




# Pointwise Mutual Information as a Performance Gauge for Retrieval-Augmented Generation

Tianyu Liu \* Jirui Qi \* Paul He 

Arianna Bisazza  Mrinmaya Sachan  Ryan Cotterell 

 ETH Zürich  CLCG, University of Groningen  University of Toronto  
{tianyu.liu, mrinmaya.sachan, ryan.cotterell}@inf.ethz.ch  
{j.qi, a.bisazza}@rug.nl, hepaul@cs.toronto.edu

## Abstract

Recent work suggests that large language models enhanced with retrieval-augmented generation are easily influenced by the order in which the retrieved documents are presented to the model when solving tasks such as question answering (QA). However, there is no method to date that exploits this phenomenon to improve generation. To fill this gap, in this study, we show that the pointwise mutual information between a context and a question is an effective gauge for language model performance. Importantly, this gauge does not depend on knowing the answer to the question *a priori*. Through experiments on two question-answering datasets using a variety of large language models, we find evidence for an empirical correlation between answer accuracy and pointwise mutual information. Additionally, we propose two methods that use the pointwise mutual information between a document and a question as a gauge for selecting and constructing prompts that lead to better performance, whose effectiveness we demonstrate through experimentation.<sup>1</sup>

## 1 Introduction

Prompt design is an important factor when applying language models (LMs) to downstream tasks, including LMs that make use of retrieval-augmented generation (RAG; Lewis et al., 2020). Well-constructed prompts can improve LMs’ answers to user-input questions and help generate responses that better align with user expectations (Gao et al., 2021; Izacard et al., 2024; Liu et al., 2024; Schulhoff et al., 2024; Ma et al., 2024, *inter alia*).

Under the RAG framework, a prompt typically consists of three components. First, an instruction provides a textual description of the overall task and general guidance for the language model. Second, a specific question encodes the precise task or query the model should perform. Third, a context

\*The first two authors contributed equally.

<sup>†</sup>Work performed while at ETH Zürich.

<sup>1</sup>Our code is available at <https://github.com/lyutyuh/poptimizer>.

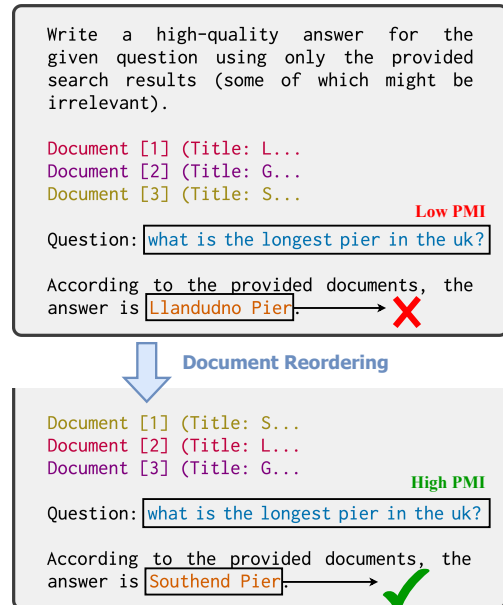


Figure 1: For the *same* question, a permutation of documents with a higher  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  tends to lead to a better answer.

encodes a set of documents retrieved from an external source by a retriever (Karpukhin et al., 2020; Ni et al., 2022). Then, an answer is sampled from the language model. Previous work has explored various empirical approaches to prompt engineering, including the manual design of prompts that mimic human reasoning (Wei et al., 2023; Yao et al., 2023). Recently, Liu et al. (2024) demonstrated that language model performance is significantly influenced by the *order* of the retrieved documents that comprise the context. Specifically, QA accuracy peaks when the gold document<sup>2</sup> is positioned at the beginning or end of the context. Although extensive experimental evidence was adduced to validate this phenomenon (Liu et al., 2024), the underlying mechanisms remain poorly understood. This gap in understanding limits the applicability of these findings in the design and optimization

<sup>2</sup>In factual QA tasks, the document containing the ground truth answer is referred to as the gold document.

of prompts for real-world applications.

While Liu et al.’s (2024) results are interesting, choosing the optimal permutation of the documents requires knowledge of the **answer** and, thus, cannot be directly used to improve RAG. In this article, we develop a proxy for the optimal permutation: We show that the pointwise mutual information between the **question** and the **context** under an LM acts as a useful proxy in determining the optimal permutation. To our knowledge, ours is the first to present in-depth analyses of the relation between question likelihood and model performance under the RAG framework.

Our findings in this paper are summarized in the following list:

- We show that the pointwise mutual information between the **question** and the **context** positively correlates with answer accuracy at the corpus level on NQ-Open (Kwiatkowski et al., 2019; Lee et al., 2019) and ELI5 (Fan et al., 2019).
- Given a question and a fixed set of documents, we demonstrate a strong correlation between the position of the gold document, the PMI between the **question** and the **context**, and QA accuracy.
- We validate the effectiveness of using question likelihood as a gauge for prompt optimization and demonstrate that likelihood-based prompt optimization is a promising direction for future study.

## 2 Setting the Stage

### 2.1 Language Modeling and RAG

**Language Modeling Background.** Let  $\Sigma$  be an **alphabet**, i.e., a finite, non-empty set of **tokens**. A **language model**  $p$  is a distribution over  $\Sigma^*$ , the set of all strings with tokens drawn from  $\Sigma$ . With  $\sigma \preceq \sigma'$ , we denote that  $\sigma$  is a proper prefix of  $\sigma'$ . Let  $Y$  be a  $\Sigma^*$ -valued random variable distributed according to  $p$  and  $\sigma \in \Sigma^*$ . We define the **prefix probability**<sup>3</sup>  $\vec{p}(\sigma)$  as the probability that  $Y$  has  $\sigma$  as a prefix:

$$\vec{p}(\sigma) \triangleq \mathbb{P}_{Y \sim p}(Y \succeq \sigma) \quad (1a)$$

$$= \sum_{\sigma' \in \Sigma^*} \mathbb{1}\{\sigma' \succeq \sigma\} p(\sigma') \quad (1b)$$

The conditional prefix probability  $\vec{p}(\sigma' | \sigma) = \frac{\vec{p}(\sigma \cdot \sigma')}{\vec{p}(\sigma)}$  tells us how certain the model is that  $\sigma'$

<sup>3</sup>See Vieira et al. (2024) for a more in-depth discussion.

naturally follows from its preceding string  $\sigma$ . Finally, we define an **infix probability**, i.e., the probability of generating a string that contains  $\sigma \square \sigma'$ ; where  $\square$  is a gap, as follows

$$\vec{p}(\sigma \square \sigma'') \triangleq \mathbb{P}_{Y \sim p}(Y \succeq \sigma \square \sigma'') \quad (2a)$$

$$= \sum_{\sigma''' \in \Sigma^*} \sum_{\sigma' \in \Sigma^*} \mathbb{1}\{\sigma''' \succeq \sigma \sigma' \sigma''\} p(\sigma''') \quad (2b)$$

**Retrieval-augmented Generation.** Modern language models are often used to perform question-answering tasks. When solving such a task with a language model, string encoding the question **question**  $q \in \Sigma^*$  is given to the model. We assume each question  $q$  has a unique correct **answer**, which we will denote  $a$ . This is, of course, a simplifying assumption, but it does jibe with how question-answering is typically evaluated. We will adorn symbols with  $\tilde{\cdot}$ , e.g.,  $\tilde{a}$ , to indicate an **answer** generated from  $\vec{p}(\cdot | q)$  that may or may not be correct. Generating  $\tilde{a}$  from  $\vec{p}(\cdot | q)$  may be done using either a deterministic method, e.g., beam search, or a stochastic method, e.g., ancestral sampling.<sup>4</sup> In RAG, the model is additionally given a set of documents  $\mathcal{D} = \{d_k\}_{k=1}^K$ , where  $d_k \in \Sigma^*$ , and a permutation of the documents  $\pi: \{1, \dots, K\} \rightarrow \{1, \dots, K\}$ . Given  $\mathcal{D}$  and  $\pi$ , a **context**  $c$  is constructed by concatenating the documents in the order defined by  $\pi$ , i.e.,  $c_{\mathcal{D}}(\pi) \triangleq d_{\pi(1)} \cdot \dots \cdot d_{\pi(K)}$ . Then, we generate an answer from  $\vec{p}(\cdot | c \cdot q)$ . We provide an example below.

**Example 2.1.** Consider  $\mathcal{D} = \{“Llandudno Pier is a Grade II* listed pier...”, “Garth Pier is a Grade II listed structure...”, “Southend Pier is a...”\}$ , and  $\pi(1) = 2, \pi(2) = 1, \pi(3) = 3$ . We have  $c_{\mathcal{D}}(\pi) = “Garth Pier is... Llandudno Pier is a Grade II* listed pier... Southend Pier is...”$ .

Let  $\tilde{a}_{\pi}$  denote an answer generated from  $\vec{p}(\cdot | c_{\mathcal{D}}(\pi) \cdot q)$ . To evaluate the quality of  $\tilde{a}_{\pi}$ , we define an evaluation metric  $g(\tilde{a}_{\pi}, a)$ . In addition, we assume the ground truth answer  $a$  to be unique for a **question–context** pair  $(q, c)$ .

**Pointwise Mutual Information.** In RAG question answering, we consider the following **pointwise mutual information**

$$\text{PMI}(q, c) \triangleq \log \frac{\vec{p}(q | c)}{\vec{p}(q)} \quad (3)$$

<sup>4</sup>In this study, we consider greedy decoding, i.e., beam search with a beam of size 1.

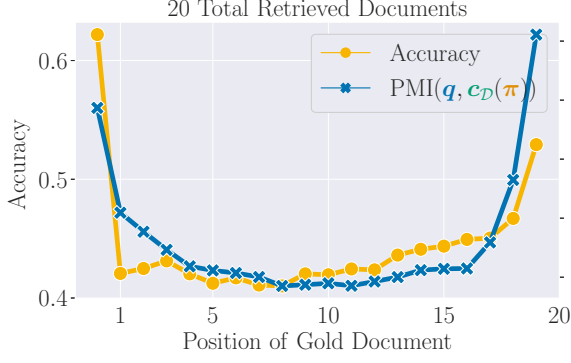


Figure 2: We observe that the PMI and QA accuracy trace a U-shaped curve—nearly in lockstep—as the gold document position within the context changes. The result is computed with LLaMA-3-8B.

between  $q$  and  $c$ , where  $c = d_{\pi(1)} \cdot \dots \cdot d_{\pi(K)}$ . In other words, Eq. (3) measures the degree of association of  $q$  with  $c$ .

## 2.2 A Concrete Hypothesis

Returning to the central goal of this paper, i.e., trying to find a proxy that helps determine the optimal permutation of the documents for RAG, we hypothesize that, given a question  $q$ , a set of documents  $\mathcal{D}$ , and the ground truth answer  $a$ , the pointwise mutual information  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  correlates with  $\log \frac{\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}{1 - \vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}$ , the log odds ratio, and can be deemed a gauge for the expected accuracy of the generated answer. In symbols, our hypothesis is as follows.

**Hypothesis 2.1.** *In RAG question answering, for a fixed question  $q$ , a set of documents  $\mathcal{D}$  permuted by  $\pi$ , and the ground truth answer  $a$ , we have the following relation between  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  and  $\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))$*

$$\begin{aligned} \text{PMI}(q, c_{\mathcal{D}}(\pi)) \\ = a \log \frac{\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}{1 - \vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))} + b \end{aligned} \quad (4)$$

for constants  $a \in \mathbb{R}_{>0}$ ,  $b \in \mathbb{R}$ .

## 2.3 A Bit of Analysis

In words, Hypothesis 2.1 says that when  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  is high, we expect the LM to answer the question  $q$  with high accuracy, and, moreover, we expect the relationship between  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  and accuracy to be affine. Although this is empirically true in many cases (Gonen et al., 2023), we offer an assumption that enables a derivation of this property.

**Assumption 2.1.** *For all question–context pairs  $(q, c)$ , let  $a$  be the correct answer, then we have*

$$\vec{p}(q|c_{\mathcal{D}}(\pi) \square a) = \vec{p}(q|c_{\mathcal{D}}(\pi)) \quad (5a)$$

$$\vec{p}(q|c_{\mathcal{D}}(\pi) \square \bar{a}) = \vec{p}(q) \quad (5b)$$

for any  $\bar{a} \in \Sigma^*$  such that  $\bar{a} \not\sqsubseteq a$ .<sup>5</sup>

We now give a brief qualitative justification of Assumption 2.1. Conditioned on the event that the model *incorrectly* answers the question given the context, Eq. (5b) says that the question  $q$  is not dependent on the provided context. Because, in RAG, we assume the correct answer is given to the model in the context and the model’s job is to retrieve it, our assumption corresponds to the notion that an incorrect response by the model should *not* be influenced by the context. Eq. (5a) corresponds to the notion that since the correct answer is already contained in the context, conditioning on the correct answer does not provide any new information to generating  $q$ .

**Proposition 2.1.** *Under assumptions given in Assumption 2.1, we have*

$$\begin{aligned} \log \frac{\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}{1 - \vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))} \\ = \text{PMI}(q, c_{\mathcal{D}}(\pi)) + C(a, c_{\mathcal{D}}(\pi)) \end{aligned} \quad (6)$$

for an answer-dependent constant  $C(a, c_{\mathcal{D}}(\pi))$ .

*Proof.* See Appendix G. ■

In other words, the pointwise mutual information  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  is equivalent up to an additive constant to the log odds ratio  $\log \frac{\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}{1 - \vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))}$ .

**Foreshadowing the Results.** In the empirical portion of this paper, we test Hypothesis 2.1 through experiments on two QA benchmarks—NQ-Open and ELI5—using a range of state-of-the-art open LMs, including LLaMA-2, LLaMA-3, LLaMA-3.1, Mistral-v0.3, and MPT. Our findings demonstrate that, as the position of relevant information within the input context  $c$  varies, the pointwise mutual information,  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$  and expected answer accuracy  $\vec{p}(a|q \cdot c_{\mathcal{D}}(\pi))$  vary in tandem, i.e., they are strongly correlated. This correlation is illustrated in Figure 2. Specifically, LMs tend to provide better responses to questions where the documents in the context are permuted so as to have higher  $\text{PMI}(q, c_{\mathcal{D}}(\pi))$ . These results

<sup>5</sup>Eq. (5a) and Eq. (5b) can be seen as a form of context-specific conditional independence (Boutillier et al., 1996).

suggest that PMI serves both as a *performance gauge* and as a strong indicator of the position of task-relevant information within the input context. Building on this insight, we propose a direction for prompt optimization through two specific methods. The first selects a permutation  $\pi$  of the documents that maximizes  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  to construct the context  $\mathbf{c}$ . The second builds on the findings of Liu et al. (2024) that the curve traced by permuting the position of the gold document results in a U-shaped curve. We exploit this finding to develop an efficient prompt ordering algorithm. Further experimentation demonstrates that our methods enhance answer accuracy across both datasets for instruction-tuned and base models alike, with the second approach achieving even greater gains.

### 3 PMI Correlates with Performance

As discussed in §2, our first goal is to determine how  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  changes as a function of the permutation  $\pi$  of the documents  $\mathcal{D}$ . Due to Hypothesis 2.1, we expect a strong correlation between  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  and expected answer accuracy.

#### 3.1 Experimental Setup

**Datasets.** We run experiments on two question-answering datasets, namely NQ-Open and ELI5. Details of the datasets are given in Appendix C. Let  $\mathcal{C} \triangleq \{(\mathbf{q}_m, \mathcal{D}_m, \mathbf{a}_m)\}_{m=1}^M$  be a dataset of triples, where each  $\mathbf{a}_m$  represents the ground truth answer to  $\mathbf{q}_m$ .

**Empirical Metrics for LM Evaluation.** In practice, LM performance is often evaluated with rule-based empirical metrics, denoted  $g$ , such as accuracy, instead of the conditional likelihood  $\vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))$ . Although mathematically quantifying the relation between  $g$  and PMI is difficult, we contend that they positively correlate due to recent progress on language model calibration (Zhao et al., 2023), i.e., the alignment between  $\vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))$  and  $g(\tilde{\mathbf{a}}, \mathbf{a})$ . On NQ-Open, the ground truth answer for each question is either a word or a short phrase. The accuracy is 1 when the LM response contains the correct answer as a substring; otherwise, the accuracy is 0. Following Liu et al. (2024), we compute the model’s average accuracy over the entire dataset. On ELI5, the correct answer for each question comprises three sub-claims, and a correct answer is expected to include all of these sub-claims. Examples are illustrated in Appendix B. We follow Gao et al. (2023) and take the recall rate of

sub-claims to be the evaluation metric, which takes values from  $\{0, 1/3, 2/3, 1\}$ . The TRUE model,<sup>6</sup> a T5-XXL model fine-tuned on natural language inference (NLI) tasks, is used to automatically evaluate whether a response entails a sub-claim.

**Language Model Settings.** Most state-of-the-art closed LMs, such as OpenAI’s ChatGPT and Anthropic’s Claude, do not provide direct access to the likelihood of either input or output tokens. Thus, we select leading open LMs for our experiments, focusing on three families: LLaMA-2, LLaMA-3 (Touvron et al., 2023), and Mistral-v0.3 (Jiang et al., 2023). We also evaluate MPT on NQ-Open.<sup>7</sup> Following the settings of Liu et al. (2024), we adopt greedy decoding for all models when generating responses. We set the maximum number of decoded tokens to 100 on NQ-Open and 300 on ELI5.

**Prompt Templates.** We follow the suggested usage and prompt formatting instructions of each LM we use. For chat and instruction-tuned models, we present the context and query to the LM in the role of user, and elicit the response from LMs in the role of assistant. For base models, we elicit responses from LMs as sentence completion.

#### 3.2 Technical Interlude: Sets of Permutations

In many of our experiments, we would like to take a sum or a max over all permutations of  $K$  items, i.e., take a sum or max over the symmetric group  $\mathbb{S}_K$ . However,  $|\mathbb{S}_K| = K!$ , which grows too large to enumerate efficiently. To cope with the size of  $\mathbb{S}_K$ , in this paper, we perform computations over a subset of  $\mathbb{S}_K$ . Specifically, starting a user-specified permutation  $\pi$ , we consider the cyclic group generated by  $(\pi)$  where the group operation is functional composition, as is standard. Let  $\sigma = (1, 2, \dots, K)$  be a shifting permutation. It is easy to see that  $|(\pi)| = K$ , and the  $k^{\text{th}}$  element of  $(\pi)$  is given by  $\tilde{\pi}_k \triangleq \sigma^{k-1} \circ \pi = \underbrace{\sigma \circ \dots \circ \sigma}_{(k-1) \text{ times}} \circ \pi$ ,<sup>8</sup> or,

equivalently we have

$$\tilde{\pi}_k(i) = (i + k - 1) \pmod K, \quad (7)$$

<sup>6</sup>[https://huggingface.co/google/t5\\_xxl\\_true\\_nli\\_mixture](https://huggingface.co/google/t5_xxl_true_nli_mixture)

<sup>7</sup>In our preliminary experiments, MPT fails to generate sufficiently long responses on ELI5, resulting in performance that is not directly comparable to other LMs.

<sup>8</sup>Here,  $\circ$  denotes function composition, the standard product of permutations.

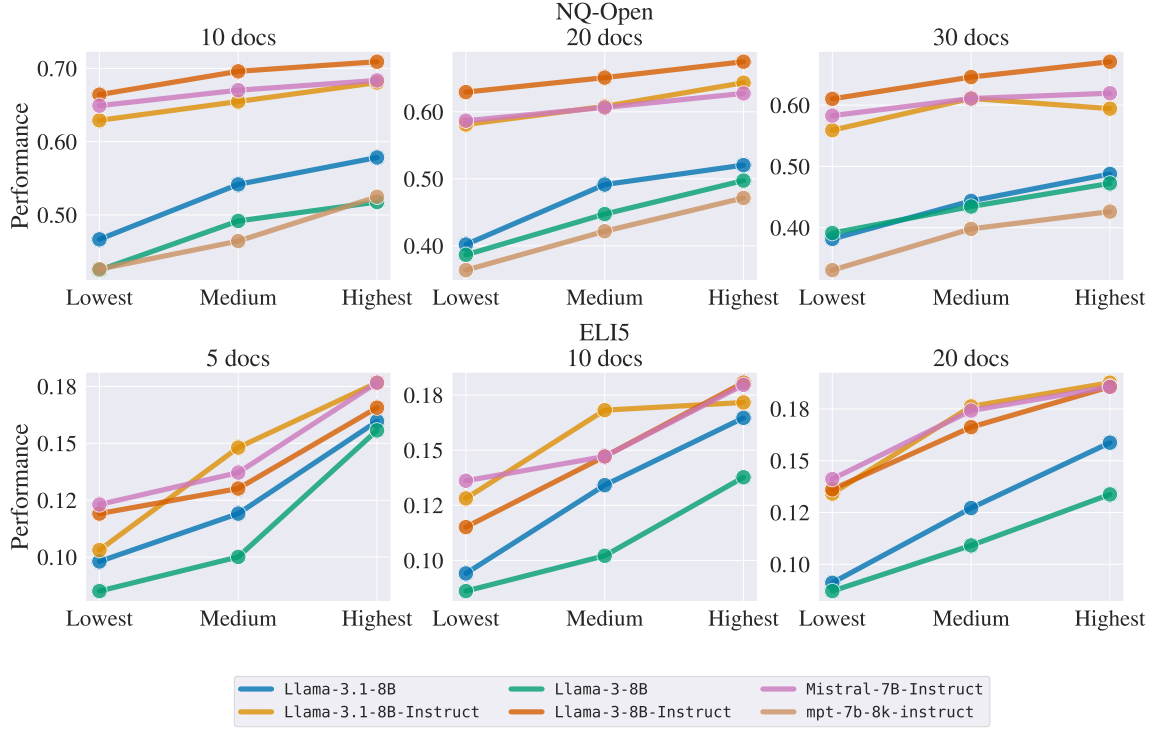


Figure 3: Corpus-level correlation between  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  and answer accuracy on NQ-Open and ELI5.

**Example 3.1.** Given a permutation  $\pi = (1, 2, 3)$ . The cyclic group  $(\pi)$  generated by  $\pi$  is equal to  $\{\tilde{\pi}_1, \tilde{\pi}_2, \tilde{\pi}_3\}$  where  $\tilde{\pi}_1 = (1, 2, 3)$ ,  $\tilde{\pi}_2 = (2, 3, 1)$ , and  $\tilde{\pi}_3 = (3, 1, 2)$ .

### 3.3 Results

We now discuss our empirical findings.

**Corpus-Level Correlation.** As our first evaluation metric, we consider a corpus-level correlation. For each  $\mathbf{q}_m, \mathcal{D}_m$  in a corpus  $\mathcal{C}$ , we compute the average PMI for the  $m^{\text{th}}$  instance as follows

$$\rho_m \triangleq \frac{1}{K} \sum_{k=1}^K \text{PMI}(\mathbf{q}_m, \mathbf{c}_{\mathcal{D}_m}(\tilde{\pi}_k)) \quad (8)$$

We then bin the elements of  $\{\rho_m\}_{m=1}^M$  into three bins according to which tertile they fall in when  $\{\rho_m\}_{m=1}^M$  are arranged into a histogram. Then, we compute the average sub-claim recall rate (ELI5) and accuracy (NQ-Open) for each bin. Our results, shown in Figure 3, demonstrate that LMs tend to perform better on the prompts with a higher  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  compared to those with lower  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$ .

**Instance-Level Correlation.** We further analyze the instance-level correlation between  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  and accuracy by varying context while keeping the question fixed. In symbols, we

compute

$$\eta_k \triangleq \frac{1}{M} \sum_{m=1}^M \text{PMI}(\mathbf{q}_m, \mathbf{c}_{\mathcal{D}_m}(\tilde{\pi}_k)) \quad (9)$$

where  $\tilde{\pi}_k$  is the permutation in which the  $k$ -th document  $\mathbf{d}_k$  contains relevant information. We then plot the curve of  $\{\eta_k\}_{k=1}^K$ , to see how PMI is affected by the position of a relevant document within a context.

**Revisiting Liu et al. (2024).** We now revisit the findings of Liu et al. (2024), who observed a drop in QA accuracy when the gold document is positioned within the middle of  $\mathbf{c}$ . We first experiment on NQ-Open by varying the position of the gold document<sup>9</sup> in  $\mathbf{c}$ . The set of retrieved documents and the order of non-gold documents remain the same. As the gold document is placed in different positions in  $\mathbf{c}$ , we find that both  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  and QA accuracy fluctuate—nearly in lockstep. To further explore this correlation between  $\text{PMI}(\mathbf{q}, \mathbf{c})$  and QA accuracy, we calculate the expected accuracy with the prompt of the highest and lowest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$ . Results are given in Table 1, showing that LMs perform better when the document order in the prompt leads to the highest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$ , while

<sup>9</sup>In NQ-Open, exactly one retrieved document is marked as the gold document for each question.

the prompt with the lowest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  results in inferior performance.

**Experiments on ELI5.** Compared to NQ-Open, ELI5 is a more challenging long-form QA dataset where questions are mostly about *how/why/what*, and the answers are expected to be more comprehensive and cover multiple aspects. Due to the lack of gold document annotations on ELI5, we adopt permutations from the cyclic group ( $\pi$ ) and random shuffling. In random shuffling for  $K$  documents, we randomly shuffle the document set  $K$  (i.e., same as the number of documents) times and obtain  $K$  document sequences for consistency. Given multiple prompts for a question, among which only the document orders in the **context** are different, we calculate the average performance of the prompts with the highest and lowest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  for each question in the same fashion as described in §3.3 for NQ-Open. Results in Table 2 show that LMs achieve higher answer accuracy on the prompts with the highest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$ , compared with the prompts with the lowest  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$ . This indicates LMs can better answer questions with higher question likelihood through document shuffling, demonstrating the strong instance-level correlation between  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  with answer accuracy.

## 4 Improving RAG via Reordering

In §3.3, we offered evidence for Hypothesis 2.1, i.e., that  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  correlates with model performance. In light of this finding, we propose two methods to permute the documents presented to the LM in RAG *without* knowledge of the **answer**.

### 4.1 Method 1: Search by PMI

Our empirical findings showed that the permutation of the documents in the **context** that leads to the highest value of  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  leads to superior performance on QA tasks. This suggests a natural algorithm

$$\pi^* = \arg \max_{\pi \in \mathbb{S}_K} \text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi)) \quad (10)$$

However, as discussed in §3.2, the set of all permutations (the symmetric group)  $\mathbb{S}_K$  is too large to enumerate. Thus, we fall back on a simple approximation. Given a user-provided permutation  $\pi$ , we search over the cyclic group generated by  $\pi$ , denoted as  $\langle \pi \rangle$ . Using the notation introduced in

§3.2, we choose  $\tilde{\pi}_{k^*}$  where we select

$$k^* = \arg \max_{k=1}^K \text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\tilde{\pi}_k)), \quad (11)$$

where  $\tilde{\pi}_k$  is defined in §3.2.

### 4.2 Method 2: Search by Curvature

We now develop a second algorithm based on the observation in Figure 4 that accuracy and PMI change simultaneously and exhibit a U-shaped curve as the gold document position within the permutation of documents in  $\mathbf{c}$ . Our algorithm is based on a discrete notion of convexity and an assumption based on our findings in §3.3, which we introduce in the abstract below.

**Technical Interlude Discrete Convexity.** A sequence of real values  $\{a_n\}_{n=1}^N$  is called **convex** if we have

$$\Delta_n^2 \triangleq 2a_n - a_{n+1} - a_{n-1} \leq 0 \quad (12)$$

for all  $n \in \{2, \dots, N-1\}$ . In the abstract, the problem we wish to solve is this: Given an arbitrary finite sequence of reals  $\{b_n\}_{n=1}^N$ , find a permutation  $\tau: [N] \rightarrow [N]$  that renders  $\{b_n\}_{n=1}^N$  convex, i.e., that Eq. (12) holds after applying the permutation to the sequence’s indices. We call such a choice of  $\tau$  a **convex permutation**. Note that convex permutations may not always exist.<sup>10</sup> To achieve a U-shape curve, we do not just want a convex permutation, but in addition the one that results in a convex sequence that has as much upwards curvature as possible. In other words, if  $\tau$  is a convex permutation, then, *in addition* we want the following sum to be *minimized*

$$\sum_{n=2}^{N-1} \Delta_n^2 = -(b_1 + b_N) + \sum_{n=2}^{N-1} b_m \quad (13a)$$

$$= -2(b_1 + b_N) + B \quad (13b)$$

$$\leq 0, \quad (13c)$$

where  $B \triangleq \sum_{n=1}^N b_n$ . However, because  $B$  is constant, the total curvature induced by a convex permutation  $\tau$  *only* depends on  $b_{\tau(1)}$  and  $b_{\tau(N)}$ . This implies that we simply need to choose the endpoints to be those elements of  $\{b_{\tau(n)}\}_{n=1}^N$  that are largest; we can always permute the remaining  $(N-2)$  elements to ensure the permutation is convex afterward. Thus, relaxing the requirement that

<sup>10</sup>E.g., the sequence  $[0, 1, 1, 1]$  has no convex permutation.

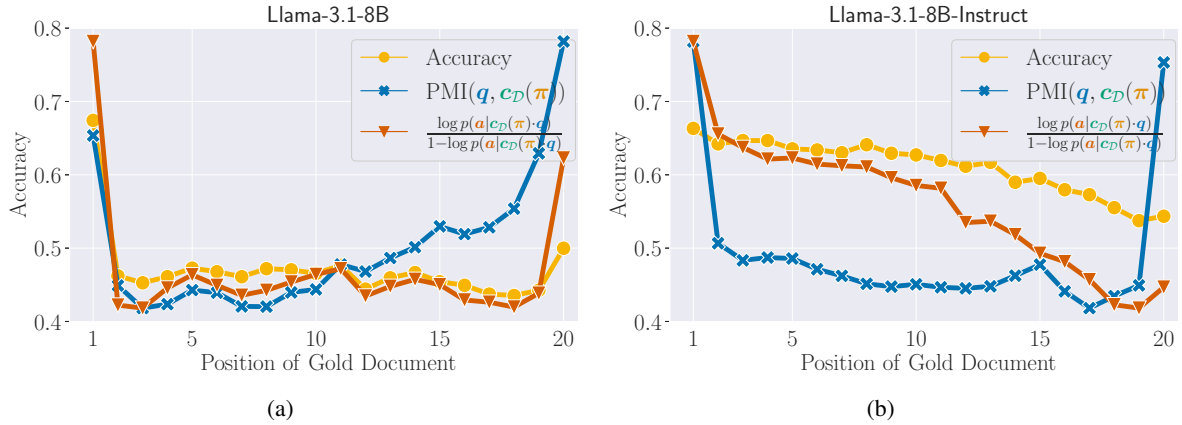


Figure 4: QA accuracy, PMI, and log odds ratio of answer likelihood on 20 docs evaluated on LLaMA-3.1-8B and LLaMA-3.1-8B-Instruct.

#Doc	PMI( $q, c_{\mathcal{D}}(\pi)$ )	Mistral-7B-Inst	LLaMA-3-8B	LLaMA-3.1-8B	LLaMA-3-8B-Inst	LLaMA-3.1-8B-Inst	MPT-7B-8K-Inst
NQ-Open							
10	Highest	<b>68.69</b> (-2.52)	<b>54.04</b> (-1.84)	<b>56.72</b> (-2.41)	<b>71.58</b> (-1.80)	<b>66.13</b> (-2.16)	<b>48.93</b> (-2.80)
	Lowest	66.98 (-2.89)	49.30 (-2.03)	53.29 (-2.72)	71.29 (-2.01)	65.70 (-2.43)	46.97 (-3.38)
20	Highest	<b>64.86</b> (-2.45)	<b>52.05</b> (-1.99)	<b>52.50</b> (-2.40)	<b>69.00</b> (-1.83)	<b>62.97</b> (-2.21)	<b>42.25</b> (-2.70)
	Lowest	62.60 (-2.83)	46.91 (-2.03)	48.51 (-2.72)	67.68 (-2.01)	61.05 (-2.43)	42.09 (-3.23)
30	Highest	<b>57.70</b> (-2.52)	<b>50.30</b> (-1.88)	<b>50.00</b> (-2.60)	64.36 (-1.84)	<b>60.95</b> (-2.41)	<b>39.31</b> (-2.56)
	Lowest	53.96 (-2.92)	45.27 (-2.03)	46.42 (-2.83)	<b>65.12</b> (-2.03)	59.55 (-2.65)	39.12 (-3.05)

Table 1: The table displays instance-level correlations between PMI( $q, c_{\mathcal{D}}(\pi)$ ) and answer accuracy. We also compute the average answer accuracy over prompts that yield the highest and lowest PMI( $q, c_{\mathcal{D}}(\pi)$ ) as the gold document placed at different positions in the document sequence for each instance. The answer accuracy and the average PMI( $q, c_{\mathcal{D}}(\pi)$ ) are reported in the table.

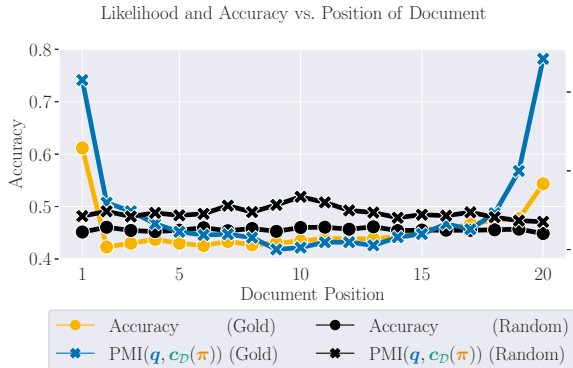


Figure 5: When the position of the gold document changes, both PMI( $q, c$ ) and accuracy curves are U-shaped. In contrast, both curves are flat for randomly chosen documents.

the permutation be convex, we choose a permutation  $\tau$  such that  $b_{\tau(1)} + b_{\tau(N)}$  is maximized. This definition motivates a new definition: We call a sequence  $\{b_n\}_{n=1}^N$  **U-shaped** iff  $b_1 \geq b_i$  and  $b_N \geq b_i$  for  $i \in \{2, 3, \dots, N-1\}$ .

**A Simple Algorithm.** The abstract discussion in the previous paragraph suggests a simple algorithm. First, we construct a real-valued sequence

$$b_{\tau(k)} \triangleq \text{PMI}(q, c_{\mathcal{D}}(\tilde{\pi}_k)) \quad (14)$$

of length  $K$  where  $\tau: [K] \rightarrow [K]$  is a permutation and  $\text{PMI}(q, c_{\mathcal{D}}(\tilde{\pi}_k))$  is defined in §3.2. Then, relaxing the requirement that the permutation be convex, we optimize

$$\tau^* = \arg \max_{\tau \in (\pi)} b_{\tau(1)} + b_{\tau(K)} \quad (15)$$

While, in general,  $\tau^*$  may not be convex, we do have a guarantee that  $\tau^*$  will induce a U-shaped sequence. To compute the optimization problem given in Eq. (15), we sort  $\tau \in (\pi)$  according to  $b_{\tau(1)} + b_{\tau(K)}$  in descending order, obtaining the sequence  $\{\tau_k\}_{k=1}^K$ . Then, we construct the resulting permutation  $\tau' = (\tau_1(1), \tau_2(1), \dots, \tau_K(1))$ , among which  $d_{\tau'(1)}$  is most likely to be the gold document.

### 4.3 Results and Analysis

Shown in Table 3, both search by PMI and search by curvature can boost answer accuracy. On NQ-Open, where only one document in the sequence is relevant to the question, gold document reordering significantly improves the answer accuracy and narrows the gap to the upper bound. Furthermore, on the more challenging and practical QA benchmark ELI5, we also observe a modest improvement

#Doc	PMI( $q, c_D(\pi)$ )	Mistral-7B-Inst	LLaMA-3-8B	LLaMA-3.1-8B	LLaMA-3-8B-Inst	LLaMA-3.1-8B-Inst
ELI5 with Rotational permutation						
5	Highest	<b>13.97</b> (-3.72)	<b>11.37</b> (-2.23)	<b>12.60</b> (-2.28)	<b>14.23</b> (-2.21)	<b>13.97</b> (-2.26)
	Lowest	13.50 (-4.06)	11.10 (-2.39)	12.50 (-2.43)	13.17 (-2.54)	13.93 (-2.48)
10	Highest	<b>15.23</b> (-3.53)	11.27 (-2.19)	12.50 (-2.29)	<b>14.50</b> (-2.10)	<b>16.17</b> (-2.23)
	Lowest	14.47 (-3.99)	<b>11.50</b> (-2.39)	<b>13.10</b> (-2.48)	14.07 (-2.55)	15.77 (-2.54)
20	Highest	<b>16.20</b> (-2.13)	11.13 (-2.19)	<b>12.77</b> (-2.28)	<b>16.20</b> (-2.13)	<b>17.17</b> (-2.18)
	Lowest	15.80 (-2.73)	<b>11.20</b> (-2.42)	12.13 (-2.48)	15.80 (-2.73)	15.67 (-2.54)
ELI5 with Random Shuffling						
5	Highest	<b>14.27</b> (-3.73)	10.73 (-2.24)	<b>12.57</b> (-2.28)	<b>14.10</b> (-2.23)	<b>14.20</b> (-2.27)
	Lowest	14.10 (-4.04)	<b>11.20</b> (-2.39)	12.33 (-2.42)	12.77 (-2.52)	14.00 (-2.48)
10	Highest	<b>15.63</b> (-3.54)	<b>11.47</b> (-2.19)	<b>12.73</b> (-2.29)	<b>15.70</b> (-2.11)	<b>16.90</b> (-2.23)
	Lowest	15.07 (-3.97)	11.23 (-2.39)	12.20 (-2.48)	14.57 (-2.52)	16.70 (-2.53)
20	Highest	16.10 (-3.44)	10.83 (-2.19)	<b>12.60</b> (-2.28)	<b>16.13</b> (-2.14)	<b>17.20</b> (-2.18)
	Lowest	<b>16.53</b> (-4.00)	<b>11.20</b> (-2.42)	11.87 (-2.49)	15.53 (-2.71)	17.10 (-2.54)

Table 2: Instance-level correlation between PMI( $q, c$ ) and answer accuracy on ELI5. The average QA accuracy is computed over prompts that yield the highest and lowest PMI( $q, c$ ) as the input documents are reordered with (1) rotational reordering and (2) random shuffling as introduced in §3.3. We report the QA accuracy and the average PMI( $q, c$ ) in parentheses.

in answer accuracy, indicating that improving question likelihoods via document reordering can effectively obtain better LM responses.

Regarding efficiency, our proposed methods are mildly time-dependent thanks to the parallelizable computation of question likelihoods, where only the LM encoding module is used, with no reliance on LM decoding.<sup>11</sup> Shown in Table 4, in our experiments, the average runtime for decoding a response of an instance in ELI5 is 10 seconds, while it only takes an extra 0.8 seconds and 2 seconds, respectively, to encode the input prompts of naïve likelihood-based selection and gold document reordering. The increment in timely cost is marginal compared with heuristic prompt engineering which requires whole decoding to judge the prompt quality (e.g., an extra 10 seconds for decoding another candidate prompt).

In summary, both proposed methods are effective and efficient. Although the improvement on ELI5 is relatively marginal compared to that on NQ-Open, given the more challenging nature of long answers and no specified gold document on ELI5, it still indicates that optimizing prompts with  $\vec{p}(q | c)$  as a gauge is a promising direction.

**Experimental Setup.** We experiment on the ELI5 dataset and a subset of 500 questions from NQ-Open, using Mistral-7B-Inst-v0.3,

<sup>11</sup>LM decoding (i.e., generation) requires a runtime approximately proportional to the number of generated tokens. Empirically, the only extra computational time for our methods is on the encoding phase for calculating likelihoods, so the overall runtime is the vanilla  $\text{Runtime}_{\text{LM}}$  plus *one* extra LM going through, which is equivalent to generating the response with *one* additional token.

Model	Baseline	PMI	Curvature	Upper Bound
NQ-Open (Answer Accuracy)				
Mistral	62.89	65.18	<b>65.72</b>	69.24
LLaMA-3.1	47.74	51.29	<b>51.36</b>	66.88
LLaMA-3.1-Inst	61.49	63.34	<b>63.56</b>	66.35
ELI5 (Answer Accuracy)				
Mistral	15.35	<b>15.63</b>	15.40	-
LLaMA-3.1	12.61	12.73	<b>13.33</b>	-
LLaMA-3.1-Inst	16.14	<b>16.90</b>	16.83	-

Table 3: Performance of our methods on NQ-Open and ELI5, the number of documents  $K$  is set to 20 and 10. Mistral, LLaMA and LLaMA-Inst stands for Mistral-7B-Inst-v0.3, LLaMA-3.1-8B and LLaMA-3.1-8B-Inst respectively. Baseline refers to the mean performance over  $K$  random document shuffling on each instance. The upper bound on NQ-Open is calculated as the performance when positioning the gold document at the beginning of the document sequence, which is not applicable for ELI5 since no gold document is marked in this practical dataset.

LLaMA-3.1-8B, and LLaMA-3.1-8B-Inst. Each question is associated with 10 (on ELI5) and 20 (NQ-Open) retrieved documents.

#### 4.4 Synthetic Experiment

Real-world datasets might have been used during the training of LLMs. Thus, their likelihoods might exhibit an *exposure bias* (Bengio et al., 2015; Ranzato et al., 2016; Cotterell et al., 2024). To avoid such potential bias, we follow Liu et al. (2024) and conduct a synthetic key-value retrieval experiment.

**Key-Value Retrieval.** To imitate question-answering tasks on random strings, we construct Python-style key-value pairs in which the keys and



Decoding	Likelihood Based	Gold Document
10s	0.8s	2s

Table 4: The average runtime for decoding an LLM response vs. the extra time for the two proposed methods.

values are UUID strings of 32 hexadecimal digits. An example is given in Figure 14. In Tables 5 to 7, we observe that both  $\text{PMI}(\mathbf{q}, \mathbf{c})$  and  $\vec{p}(\mathbf{a} | \mathbf{c} \cdot \mathbf{q})$  show synchronous U-shaped patterns as the location of the key in context changes, consistent with the RAG-based QA experiments in §3.3, indicating the generalizability of the findings on unseen data.

## 5 Discussion

### 5.1 Instruction-tuned vs. Base Models

In our analysis, we find base LMs, e.g., LLaMA-3-8B, tend to be more sensitive to the permutation of the documents. Specifically, we observe that QA performance drops when the gold document is placed in the middle of the document sequence. On the other hand, the performance of instruction-tuned models is more robust to permutations of the documents in the context, as shown in Figure 4. However, we still do observe a U-shaped curve, but the drop in QA performance is less significant for the instruction-tuned model when the gold document is positioned at the middle.

The fact that PMI serves as a useful gauge for both the base and instruction-tuned models suggests that  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  is affected little by the instruction tuning.

### 5.2 Placing the Context after the Question

In our experiments, we only explore the correlation between PMI and accuracy when the question follows the context. However, one could also use a prompt template in which the context follows the question. We remark that in this case, PMI can be computed according to the equation

$$\text{PMI}(\mathbf{q}, \mathbf{c}) = \log \frac{\vec{p}(\mathbf{c} | \mathbf{q})}{\vec{p}(\mathbf{c})}. \quad (16)$$

## 6 Conclusion

In our study, we analyzed the relationship between the PMI between question and context  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  and question-answering performance under the retrieval-augmented generation framework. Through experimentation, we demonstrated that  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  is affected by the order

of documents in the input context. We find evidence for a positive correlation between question likelihood and answer accuracy at both the corpus level and instance level. Our findings show that it is possible to use  $\text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi))$  to gauge language model performance and improve the quality of input prompts. We propose two practical methods for prompt optimization based on these findings. Experimental results show that both effectively and efficiently improve LM’s accuracy on QA tasks, demonstrating that using PMI as a gauge for optimizing prompts is a promising direction.

## Limitations

One major limitation of our work is that only open-source LMs are studied in this work since we need full access probabilities under the LM. Thus, closed language models such as GPT-4 cannot be used for selecting permutations. In addition, our prompt modification is limited to document permutation in this work. Other prompt modifications may also contribute to obtaining a higher  $\text{PMI}(\mathbf{q}, \mathbf{c})$  and improve QA performance.

## References

- Akari Asai, Xinyan Yu, Jungo Kasai, and Hanna Hajishirzi. 2021. [One question answering model for many languages with cross-lingual dense passage retrieval](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 7547–7560. Curran Associates, Inc.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. [Scheduled sampling for sequence prediction with recurrent neural networks](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, page 1171–1179, Cambridge, MA, USA. MIT Press.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George 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*, pages 2206–2240. PMLR.
- Craig Boutilier, Nir Friedman, Moises Goldszmidt, and Daphne Koller. 1996. Context-specific independence in Bayesian networks. In *Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence*, page 115–123.
- Ryan Cotterell, Anej Svete, Clara Meister, Tianyu Liu, and Li Du. 2024. [Formal aspects of language modeling](#). *Preprint*, arXiv:2311.04329.

- Sabit Ekin. 2023. [Prompt engineering for chatgpt: a quick guide to techniques, tips, and best practices](#). *Authorea Preprints*.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.
- Louie Giray. 2023. [Prompt engineering with chatgpt: a guide for academic writers](#). *Annals of biomedical engineering*, 51(12):2629–2633.
- Hila Gonen, Srinu Iyer, Terra Blevins, Noah Smith, and Luke Zettlemoyer. 2023. [Demystifying prompts in language models via perplexity estimation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10136–10148, Singapore. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Preprint*, arXiv:2112.09118.
- Gautier Izacard and Edouard Grave. 2021. [Leveraging passage retrieval with generative models for open domain question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2024. [Atlas: Few-shot learning with retrieval augmented language models](#). *J. Mach. Learn. Res.*, 24(1).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *Preprint*, arXiv:2310.06825.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich K  ttler, Mike Lewis, Wen-tau Yih, Tim Rockt  schel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS ’20*, Red Hook, NY, USA. Curran Associates Inc.
- Nelson Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7001–7025, Singapore. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Lijia Ma, Xingchen Xu, and Yong Tan. 2024. [Crafting knowledge: Exploring the creative mechanisms of chat-based search engines](#). *arXiv preprint arXiv:2402.19421*.
- Ggaliwango Marvin, Nakayiza Hellen, Daudi Jjingo, and Joyce Nakatumba-Nabende. 2023. [Prompt engineering in large language models](#). In *International conference on data intelligence and cognitive informatics*, pages 387–402. Springer.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. 2022. [Teaching language models to support answers with verified quotes](#). *arXiv preprint arXiv:2203.11147*.

- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. [Webgpt: Browser-assisted question-answering with human feedback](#). *arXiv preprint arXiv:2112.09332*.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, and Yinfei Yang. 2022. [Large dual encoders are generalizable retrievers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9844–9855, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. [How context affects language models’ factual predictions](#). In *Automated Knowledge Base Construction*.
- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick Lewis, Barlas Oğuz, Edouard Grave, Wen-tau Yih, et al. 2021. [The web is your oyster-knowledge-intensive nlp against a very large web corpus](#). *arXiv preprint arXiv:2112.09924*.
- Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang Zhu, and Michael Zeng. 2023. [Automatic prompt optimization with “gradient descent” and beam search](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, Singapore. Association for Computational Linguistics.
- Jirui Qi, Gabriele Sarti, Raquel Fernández, and Arianna Bisazza. 2024. [Model internals-based answer attribution for trustworthy retrieval-augmented generation](#). *arXiv preprint arXiv:2406.13663*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yin-heng Li, Aayush Gupta, HyoJung Han, Sevien Schulhoff, et al. 2024. [The prompt report: A systematic survey of prompting techniques](#). *arXiv preprint arXiv:2406.06608*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint, arXiv:2302.13971*.
- Tim Vieira, Ben LeBrun, Mario Giulianelli, Juan Luis Gastaldi, Brian DuSell, John Terilla, Timothy J. O’Donnell, and Ryan Cotterell. 2024. [From language models over tokens to language models over characters](#). *Preprint, arXiv:2412.03719*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint, arXiv:2201.11903*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). *Preprint, arXiv:2305.10601*.
- Yao Zhao, Mikhail Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu. 2023. [Calibrating sequence likelihood improves conditional language generation](#). In *The Eleventh International Conference on Learning Representations*.
- Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. [Large language models are human-level prompt engineers](#). In *The Eleventh International Conference on Learning Representations*.

## A Related Work

### A.1 Prompt Engineering

Prompt engineering is important for making the best use of LMs in real-world applications (Giray, 2023; Ekin, 2023; Gonen et al., 2023). The most straightforward prompt engineering method is to manually design prompts using heuristics, which requires human experts to design prompts based on domain-specific knowledge and select the prompts that lead to better performance on downstream tasks (Zhou et al., 2023; Marvin et al., 2023). Meanwhile, another line of work explores automatic approaches for prompt engineering (Gao et al., 2021; Pryzant et al., 2023). However, they both require decoding for outputs from LMs to evaluate the quality of prompts, thus incurring high computational costs.

### A.2 Retrieval-Augmented Generation

Retrieval-augmented generation is a technique for improving LMs’ ability to solve knowledge-intensive tasks (Lewis et al., 2020; Asai et al., 2021; Borgeaud et al., 2022). In the RAG framework, a set of documents relevant to a user query is retrieved from an external source and inserted into prompts as a **context**, to provide additional information to the LM and improve response quality (Petroni et al., 2020; Lewis et al., 2020). RAG tasks can be divided into two types: short-form and long-form, depending on the topic of the questions and the format of the expected answers. Short-form QA (Izacard and Grave, 2021; Liu et al., 2024) usually concerns factual questions about real-world facts. The expected answers are often unambiguous and concrete words or short phrases. Long-form QA (Fan et al., 2019; Gao et al., 2023) involves *how*, *why*, and *what* questions that seek more comprehensive responses.

### A.3 Effect of Document Order in RAG

Liu et al. (2024) finds that LMs perform better when the document with relevant information is positioned at the beginning or the end of the prompt.<sup>12</sup> Specifically, when moving the task-relevant information from the beginning to the end of the document sequence, answer accuracy exhibits a U-shaped trend on a multi-document QA task and a synthetic key–value retrieval task, both using RAG pipelines. However, Liu et al. (2024) mainly focuses on an empirical study with less in-depth analysis, resulting in a gap between the phenomenon and its practical implications. In this work, we attempt to bridge this gap.

## B Illustration of Evaluation Metrics

The evaluation metrics for NQ-Open and ELI5 are illustrated with two examples in Figure 6.

## C Datasets

**NQ-Open.** We first experiment on the NQ-Open dataset following Liu et al. (2024). This dataset covers 2655 factual questions curated from the Natural Questions dataset (Kwiatkowski et al., 2019; Lee et al., 2019) under CC-BY-SA-3.0 license. Each question is accompanied by  $K$  documents retrieved from Wikipedia, among which *exactly one* contains the answer to the question, namely the gold document. The remaining  $k - 1$  documents are termed **distractors**, which are relevant to the topic of the question but do not contain any ground truth answers, retrieved using Contriever (Izacard et al., 2022). In our experiments, the total number of documents  $K$  is drawn from  $\{10, 20, 30\}$ .<sup>13</sup>

**ELI5.** To validate the generality of our findings, we also experiment on an open-ended non-factual QA dataset ELI5 (Fan et al., 2019) under the BSD license. ELI5 consists of questions beginning with *how*, *why* or *what* curated from the Reddit forum “Explain Like I’m Five”<sup>14</sup>, where the answers are expected to be more comprehensive and diverse. Each question is accompanied by  $K$  documents retrieved from Sphere (Piktus et al., 2021)—a filtered version of Common Crawl<sup>15</sup>, where  $K$  is taken to be  $\{5, 10, 20\}$

<sup>12</sup>In Liu et al.’s (2024) experimental settings, the gold document is unique in a prompt for each question.

<sup>13</sup>We remark that NQ-Open was specifically synthesized to examine how answer accuracy is affected by changing the position of relevant information. In real-world applications, the retrieved documents for one question may contain multiple gold documents or none. Nevertheless, it mimics the RAG setup underlying many commercial generative search and QA systems.

<sup>14</sup><https://www.reddit.com/r/explainlikeimfive/>

<sup>15</sup><https://commoncrawl.org>

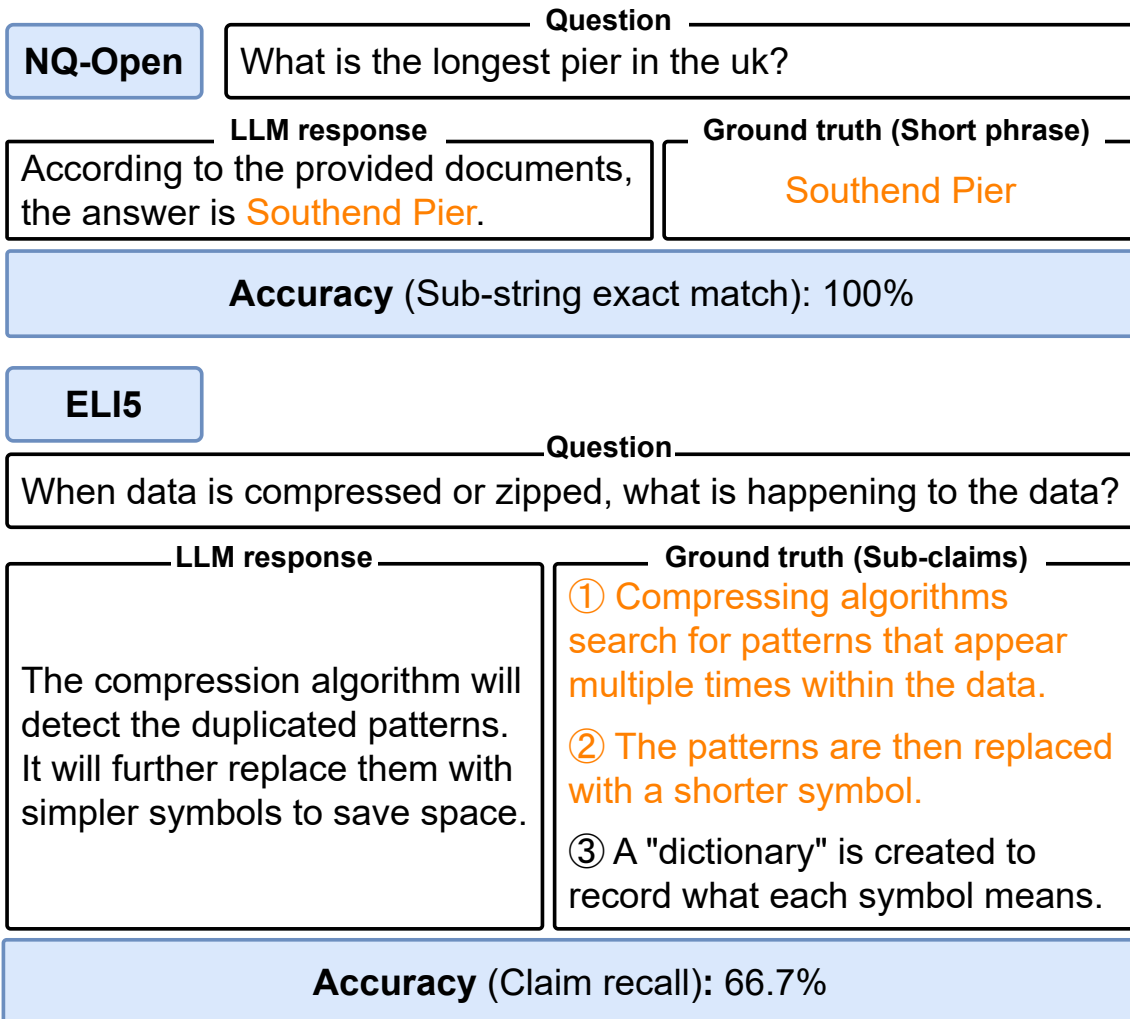


Figure 6: Evaluation metrics used in our experiments. On NQ-Open, the evaluation metric is exact string match. On ELI5, a pretrained NLI model is used to evaluate whether the LM output entails the reference claims.

Write a high-quality answer for the given question using only the provided search results (some of which might be irrelevant).

Document [1](Title: Southend Pier) Southend Pier is a major landmark in . . .

Document [2](Title: Llandudno Pier) Llandudno Pier Llandudno Pier is a Grade II\* listed pier in the seaside resort of Llandudno. . .

Document [3](Title: Garth Pier) Garth Pier Garth Pier is a Grade II listed structure in Bangor. . .

. . .

Question: what is the longest pier in the uk

According to the provided documents, the answer is Southend Pier.

Figure 7: An example prompt and LM output on NQ-Open. The prompt comprises (1) an instruction that describes the task to be solved, (2) a **context** that contains the information for solving the task, in which the **gold document** contains the ground truth answer, and (3) a **question** that describes the specific query. At the end of the prompt, we append an exemplar output that gives the ground truth **answer** to the **question** for evaluating the likelihood of the answer.

Instruction: Write an accurate, engaging, and concise answer for the given question using only the provided search results (some of which might be irrelevant). Use an unbiased and journalistic tone.

Document [1](Title: Trash Islands - the Ocean Garbage Patch): Trash Islands Trash Islands of the Pacific and Atlantic Oceans. . .

Document [2](Title: Where does our garbage go? - Sea Turtle Camp): Pacific Garbage Patch Landfills are a common human solution for disposing of trash on land. . .

Document [3](Title: Plastic pollution crisis: How waste ends up in our oceans - Y108): our ecosystems as a whole. Plastic is non-biodegradable. Every year, about 8-million tons of plastic. . .

. . .

Question: how does so much of our trash end up in the ocean?

According to various sources, a significant portion of the world’s trash ends up in the ocean due to a combination of factors. While it’s often. . . individuals is necessary to mitigate the problem of plastic pollution in the world.  
[Answer length: 242 words]

Figure 8: An example prompt and LM output on ELI5. The prompt comprises (1) an instruction that describes the task to be solved, consistent with previous works on ELI5 (Gao et al., 2023; Qi et al., 2024), (2) a **context** that contains the information for solving the task, but *no gold document* is marked, and (3) a **question** that describes the specific query. At the end of the prompt, we append an exemplar output that gives the ground truth **answer** to the **question** for evaluating the likelihood of the answer.

to avoid truncation due to the long questions and LMs responses for the long-form QA task. In contrast to NQ-Open, ELI5 does not provide the annotations of gold documents, which aligns with real-world RAG application scenarios, making it a more practical and challenging dataset (Nakano et al., 2021; Menick et al., 2022; Liu et al., 2023).

### D Prompt Templates

The prompt templates used for our experiments are given in Figures 7–9.

```

{
  "749d280d-8d74-4a2b-87fa-e2a13b689892":
    "51f95eb8-1f16-4bbf-a7be-6109e581fc04",
  "6618b34a-08b6-46a8-a438-aedc1a2a4635":
    "3e93dc61-1e82-46b1-94be-7ef2e63746e5",
  ...
}
Key: "749d280d-8d74-4a2b-87fa-e2a13b689892"
Value:

```

Figure 9: An example of synthetic data for key-value retrieval.

Key Location	$\vec{p}(q   c)$	$\vec{p}(a   c \cdot q)$
0	-3.96	-0.76
34	-6.60	-0.87
69	-7.87	-0.82
104	-8.70	-1.08
139	-8.03	-0.76

Table 5: Question likelihood and answer likelihood on synthetic key-value retrieval task using Llama-3.1-8B-Instruct model.

Key Location	$\vec{p}(q   c)$	$\vec{p}(a   c \cdot q)$
0	-3.01	-0.08
34	-6.22	-0.15
69	-6.86	-0.31
104	-7.87	-0.27
139	-7.33	-0.07

Table 6: Question likelihood and answer likelihood on synthetic key-value retrieval task using Llama-3.1-8B model.

Key Location	$\vec{p}(q   c)$	$\vec{p}(a   c \cdot q)$
0	-4.03	-0.00
34	-6.31	-0.12
69	-8.15	-0.23
104	-9.67	-0.15
139	-8.77	-0.04

Table 7: Question likelihood and answer likelihood on synthetic key-value retrieval task using Mistral-7B-Instruct model.

## E Full Results on NQ-Open

We show the full results on NQ-Open in Figures 10–13.

## F Synthetic Experiment

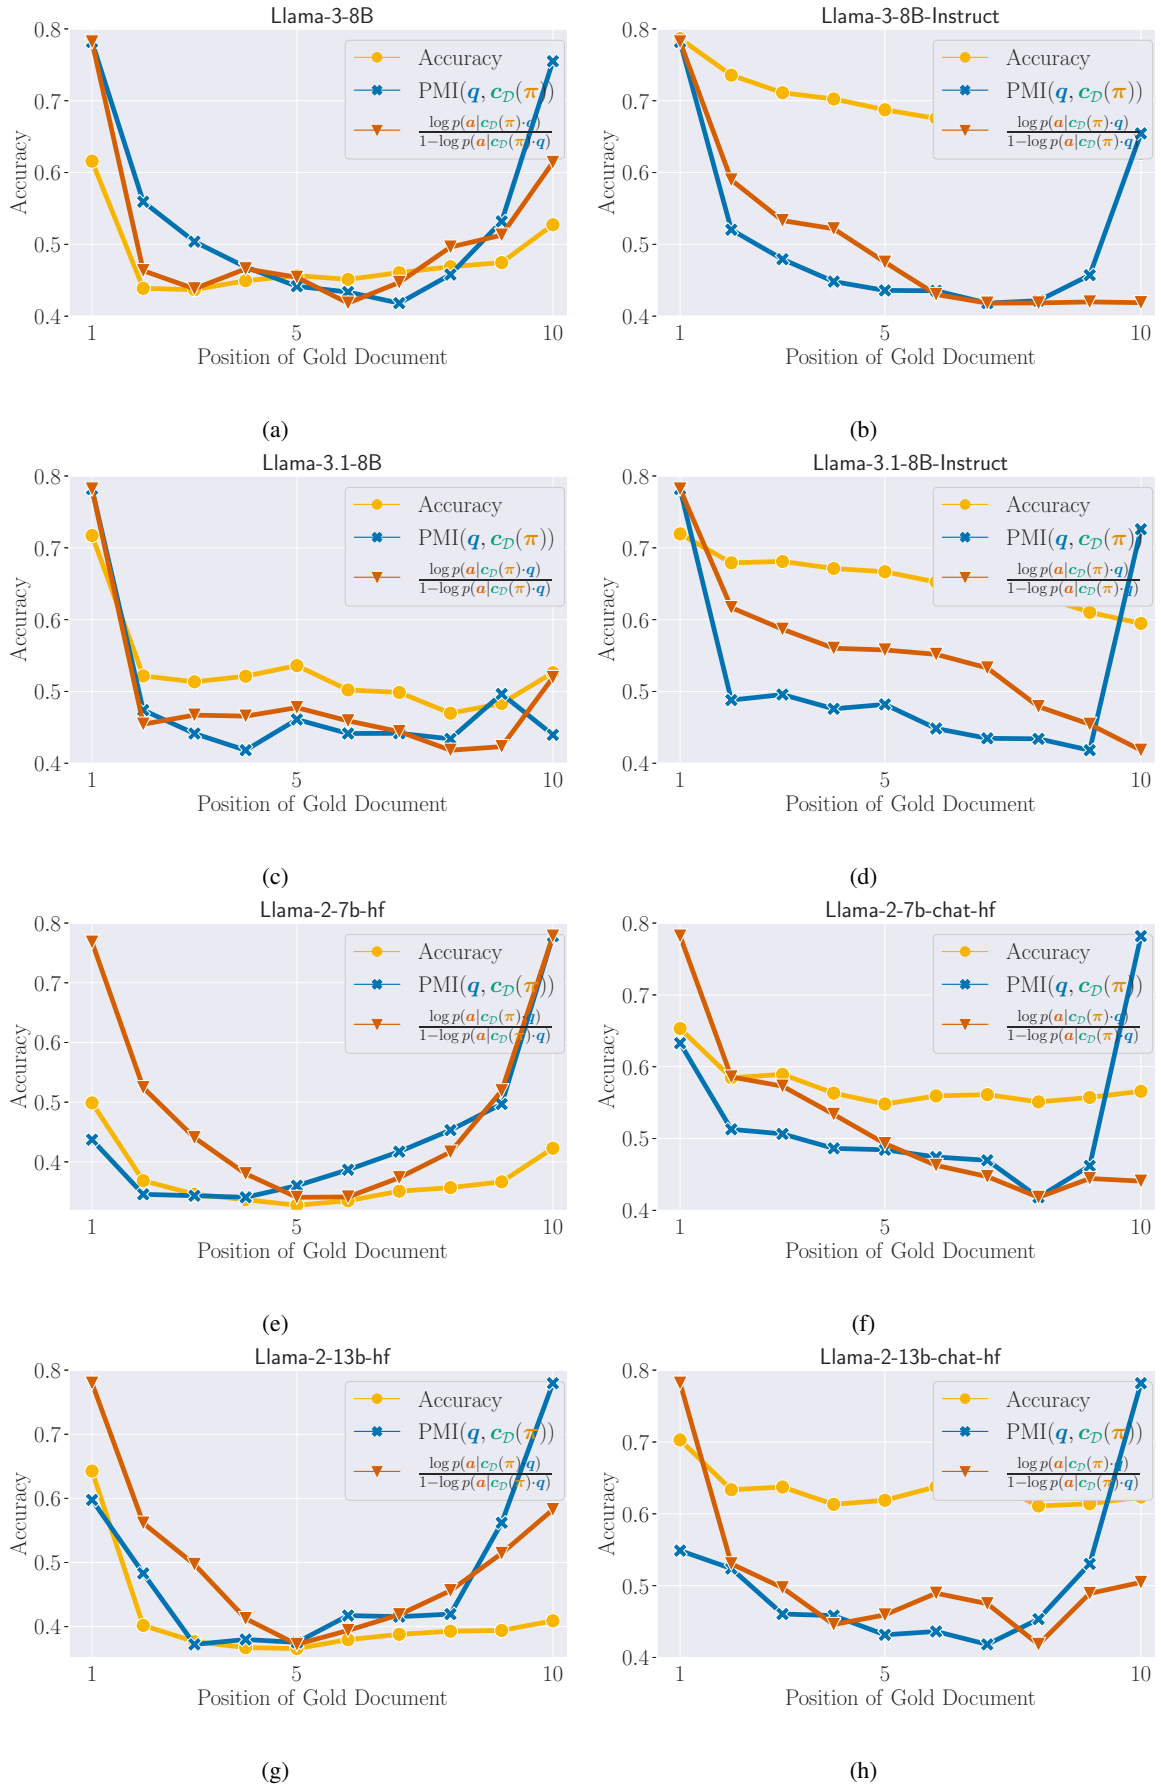


Figure 10: QA accuracy, PMI, and log odds ratio of answer likelihood on 10 documents.



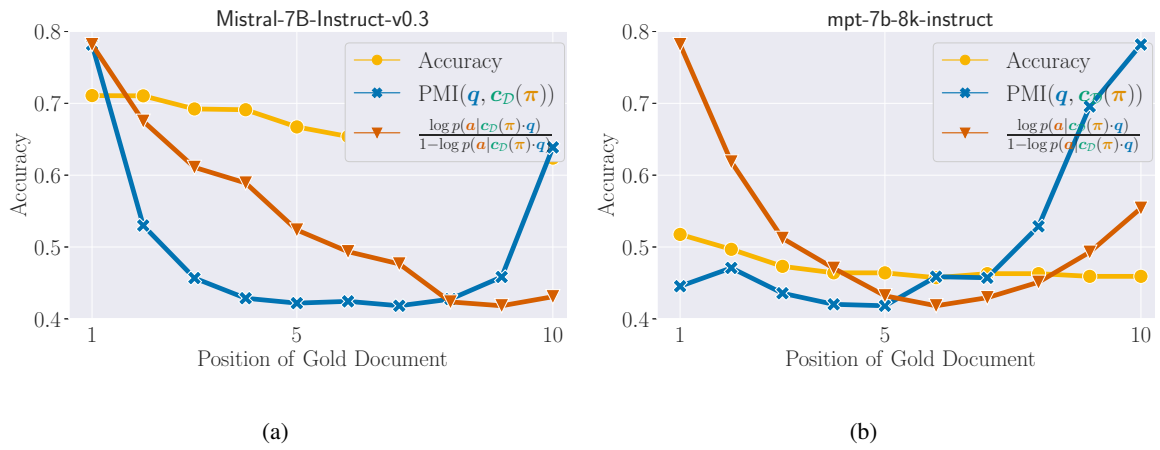


Figure 11: QA accuracy, PMI, and log odds ratio of answer likelihood on 10 documents.

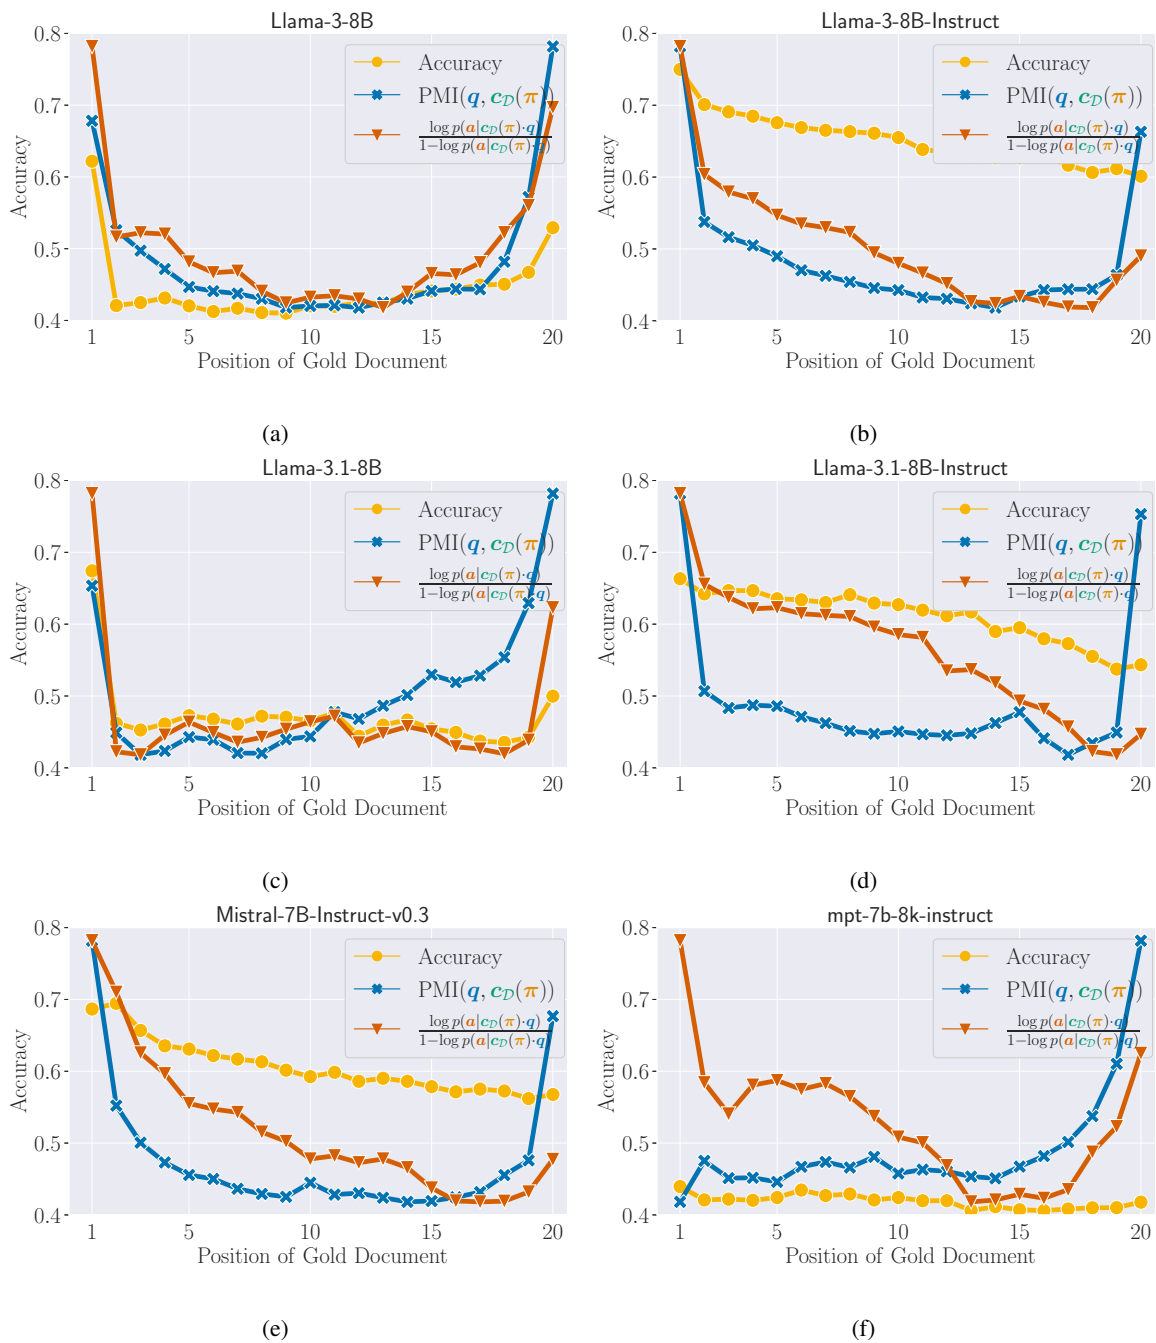


Figure 12: QA accuracy, PMI, and log odds ratio of answer likelihood on 20 documents.

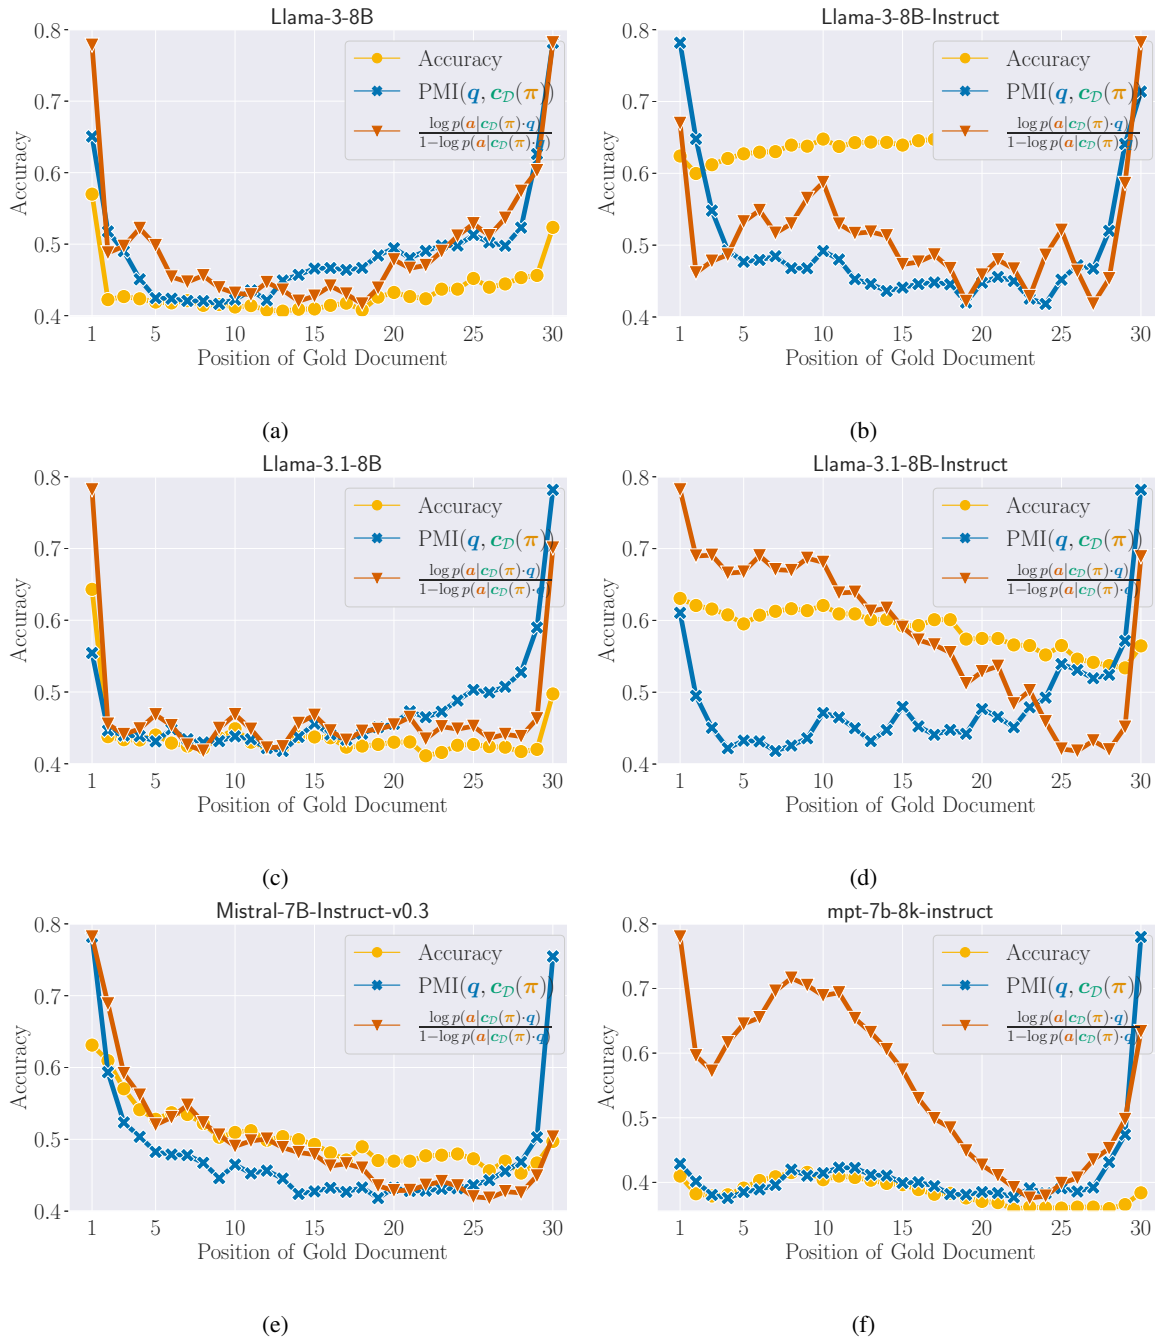


Figure 13: QA accuracy, PMI, and log odds ratio of answer likelihood on 30 documents.

```

{
  "ebad6435-1e86-4b9e-836a-9a88a8c93743":
    "c13ac8fc-81fe-408a-bf8f-914b6b8dc310",
  "33e652a0-fbcd-4abd-9935-14043ef82de9":
    "339ffb66-ec38-4d2a-a99f-67755d87eec3",
  "7a990232-7ddd-41b6-a8eb-1c61dc96da3c":
    "0d233f17-9d85-441e-868c-aa682d3dbbe7",
  ...
}
Key: "7a990232-7ddd-41b6-a8eb-1c61dc96da3c"
Value: "0d233f17-9d85-441e-868c-aa682d3dbbe7"

```

Figure 14: Example input for key-value retrieval task.

## G Proof of Proposition 2.1

**Proposition 2.1.** *Under assumptions given in Assumption 2.1, we have*

$$\begin{aligned} \log \frac{\vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))}{1 - \vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))} & \\ &= \text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi)) + C(\mathbf{a}, \mathbf{c}_{\mathcal{D}}(\pi)) \end{aligned} \quad (6)$$

for an *answer-dependent constant*  $C(\mathbf{a}, \mathbf{c}_{\mathcal{D}}(\pi))$ .

*Proof.* First note that, by Bayes' rule, we have

$$\vec{p}(\mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi) \cdot \mathbf{q}) = \frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi) \square \mathbf{a})}{\vec{p}(\mathbf{c}_{\mathcal{D}}(\pi) \cdot \mathbf{q})}. \quad (17)$$

Then,

$$\log \frac{\vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))}{1 - \vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))} = \log \frac{\vec{p}(\mathbf{a} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))}{\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \vec{p}(\bar{\mathbf{a}} | \mathbf{q} \cdot \mathbf{c}_{\mathcal{D}}(\pi))} \quad (18a)$$

$$= \log \frac{\frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi) \square \mathbf{a})}{\vec{p}(\mathbf{c}_{\mathcal{D}}(\pi) \cdot \mathbf{q})}}{\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \frac{\vec{p}(\square \bar{\mathbf{a}} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi) \square \bar{\mathbf{a}})}{\vec{p}(\mathbf{c}_{\mathcal{D}}(\pi) \cdot \mathbf{q})}} \quad (\text{Bayes' rule}) \quad (18b)$$

$$= \log \frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi) \square \mathbf{a})}{\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \vec{p}(\square \bar{\mathbf{a}} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi) \square \bar{\mathbf{a}})} \quad (18c)$$

$$= \log \frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi)) \vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi))}{(\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \vec{p}(\square \bar{\mathbf{a}} | \mathbf{c}_{\mathcal{D}}(\pi))) \vec{p}(\mathbf{q})} \quad (\text{Assumption 2.1}) \quad (18d)$$

$$= \log \frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi))}{\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \vec{p}(\square \bar{\mathbf{a}} | \mathbf{c}_{\mathcal{D}}(\pi))} \frac{\vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi))}{\vec{p}(\mathbf{q})} \quad (18e)$$

$$= \log \underbrace{\frac{\vec{p}(\square \mathbf{a} | \mathbf{c}_{\mathcal{D}}(\pi))}{\sum_{\bar{\mathbf{a}} \in \Sigma^*} \mathbb{1}\{\bar{\mathbf{a}} \neq \mathbf{a}\} \vec{p}(\square \bar{\mathbf{a}} | \mathbf{c}_{\mathcal{D}}(\pi))}}_{\triangleq C(\mathbf{a}, \mathbf{c}_{\mathcal{D}}(\pi))} + \log \frac{\vec{p}(\mathbf{q} | \mathbf{c}_{\mathcal{D}}(\pi))}{\vec{p}(\mathbf{q})} \quad (18f)$$

$$= \text{PMI}(\mathbf{q}, \mathbf{c}_{\mathcal{D}}(\pi)) + C(\mathbf{a}, \mathbf{c}_{\mathcal{D}}(\pi)), \quad (18g)$$

where  $C(\mathbf{a}, \mathbf{c}_{\mathcal{D}}(\pi))$  is constant with respect to  $\mathbf{q}$ . ■