following the prompt template of [46] and [15], we add 8 example query-answer pairs to the input as prompts. Since this dataset only contains true/false questions, we directly reported the accuracy.
- **NQ** [17]. Since the questions in NQ are relatively simple, we did not include example query-answer pairs as prompts in the input to LLM. For the evaluation metrics, we use the same parameters as we did in our 2WikiMultihopQA experiment.

For all three datasets, we chose BM25 as our retrieval model based on findings from [29], which demonstrated its superior performance in RAG, even outperforming SOTA dense retrieval models. For the external corpus, we use Wikipedia passages. The Top 3 relevant passages retrieved by BM25 were added to the prompt.

### 4.3 Baselines

**4.3.1 Hallucination Detection.** We chose the following hallucination detection methods as baselines:

- **SCG-BERTScore** [27] uses the LLM to produce multiple samples for a single input. BERTScore calculates the average similarity between a sentence and its most similar counterpart in various samples. If a sentence is highly similar to many samples, it’s probably factual. If not, it might be a fabricated detail or "hallucination".
- **SCG-QA** [27] generates multiple samples from the LLM for a single input. Based on one of these samples, a query generation system creates multiple-choice questions. Then a QA system uses the other samples to answer the question. The consistency of these answers is used to measure the probability of hallucination.
- **SCG-Ngram** [27] involves training an n-gram model, on samples generated by the selected LLM. As the sample size escalates, the behavior of this new model progressively approximates that of the original LLM. Subsequently, the average log probabilities for a given response,  $R$ , are computed

**Table 1: A comparative overview of our selected Retrieval-Augmented Generation baselines.**

	Timing for Retrieval	Query Formulation
<b>SRR</b>	Before Generation	Initial Input
<b>FLR</b>	Per Sentence	Last Generated Sentence
<b>TPR</b>	Any Token Probability Below Threshold	Last Generated Sentence Exclude Uncertain Tokens
<b>DRAD</b>	Entity-based Hallucination Detection	Context of Hallucination

based on the n-gram model. This average of log probabilities is then employed as a metric to quantify the likelihood of hallucination.

- **SCG-Ensemble** [27] is a simple combination of the normalized scores of SCG-BERTScore, SCG-QA, and SCG-Ngram.
- **Predictive Probability** [15] involves leveraging the probabilities of tokens generated by LLMs as a metric to measure hallucination. FLARE [15] adopts this simple methodology to determine the appropriate timing for invoking retrieval augmentation.
- **Predictive Entropy** utilizes the entropy of the output distribution of tokens generated in the text by large models as a metric to gauge hallucination.

**4.3.2 Text Generation.** We choose the following Text Generation baselines for comparison. To ensure a fair comparison, each sentence in every multi-round RAG method can only be revised once.

- **NOR (NO Retrieval).** Directly generating answers without retrieval augmentation.
- **SRR (Single-Round Retrieval).** In this setting, we retrieve relevant external documents once based on the initial question. Most of the current retrieval-augmented LLMs adopt this setting.
- **FLR (Fix Length Retrieval)** [44]. A multi-round retrieval augmentation method that triggers the retrieval module for every sentence.
- **TPR (Token Probability Retrieval)** [15]. A multi-round retrieval augmentation method that triggers retrieval each time it encounters an uncertain token.

For multi-round Retrieval-Augmented Generation, the two most critical aspects are the timing for retrieval and the method of query formation when triggering retrieval. To visually demonstrate the differences between our selected baselines, we present the timing for retrieval and the method of query formation for each baseline in Table 1.

## 4.4 Implementation Details

In this subsection, we provide a comprehensive overview of our implementation details for the major components of our study: configurations for LLMs, Named Entity Recognition (NER), and the External Knowledge Corpus.

- **LLM Configuration:** We follow settings of Jiang et al. [15] and validate our retrieval augmentation approach based on

**Table 2: Experimental results of our hallucination detection technique compared to other baseline methods on the SeflCheckGPT dataset. The best performances are highlighted in bold. The performance of each SeflCheckGPT method is directly sourced from their original paper.**

		AUC
<b>Multi Generation</b>	SCG_BERTScore	81.96
	SCG_QA	84.26
	SCG_n-gram	85.63
	SCG_ensemble	87.33
<b>Single Generation</b>	Avg Token Entropy	80.73
	Avg Token Prob	83.21
	Max Token Entropy	85.75
	Min Token Prob	87.51
	<b>RHD(ours)</b>	<b>89.31</b>

**Table 3: Comparative experimental results for various token pooling methods of RHD. Since low probability and high entropy indicate low confidence, we employ min pooling for probability and max pooling for entropy.**

AUC		
Pooling Method	Probability	Entropy
<b>Max</b>	-	<b>88.00</b>
<b>Min</b>	88.75	-
<b>First</b>	87.60	87.46
<b>Average</b>	<b>89.31</b>	87.91

the GPT-3.5 language model, specifically the "text-davinci-003" variant, by iteratively querying its API <sup>3</sup>. Notably, this API allows access to the probabilities of the generated tokens.

- **Name Entity Recognition (NER):** For the Named Entity Recognition (NER) component of RHD, we follow the methodologies in prior studies [23, 42]. Specifically, we utilized the Spacy library, a tool recognized for its efficacy in NER as evidenced by previous research [33].
- **External Knowledge Corpus:** We adopt Wikipedia as our external knowledge corpus. Each article is segmented into 100-token passages.

## 5 EXPERIMENTAL RESULTS

### 5.1 Hallucination Detection

In this section, we present the experimental findings obtained from the WikiBio GPT-3 dataset. This dataset has been specifically crafted for a binary classification task, with the objective of measuring the likelihood of hallucinated text produced by the GPT3 (text-davinci-003) model. The AUC (Area Under the Curve) metric was adopted to evaluate the performance of each hallucination detection method.

<sup>3</sup><https://api.openai.com/v1/chat/completions>

**Table 4: Experimental results of DRAD and other baselines on 2WikiMultihopQA, NQ, and StrategyQA. The best results are in bold. #Num indicates the times of retrieval module calls, smaller means more efficient.**

		2WikiMultihopQA					NQ		StrategyQA	
		F1	EM	Prec.	Recall	#Num	F1	#Num	Accuracy	#Num
<b>Without Retrieval</b>	<b>NOR</b>	0.2939	0.22	0.2879	0.3001	0	0.283	0	0.64	0
<b>Single-round Retrieval</b>	<b>SRR</b>	0.3768	0.26	0.3622	0.3927	<b>1</b>	0.293	1	0.65	1
	<b>FLR</b>	0.4788	0.38	0.4693	0.4886	9.61	0.291	3.41	0.62	5.48
<b>Multi-round Retrieval</b>	<b>TPR</b>	0.3866	0.24	0.3681	0.4071	5.29	0.252	1.67	0.60	2.25
	<b>DRAD</b>	<b>0.4856</b>	<b>0.39</b>	<b>0.4741</b>	<b>0.4976</b>	1.40	<b>0.339</b>	<b>0.48</b>	<b>0.76</b>	<b>0.53</b>

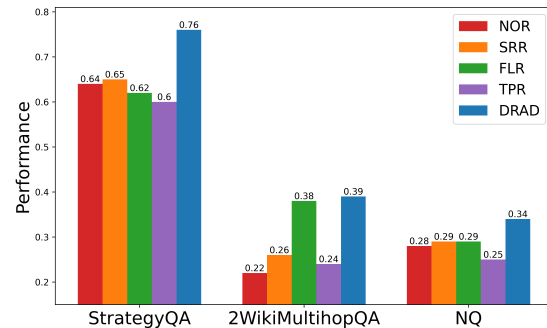
**5.1.1 Overall Results.** The experiment results of RHD and other baselines are shown in Table 2. We have the following observations. (1) Firstly, methods based on token probability are shown effective in detecting hallucinations. Among these, the minimum token-level probability has superior performance among all token-generation probability techniques. However, the performance of token-level hallucination detection methods generally falls short when compared to RHD. **This also confirms our assumption that indiscriminately considering the probability of every token to detect hallucinations is unreasonable, as many tokens are meaningless words.** (2) Furthermore, the SCG\_ensemble method, while effective among multi-round generation techniques, has drawbacks due to the need for LLMs to create multiple responses, making it impractical for real-time use. Despite its high computational complexity, its performance advantage over probability-based methods is not substantial. Thus, we recommend using the token probability methods when possible, due to their efficiency and effectiveness. (3) Lastly, our proposed hallucination detection method RHD outperforms all other baselines. This method enables real-time detection of hallucinations and exhibits excellent performance, providing a solid foundation for our proposed Detect-Retrieve-Revise paradigm.

**5.1.2 Ablation Study on Pooling Methods.** In the output generated by LLM, numerous entities are composed of multiple tokens. Thus, we explore diverse pooling strategies for these entities that encompass multiple tokens which are shown in Table 3. Through the experimental results, we have the following observations: Firstly, different pooling methods have minimal impact on our RHD method, suggesting that RHD is not sensitive to pooling strategies. Additionally, the average pooling of Probability achieves the highest AUC with a score of 89.31, indicating it's the most effective method for representing an entity. The following experiments in this section all follow this setting.

## 5.2 Downstream Question-Answering Tasks

**5.2.1 Overall Performance.** The overall experimental results of DRAD and other baselines<sup>4</sup> are presented in Table 4 and Figure 3. The experimental results show that DRAD outperforms all baseline

<sup>4</sup>The implementation method of the TPR (FLARE) official Github repository is as follows: as long as the probability of all the tokens generated by LLM is not higher than the threshold, retrieval is continuously triggered to modify this sentence until the condition is met. To ensure a fair comparison, we stipulate that each sentence in every multi-round RAG method can only be revised once. That's why the performance reported in the original FLARE paper is higher than we reproduce.



**Figure 3: Visual comparison between DRAD and baselines across all datasets. For the evaluation metric, we choose Accuracy for StrategyQA, EM for 2WikiMultihopQA, and F1 for NQ.**

methods across various benchmarks. Specifically, in the context of StrategyQA and NQ datasets, DRAD demonstrates a remarkable improvement in performance relative to preceding retrieval-augmented models. For the 2WikiMultihopQA dataset, while the results are commensurate with those of FLR, it is noteworthy that the number of retrieval invocations for DRAD constitutes only 14.7% of those necessitated by FLR. For the FLR method, though the retrieval module is invoked for every individual sentence, our experimental results suggest that such an approach does not invariably enhance the performance of Language Learning Models (LLMs). Particularly, when tested on the StrategyQA dataset, the performance of the FLR method was found to be subpar in comparison to the single-round retrieval methods. Additionally, despite utilizing the identical hyperparameters as delineated in the original paper on TPR, the performance of the TPR model was observed to be inferior to other baseline models. On another note, on datasets such as NQ and StrategyQA, the DRAD model necessitated even fewer invocations of the retrieval module than methods supplemented by single-round retrieval. This observation emphatically underscores the superior efficiency and effectiveness of our proposed methodology.

**5.2.2 Ablation Studies on Threshold.** This section presents the results of our ablation study on different thresholds. The comparison



**Table 5: Comparison between different thresholds of RHD on 2WikiMultihopQA and StrategyQA. The best results are in bold.**

Threshold	2WikiMultihopQA		StrategyQA
	F1	EM	Accuracy
<b>0.40</b>	<b>0.4732</b>	<b>0.39</b>	<b>0.76</b>
<b>0.42</b>	0.4708	<b>0.39</b>	0.75
<b>0.44</b>	0.4682	<b>0.39</b>	0.75
<b>0.46</b>	0.4715	0.38	<b>0.76</b>
<b>0.48</b>	0.4689	0.38	0.75
<b>0.50</b>	0.4654	0.38	0.75

**Table 6: Comparison between the efficiency of DRAD and other baselines across all datasets. The most efficient results are in bold. 2WMQA indicates the 2WikiMultihopQA dataset.**

Dataset	#Num of retrieval			
	FLR	TPR	DRAD	SRR
NQ	3.41	1.67	<b>0.48</b>	1
<b>2WMQA</b>	9.61	5.29	1.40	<b>1</b>
<b>StrategyQA</b>	5.48	2.25	<b>0.53</b>	1

between different thresholds of RHD ( $\theta_1$ ) on 2WikiMultihopQA and StrategyQA is shown in Table 5. The experimental results demonstrate that our method is not sensitive to hyperparameters. There is no significant performance difference in the range of 0.4 to 0.5.

It’s important to note that the threshold directly determines the frequency of invoking the retrieval module. A higher threshold results in fewer calls to the retrieval. In practical applications, the threshold can be adjusted based on real-world requirements to balance efficiency and effectiveness.

**5.2.3 Efficiency of DRAD.** This section presents the experimental results regarding the efficiency of various multi-round retrieval-augmented LLMs as shown in Table 6. We compared the number of retrieval calls made by SRR, FLR, TPR, and DRAD across different datasets. The results indicate that FLR, due to its requirement of invoking retrieval for every sentence, is the least efficient. For TPR’s efficiency lies between that of DRAD and FLR. DRAD, however, demonstrates the fewest retrieval calls, making it the most efficient.

It is noteworthy that DRAD makes fewer than one retrieval call on the NQ and StrategyQA datasets, yet its performance surpasses that of Single-time retrieval. This underscores the superiority of the DRAD approach over other retrieval-augmented models in both efficiency and effectiveness.

## 6 CONCLUSIONS AND FUTURE WORKS

In this study, we introduce DRAD to mitigate hallucinations of LLMs by integrating real-time hallucination detection (RHD) and self-correction based on external knowledge (SEK). RHD monitors the LLM’s output for potential hallucinations without external models, while SEK retrieves relevant information to adjust the LLM’s output, preventing hallucinations. Experiments demonstrate

DRAD’s superiority in hallucination detection and its effectiveness compared to other retrieval augmentation methods. We acknowledge the limitations of this paper, notably that the real-time hallucination detection method (RHD) depends on token probability data, which is unavailable from certain APIs. Future work aims to develop more real-time hallucination detection methods to overcome this constraint.

## REFERENCES

- [1] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*. PMLR, 2206–2240.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [3] Jia Chen, Haitao Li, Weihang Su, Qingyao Ai, and Yiqun Liu. 2023. THUIR at WSDM Cup 2023 Task 1: Unbiased Learning to Rank. *arXiv preprint arXiv:2304.12650* (2023).
- [4] Xuesong Chen, Ziyi Ye, Xiaohui Xie, Yiqun Liu, Xiaorong Gao, Weihang Su, Shuqi Zhu, Yike Sun, Min Zhang, and Shaoping Ma. 2022. Web search via an efficient and effective brain-machine interface. In *Proceedings of the fifteenth ACM international conference on web search and data mining*. 1569–1572.
- [5] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311* (2022).
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Yan Fang, Jingtao Zhan, Qingyao Ai, Jiaxin Mao, Weihang Su, Jia Chen, and Yiqun Liu. 2024. Scaling Laws For Dense Retrieval. *arXiv preprint arXiv:2403.18684* (2024).
- [8] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics* 9 (2021), 346–361.
- [9] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*. PMLR, 3929–3938.
- [10] Hiroaki Hayashi, Prashant Budania, Peng Wang, Chris Ackerson, Raj Neervannan, and Graham Neubig. 2021. WikiAsp: A dataset for multi-domain aspect-based summarization. *Transactions of the Association for Computational Linguistics* 9 (2021), 211–225.
- [11] Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060* (2020).
- [12] Gautier Izacard and Edouard Grave. 2020. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282* (2020).
- [13] 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. *Comput. Surveys* 55, 12 (2023), 1–38.
- [14] Zhengbao Jiang, Luyu Gao, Jun Araki, Haibo Ding, Zhiruo Wang, Jamie Callan, and Graham Neubig. 2022. Retrieval as attention: End-to-end learning of retrieval and reading within a single transformer. *arXiv preprint arXiv:2212.02027* (2022).
- [15] Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. *arXiv preprint arXiv:2305.06983* (2023).
- [16] Urvasi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172* (2019).
- [17] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics* 7 (2019), 453–466.
- [18] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [19] Haitao Li, Jia Chen, Weihang Su, Qingyao Ai, and Yiqun Liu. 2023. Towards better web search performance: pre-training, fine-tuning and learning to rank.



- arXiv preprint arXiv:2303.04710* (2023).
- [20] Haitao Li, Weihang Su, Changyue Wang, Yueyue Wu, Qingyao Ai, and Yiqun Liu. 2023. Thuir@ coliee 2023: Incorporating structural knowledge into pre-trained language models for legal case retrieval. *arXiv preprint arXiv:2305.06812* (2023).
- [21] Haitao Li, Changyue Wang, Weihang Su, Yueyue Wu, Qingyao Ai, and Yiqun Liu. 2023. THUIR@ COLIEE 2023: more parameters and legal knowledge for legal case entailment. *arXiv preprint arXiv:2305.06817* (2023).
- [22] Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jingyuan Wang, Jian-Yun Nie, and Ji-Rong Wen. 2023. The Web Can Be Your Oyster for Improving Large Language Models. *arXiv preprint arXiv:2305.10998* (2023).
- [23] Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2021. A token-level reference-free hallucination detection benchmark for free-form text generation. *arXiv preprint arXiv:2104.08704* (2021).
- [24] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [25] Xinyu Ma, Jiafeng Guo, Ruqing Zhang, Yixing Fan, Xiang Ji, and Xueqi Cheng. 2021. Prop: Pre-training with representative words prediction for ad-hoc retrieval. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 283–291.
- [26] Yixiao Ma, Yueyue Wu, Weihang Su, Qingyao Ai, and Yiqun Liu. 2023. CaseEncoder: A Knowledge-enhanced Pre-trained Model for Legal Case Encoding. *arXiv preprint arXiv:2305.05393* (2023).
- [27] Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896* (2023).
- [28] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661* (2020).
- [29] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *arXiv preprint arXiv:2302.00083* (2023).
- [30] Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, Vol. 242. Citeseer, 29–48.
- [31] Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval* 3, 4 (2009), 333–389.
- [32] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100* (2022).
- [33] Hemlata Shelar, Gagandeep Kaur, Neha Heda, and Poorva Agrawal. 2020. Named entity recognition approaches and their comparison for custom ner model. *Science & Technology Libraries* 39, 3 (2020), 324–337.
- [34] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652* (2023).
- [35] Ivan Stelmakh, Yi Luan, Bhuvan Dhingra, and Ming-Wei Chang. 2022. Asqa: Factoid questions meet long-form answers. *arXiv preprint arXiv:2204.06092* (2022).
- [36] Weihang Su, Qingyao Ai, Xiangsheng Li, Jia Chen, Yiqun Liu, Xiaolong Wu, and Shengluan Hou. 2023. Wikiformer: Pre-training with Structured Information of Wikipedia for Ad-hoc Retrieval. *arXiv preprint arXiv:2312.10661* (2023).
- [37] Weihang Su, Qingyao Ai, Yueyue Wu, Yixiao Ma, Haitao Li, and Yiqun Liu. 2023. Caseformer: Pre-training for Legal Case Retrieval. *arXiv preprint arXiv:2311.00333* (2023).
- [38] Weihang Su, Yiran Hu, Anzhe Xie, Qingyao Ai, Zibing Que, Ning Zheng, Yun Liu, Weixing Shen, and Yiqun Liu. 2024. STARD: A Chinese Statute Retrieval Dataset with Real Queries Issued by Non-professionals. *arXiv preprint arXiv:2406.15313* (2024).
- [39] Weihang Su, Xiangsheng Li, Yiqun Liu, Min Zhang, and Shaoping Ma. 2023. Thuir2 at ntcir-16 session search (ss) task. *arXiv preprint arXiv:2307.00250* (2023).
- [40] Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. 2024. Dragin: Dynamic retrieval augmented generation based on the real-time information needs of large language models. *arXiv preprint arXiv:2403.10081* (2024).
- [41] Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu. 2024. Unsupervised real-time hallucination detection based on the internal states of large language models. *arXiv preprint arXiv:2403.06448* (2024).
- [42] Amogh Kamat Tarcar, Aashis Tiwari, Vineet Naique Dhaimodker, Penjo Rebelo, Rahul Desai, and Dattaraj Rao. 2019. Healthcare NER models using language model pretraining. *arXiv preprint arXiv:1910.11241* (2019).
- [43] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [44] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. *arXiv preprint arXiv:2212.10509* (2022).
- [45] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171* (2022).
- [46] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [47] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* 32 (2019).
- [48] Ziyi Ye, Xiaohui Xie, Qingyao Ai, Yiqun Liu, Zhihong Wang, Weihang Su, and Min Zhang. 2024. Relevance Feedback with Brain Signals. *ACM Transactions on Information Systems* 42, 4 (2024), 1–37.
- [49] ChengXiang Zhai. 2008. Statistical language models for information retrieval. *Synthesis lectures on human language technologies* 1, 1 (2008), 1–141.
- [50] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068* (2022).
- [51] Chungting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. 2020. Detecting hallucinated content in conditional neural sequence generation. *arXiv preprint arXiv:2011.02593* (2020).