

MODEL TELLS ITSELF WHERE TO ATTEND: STEER-ABLE PROMPTING FOR RELIABLE READING COMPREHENSION OF LLMs

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) have demonstrated remarkable performance across various real-world tasks. However, they often struggle to fully comprehend and effectively utilize their input contexts, resulting in responses that are hallucinated. This difficulty increases for contexts that are long or contain distracting information, which can divert LLMs from fully capturing essential evidence. To address this issue, many works use prompting to help LLMs comprehend contextual information more reliably. For instance, iterative prompting highlights key information in two steps that first ask the LLM to identify important pieces of context and then derive answers accordingly. However, textual prompting methods are constrained to highlighting key information implicitly in token space, which is often insufficient to fully steer the model’s attention. To improve model reading comprehension, we propose SteerPrompt, a method that automatically identifies key contextual information and explicitly highlights it by steering an LLM’s attention scores. Like prompting, SteerPrompt is applied at inference time and does not require changing any model parameters. Our experiments on open-book QA demonstrate that SteerPrompt effectively enables models to grasp essential contextual information, leading to substantially improved problem-solving performance, e.g., an average improvement of 7.95% for LLAMA3-70B-Instruct. Code will be publicly available.

1 INTRODUCTION

Large language models (LLMs) exhibit remarkable performance across various natural language processing (NLP) tasks and artificial intelligence (AI) applications (Brown et al., 2020; Touvron et al., 2023; OpenAI, 2023). Despite their remarkable capabilities, recent studies reveal that LLMs often encounter challenges in fully understanding their input contexts, overlooking or showing insensitivity to crucial contextual information (Kasai et al., 2023; Li et al., 2023; Si et al., 2023; Zhou et al., 2023; Yu et al., 2024; Zhang et al., 2024). Consequently, the models tend to fabricate answers (also known as hallucination), resulting in responses that are inconsistent with the presented contexts (Zhou et al., 2023; Yu et al., 2024). This becomes particularly problematic when models are presented prompts containing lengthy background contexts (Liu et al., 2023) or complex questions, such as in open-book question answering (QA) (Kwiatkowski et al., 2019; Shi et al., 2023b; Peng et al., 2023). In these information-dense scenarios, lengthy contexts can overwhelm LLMs, which contain many details with varying degree of relevance (Wan et al., 2024; Zhang et al., 2024). Some sentences are crucial for providing the correct answer, while others, though irrelevant, can distract models from fully understanding the essential information.

To improve reading comprehension of models, most prior work explores well-designed prompts to guide the LLM to use contextual knowledge more reliably (Zhou et al., 2023; Wan et al., 2024; Radhakrishnan et al., 2023). In particular, *iterative prompting* in chain-of-thought (COT; Wei et al., 2022) fashion can help LLMs decompose complex task-solving into more interpretable and manageable intermediate steps, thus yielding better performance (Radhakrishnan et al., 2023). Motivated by this, it is natural to design multi-step iterative prompting to guide LLMs to pay more attention to relevant contextual parts and derive answers accordingly. Specifically, for open-book QA tasks, iterative prompting can be decomposed into two steps: (i) *identifying key information* and (ii) *deriving answers using the key information*.

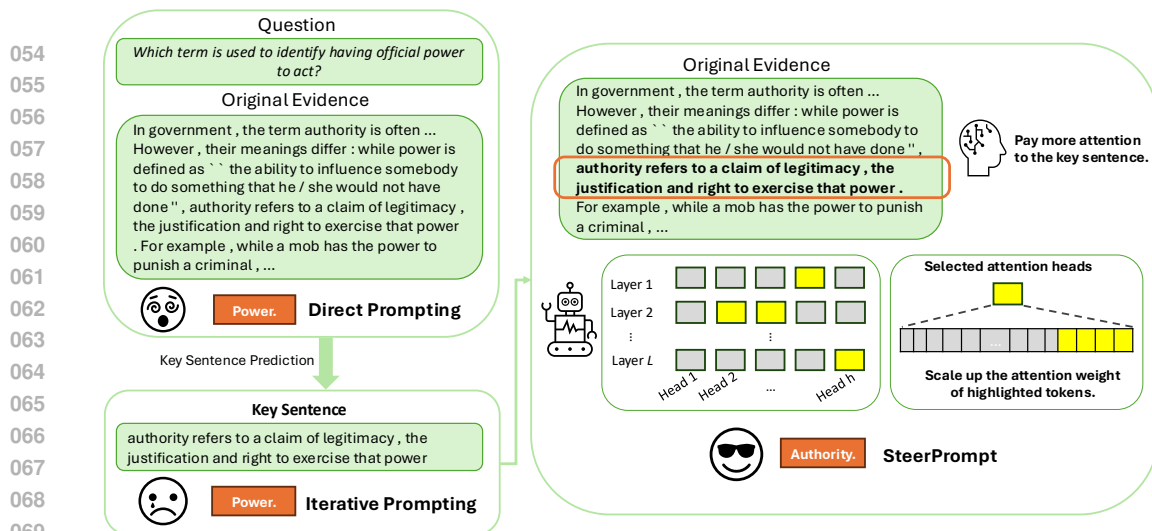


Figure 1: The illustration of SteerPrompt and alternative methods given a running example. Responses by Vicuna-7B are shown in red square where *Authority* is the label. Prompting methods (direct and iterative prompting) fail to guide a model to derive correct answers while SteerPrompt successfully steers it to answer correctly by explicitly highlighting identified key parts.

Iterative prompting can work effectively for black-box LLMs of significantly large sizes (e.g., >100B) (Radhakrishnan et al., 2023). However, for LLMs of smaller sizes (e.g., LLAMA3-70B, Meta, 2024), it remains unclear if this strategy can guide models to fully attend to the extracted key information and subsequently improve performance. First, step-by-step generations typically result in longer contexts. However, key information is only highlighted in token space by appending the short predicted key sentences, which are often not strong enough to fully steer the model’s attention. As illustrated in the left part of Figure 1, even though the model correctly predicts the key sentence which is appended in the subsequent prompt, it still fails to provide the correct answer. Moreover, errors can propagate across steps, further compromising performance. Therefore, we aim to develop an alternative inference method that emulates iterative prompting while addressing these limitations.

Motivated by this, we propose SteerPrompt, an inference-only method that (i) *automatically* identifies key contextual parts, and (ii) *explicitly* highlights them through attention score manipulation for improving models’ reading comprehension and performance on open-book QA tasks. Specifically, SteerPrompt integrates iterative question-decomposition prompting and attention steering approaches (Zhang et al., 2024). Given the original context and question, SteerPrompt first prompts an LLM to identify the key information (sentences) through free-text generation. Then, instead of appending those key sentences to the initial prompt, SteerPrompt maps those sentences back to the original context using semantic embeddings (Figure 1 right). By using the original sentences for highlighting, it avoids more lengthy input for the next step, and potentially reduces the unreliable key sentences generations, mitigating the error propagation. Finally, to guide the model to attend to the selected key sentences, SteerPrompt highlights them through attention steering that upweights their corresponding attention scores at the selected attention heads as done by Zhang et al. (2024). Unlike existing attention steering work, our method does not necessitate human annotation on the highlighting part, rectifying its critical limitation. Additionally, we also design an efficient coarse-to-fine search scheme for identifying effective attention heads for steering, which reduces the searching overhead by $4.5\times$ compared to the greedy method of previous work (Zhang et al., 2024).

We conduct experiments to evaluate the effectiveness of SteerPrompt using Vicuna-7B (Chiang et al., 2023), LLAMA3-8B-Instruct, and LLAMA3-70B-Instruct (Meta, 2024) on both single- and multi-hop open-book QA tasks from Natural Questions (NQ; Kwiatkowski et al., 2019) and HotpotQA (Yang et al., 2018b). Empirical evidence shows that it can be a much easier task for the LLMs to select the contextual key sentences than comprehending the full context (see analysis in Section 5.1). SteerPrompt converts the challenging bottleneck of contextual comprehension into an easier problem of key-sentence selection in a steerable way, consistently providing significant performance improvements over baseline prompting strategies. For example, SteerPrompt achieves an average improvement of 8.99% on exact-match (EM) score over iterative prompting for LLAMA3-

70B-Instruct across both tasks. Remarkably, the attention head sets obtained by SteerPrompt exhibit outstanding generalization ability, allowing them to be effectively used across different tasks.

2 BACKGROUND

Problem description. In standard LLM prompting, we are given a pre-trained LLM and a text prompt \mathbf{x} consisting of n tokens. In the closed-book setting, the prompt \mathbf{x} can only be a question or instruction to be completed by models. Relying solely on model parametric knowledge poses challenges in scenarios involving complex questions that entail new knowledge or private information (Zhou et al., 2023; Yu et al., 2024). Existing methods (Shi et al., 2023b; Peng et al., 2023) resort to augmenting the prompt with additional background contexts to facilitate question answering, i.e., open-book question answering. The following box presents a prompt template that we use for open-book QA:

A direct prompt template for open-book QA

Answer the question below, paired with a context that provides background knowledge. Only output the answer without other context words.

Context: {context}

Question: {question}

Answer:

Multi-head attention. A typical transformer model consists of L stacked layers, where each layer contains two submodules: a multi-head attention (MHA) and a fully connected feed-forward network (FFN). Given the input $\mathbf{X} \in \mathbb{R}^{n \times d}$, MHA of the layer l performs the attention function in parallel H heads: $\text{MHA}^{(l)}(\mathbf{X}) = \text{Concat}(\mathbf{H}^{(l,1)}, \dots, \mathbf{H}^{(l,H)})\mathbf{W}_o$ with

$$\mathbf{H}^{(l,h)} = \text{Softmax}(\mathbf{A}^{(l,h)})\mathbf{V}^{(l,h)}$$

where $\mathbf{A} = \frac{1}{\sqrt{d_n}}\mathbf{Q}\mathbf{K}^\top \in \mathbb{R}^{n \times n}$ is the scaled inner product between query \mathbf{Q} and key \mathbf{K} . $\mathbf{Q} = \mathbf{X}\mathbf{W}_{q_n}$, $\mathbf{K} = \mathbf{X}\mathbf{W}_{k_n}$, $\mathbf{V} = \mathbf{X}\mathbf{W}_{v_n}$ and $\mathbf{W}_{q_n}, \mathbf{W}_{k_n}, \mathbf{W}_{v_n} \in \mathbb{R}^{d \times d_n}$ are learnable projection matrices of head h . d_h is typically set to d/H .

Post-hoc attention steering. Zhang et al. (2024) propose PASTA, an inference-only method that applies attention reweighting to steer model attention towards user-highlighted input sets, thereby improving instruction following and contextual comprehension. Specifically, given the index set of user-specified tokens as \mathcal{G} ($\mathcal{G} \subset [n]$), PASTA highlight these tokens by upweighting their attention scores with a constant attention bias $\mathbf{B}^{(l,h)}$:

$$\mathbf{H}^{(l,h)} = \text{Softmax}(\mathbf{A}^{(l,h)} + \mathbf{B}^{(l,h)})\mathbf{V}^{(l,h)}, \quad \mathbf{B}_{ij}^{(l,h)} = \begin{cases} -\delta & \text{if } (l, h) \in \mathcal{H} \text{ and } j \notin \mathcal{G} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

where δ is a positive constant. After $\text{Softmax}(\cdot)$, the attention scores of tokens not in \mathcal{G} is scaled down by $\exp(\delta)$. Correspondingly, the others in \mathcal{G} are upweighted due to the renormalization of Softmax^1 , steering the model to pay more attention to the input spans of \mathcal{G} . Following Zhang et al. (2024), we set $\delta = \log 100$ in all of our experiments. \mathcal{H} is an index set of attention heads selected for steering. Since various heads function diversely, steering different heads yields dramatically different performance. To identify the effective heads, Zhang et al. (2024) employ a greedy search approach that evaluates the steering performance of each head on multiple tasks and selects those with best accuracy. The resulting head set \mathcal{H} can be generalized for steering across different tasks.

PASTA requires the access to user-annotated input spans for highlighting. In the case of context-specific tasks, it is generally prohibitively expensive to extract and annotate relevant sentences from lengthy contexts through humans. To address this critical limitation and improve model reading comprehension by automatic explicit highlighting, we introduce our method – SteerPrompt.

¹(1) is a simplified formula from Equation (2) in Zhang et al. (2024), which we elaborate in Appendix A.

3 METHOD

Our proposed method – Steerable Prompting (SteerPrompt), integrates iterative prompting and attention steering. This integration synergistically combines the advantages of both techniques while mitigating their respective limitations. For multi-step iterative prompting, incorporating attention steering externalizes the highlighting of key information through an inference-only operation, efficiently enhancing model reading comprehension with improved reliability and controllability. For post-hoc attention steering, equipping it with iterative prompting enables the automatic identification of contextually relevant key information, thereby addressing its significant reliance on human annotations.

Algorithm 1 SteerPrompt

Input a question q , a context c , the head set \mathcal{H} of an LLM \mathcal{M} , prompt templates $\mathcal{P}_i, \mathcal{P}_d$ and δ .

- 1: Generate $g_1 = \text{Generate}_{\mathcal{M}}(\mathcal{P}_i(q, c))$ as in (2);
- 2: Calculate $s_k = \text{Match}_e(g_1, \{s_1, \dots, s_m\})$ as in (3);
- 3: Steer $g_2 = \text{Steer}_{\mathcal{H}, s_k}(\text{Generate}_{\mathcal{M}}(\mathcal{P}_d(q, c)))$ as in (4);

Output: The final answer g_2

3.1 AUTOMATIC CONTEXTUAL HIGHLIGHTING

In the open-book QA task, an LLM \mathcal{M} is prompted to answer a question q paired with a background context c that consists of m sentences $c = s_1 || \dots || s_m$. Instead of directly prompting an LLM with (q, c) , SteerPrompt first prompts the LLM to generate a key sentence from the context c that supports answering the question:

$$g_1 = \text{Generate}_{\mathcal{M}}(\mathcal{P}_i(q, c)), \quad (2)$$

where \mathcal{P}_i is the prompt template of key sentence identification that we show in Section 4.1. Then, SteerPrompt maps g_1 back to a sentence from the original context c to avoid potential token-level generation errors in g_1 and mitigate error propagation. Specifically, it employs a small encoder e to calculate the semantic embeddings of g_1 and every $s_i (1 \leq i \leq m)$, and pick the best-matching sentence s_k with the highest similarity to g_1 :

$$s_k = \text{Match}_e(g_1, \{s_1, \dots, s_m\}) \subset c. \quad (3)$$

In the final step, SteerPrompt steers the attention scores of tokens in s_k based on (1) at the specific attention heads \mathcal{H} , when directly prompting the LLM \mathcal{M} to derive the answer for (q, c) :

$$g_2 = \text{Steer}_{\mathcal{H}, s_k}(\text{Generate}_{\mathcal{M}}(\mathcal{P}_d(q, c))) \quad (4)$$

where \mathcal{P}_d is the prompt template of direct answering as shown in Section 2, and $\text{Steer}_{\mathcal{H}, s_k}(\cdot)$ is detailed by (1) with \mathcal{G} as the index set of s_k . As such, the identified key sentence s_k is explicitly highlighted through attention score upweighting, directing the model to grasp the key information and solve the questions more reliably. Notably, SteerPrompt is applied at inference time and does not require changing any model parameters. More importantly, it does not involve human annotation on highlighted parts. The key information is automatically identified by iterative prompting the model \mathcal{M} , addressing the major limitation of existing attention steering approach.

3.2 COARSE-TO-FINE MODEL PROFILING

SteerPrompt requires carefully selecting \mathcal{H} , the set of attention heads to be steered in (1), but finding these heads can be computationally intensive. Zhang et al. (2024) propose a greedy search strategy that evaluates the steering performance of each head on small validation sets of multiple tasks and selects the heads that yield the best performance. This greedy strategy requires evaluating $L \times H$ times, resulting in non-trivial overheads especially for large models. To improve the efficiency of searching heads, we propose an alternative coarse-to-fine model profiling scheme that searches from the layer level to head level. Specifically, we first evaluate the performance of steering all attention heads of one single layer, then pick the top- l layers, and further evaluate the steering performance of each head in these layers. The head set \mathcal{H} is obtained by selecting the best-performing heads from top- l layers. Empirically, we find that a small l (e.g., $l = 6$ compared to $L = 32$) is sufficient for SteerPrompt to achieve superior performance and identify effective attention heads that can generalize across tasks, substantially reducing the searching overheads to $\frac{lH+L}{LH}$.

Table 1: Generation examples of a Vicuna-7B on NQ and HotpotQA. Texts in **blue** are predicted by the model for highlighting and texts in **bold** are highlighted by SteerPrompt.

Task	Prompt	Baseline	SteerPrompt
NQ	<p>Answer the question below, paired with a context that provides background knowledge. Only output the answer without other context words.</p> <p>Context: Although the delegates were divided early on as to whether to break from Crown rule, the second Continental Congress on July 2, 1776, passed a resolution asserting independence, with no opposing vote recorded. The Declaration of Independence was issued two days later declaring themselves a new nation: the United States of America. It established a Continental Army, giving command to one of its members, George Washington of Virginia. It waged war with Great Britain, made a militia treaty with France, and funded the war effort with loans and paper money.</p> <p>Question: when did the continental congress vote to adopt the declaration of independence?</p> <p>Answer:</p>	<p>✗ The Continental Congress voted to adopt the Declaration of Independence on July 4, 1776.</p>	<p>✓ July 2, 1776. Label: July 2, 1776.</p>
HotpotQA	<p>Answer the question below, paired with a context that provides background knowledge. Only output the answer without other context words.</p> <p>Context: [1]: Branford, Connecticut - Branford is a shoreline town located on Long Island Sound in New Haven County, Connecticut, 8 mi east of New Haven. The population was 28,026 at the 2010 census. [2]: Long Island Sound - Long Island Sound is a tidal estuary of the Atlantic Ocean, lying between the eastern shores of Bronx County, New York City, southern Westchester County, and Connecticut to the north, and the North Shore of Long Island, to the south. From east to west, the sound stretches 110 miles (177 km) from the East River in New York City, along the North Shore of Long Island, to Block Island Sound. A mix of freshwater from tributaries and saltwater from the ocean, Long Island Sound is 21 miles (34 km) at its widest point and varies in depth from 65 to.</p> <p>Question: How long is the tidal estuary in which Branford is a shoreline town?</p> <p>Answer:</p>	<p>✗ Long Island Sound.</p>	<p>✓ 110 miles. Label: 110 miles.</p>

4 EXPERIMENTS

We conduct experiments to evaluate the effectiveness of SteerPrompt using Vicuna-7B (Chiang et al., 2023), LLAMA3-8B-Instruct, and LLAMA3-70B-Instruct (Meta, 2024) on both single- and multi-hop open-book QA tasks from Natural Questions (NA; Kwiatkowski et al., 2019) and HotpotQA (Yang et al., 2018b).

4.1 EXPERIMENTAL SETUP

Datasets. We study 2 datasets: HotpotQA (Yang et al., 2018a) and the MRQA version Fisch et al. (2019) of Natural Questions (NQ) (Kwiatkowski et al., 2019). Following the filtering procedures outlined by Yu et al. (2024), duplicated and low-quality questions are removed from the NQ dataset, resulting in 7,189 instances remaining in NQ, and 5,190 instances in HotpotQA. For each dataset, we randomly select 1,000 examples as the profiling set and keep the remaining examples as the test set (see breakdown in Table 7). For all the experiments, we present two evaluation metrics: Exact Match (EM), and Token-level F1 score. We apply greedy search decoding for all experiments.

Implementation Details. We implement our experiments using Huggingface Transformers (Wolf et al., 2019) and PyTorch (Paszke et al., 2019). All the experiments are conducted on NVIDIA A6000 and A100 GPUs.

SteerPrompt Settings. For SteerPrompt, we use the following prompt template \mathcal{P}_i to prompt an LLM \mathcal{M} to identify the key information from the context that support answering the question.

Prompt template \mathcal{P}_i of key sentence identification

A question, and a passage are shown below. Please select the key sentence in the passage that supports to answer the question correctly. Only output the exactly same sentence from the passage without other additional words.

Question: {question}

Passage: {context}

Sentence:

Then, we map the predicted key sentence g_1 back to the original context by (3), which uses a small encoder models to calculate the semantic embeddings of the predicted key sentence g_1 and every

Table 2: Main evaluation results with Vicuna-7B, LLAMA-8B-Instruct, and LLAMA3-70B-Instruct on NQ and HotpotQA. "In-domain" means that the head set is selected based on the profiling set of the target task. "Out-of-domain" means that the head set is selected from the other dataset and the target task is unseen during the profiling.

Model	Method	NQ		HotpotQA		All
		EM	Token F1	EM	Token F1	Ave.
Vicuna-7B	Direct Prompting	5.88	32.21	18.11	38.77	23.74
	Iterative Prompting	4.36	31.48	14.04	34.99	21.22
	SteerPrompt _{out-of-domain generalize}	11.78	35.53	21.94	39.92	27.29
	SteerPrompt _{in-domain profiling}	15.41	38.59	29.54	47.51	32.76
LLAMA3-8B	Direct Prompting	31.62	54.19	42.58	63.30	47.92
	Iterative Prompting	32.62	55.30	42.10	63.07	48.27
	SteerPrompt _{out-of-domain generalize}	41.21	61.61	51.96	70.97	56.44
	SteerPrompt _{in-domain profiling}	33.51	56.91	58.08	75.41	55.98
LLAMA3-70B	Direct Prompting	30.00	54.18	42.63	64.26	47.77
	Iterative Prompting	30.57	54.92	42.86	65.15	48.38
	SteerPrompt _{out-of-domain generalize}	33.46	56.59	58.32	76.74	56.28
	SteerPrompt _{in-domain profiling}	40.51	63.32	50.90	70.59	56.33

sentence s_i in the context c . Specifically, we use a "all-MiniLM-L6-v2" model from Sentence-Transformer (Reimers & Gurevych, 2019) as the encoder to encode sentences. Then, we calculate the cosine similarity between semantic embeddings of g_1 and each sentence s_i in the context, and select the contextual sentence s_k with the highest similarity score as the final key sentence prediction. For multi-hop question answering, such as HotpotQA, the key sentences are identified for each individual hop separately. Finally, we highlight s_k by (4) while directly prompting the model to answer the question paired by the context with the direct prompting template shown in Section 2.

Coarse-to-fine Model Profiling. For the coarse-to-fine search strategy outlined in Section 3.2, we consider all attention heads in the top- l layers as potential candidates for selection, where l is chosen from $\{3, 4, 5, 6\}$. Subsequently, we either select top- i heads from each individual layer, or top- j heads from the pool of head candidates. Top- i is chosen from $\{4, 6, 8\}$, and top- j is chosen from $\{16, 24, 32, 64\}$. The final head set utilized in the study is determined based on the highest token-F1 performance achieved on the profiling set.

Baselines. We evaluate three open-source LLMs: Vicuna-7B Chiang et al. (2023), Llama3-8B-Instruct, and Llama3-70B-Instruct under direct prompting, iterative prompting, and SteerPrompt.

- *Direct prompting:* Models are prompted to directly answer the question q based on the provided context c . The prompt template \mathcal{P}_d is displayed in Section 2.

- *Iterative Prompting:* Models are first prompted to generate the key sentence that supports answering the question, using the same prompt template \mathcal{P}_i . For multi-hop question answering, such as HotpotQA, the key sentences are identified for each individual hop separately (see Appendix C). The predicted key sentences are also mapped back to the original context, similar as that in SteerPrompt. Then, the model are prompted to answer the question with the key sentences appended to the context.

4.2 MAIN RESULT: STEERPROMPT IMPROVES OPEN-BOOK QA.

We evaluate the performance of SteerPrompt on NQ and HotpotQA in two settings: in-domain and out-of-domain evaluation. For the in-domain setting, we perform profiling (selecting the head set to steer) and evaluate performance on the same dataset. For the out-of-domain setting, we perform profiling and evaluate performance on different datasets, where the target task is unseen during the profiling to evaluate the generalization ability of SteerPrompt.

In-domain Evaluation. Table 2 suggests that, for all the models, SteerPrompt significantly improves the model performance compared with other baselines, regardless of model size and datasets. For

example, SteerPrompt achieves 40.51% EM for LLAMA3-70B-Instruct on NQ, yielding a significant 9.94% improvement compared to the best-performing baseline. We also observe that iterative prompting can mostly improve upon the direct prompting, showcasing the performance gains from identifying key sentences and appending them to contexts. However, in certain cases, such as Vicuna-7B, iterative prompting can actually underperform direct prompting. It suggests that highlighting in token space by appending key sentences is insufficient to fully steer a model’s attention. In contrast, SteerPrompt shows a consistently substantial improvement over all baselines, demonstrating the effectiveness of automatic attention steering to improve model reading comprehension. Table 1 further illustrates this by comparing the generation examples of SteerPrompt and direct prompting.

Out-of-domain Evaluation. In this setting, given an evaluation task (e.g., NQ), we employ the head sets selected with profiling set of the other task (e.g., HotpotQA) for SteerPrompt to evaluate its generalization ability across different domains and tasks. The results in Table 2 indicate that SteerPrompt significantly outperforms all baseline methods for all models and all datasets, achieving better or comparable performance to that of in-domain profiling. Notably, for LLAMA3-8B-Instruct on NQ, the cross-domain performance surpasses the in-domain performance, compellingly demonstrating the robustness and generalization proficiency of our approach.

5 ANALYSIS

5.1 ISOLATING THE EFFECT OF STEERPROMPT’S TWO COMPONENTS

SteerPrompt consists of two primary components: automatic key sentence identification, and explicit highlighting key sentences. To underscore the necessity of both components, we conduct the comparison between following methods: (i) direct prompting with the original context; (ii) iterative prompting that appends the identified key sentences appended to the context; (iii) highlighting the entire context by attention steering approach but without key-sentence identification; (iv) SteerPrompt that highlights the identified key sentences. Moreover, we present additional results about iterative prompting to provide comprehensive evaluation for it. Specifically, within each context of the NQ dataset, there is one gold sentence that entails the answer. We also evaluate the performance of iterative prompting when appending the gold sentence.

The results in Table 3 indicate that SteerPrompt can benefit from using the identified key sentence, yielding significant performance gains. Specifically, highlighting the entire context via attention steering can improve upon direct prompting but underperforms SteerPrompt, suggesting the importance of key sentence identification. Meanwhile, the comparison between (ii) and (iv) illustrates the performance gains yielded by explicitly highlighting via attention steering. Therefore, these results suggest that both components are essential for SteerPrompt to achieve its best performance.

Table 3: Performance comparison among appending identified key sentences or gold sentence, and highlighting different parts of contexts. *N.A.* means that it cannot identify the gold sentence from HotpotQA that entails answers.

Method	LLAMA3-8B on HotpotQA		Vicuna-7B on NQ	
	EM	Token F1	EM	Token F1
Direct prompting	42.58	63.30	5.88	32.21
Iterative prompting w. identified key sentences	42.10	63.07	4.36	31.48
Iterative prompting w. gold sentences	<i>N.A.</i>	<i>N.A.</i>	5.01	33.67
Highlight the entire context	54.88	73.23	–	–
SteerPrompt highlights identified key sentences	58.08	75.41	15.41	38.59

Moreover, we can see that, even appending the fully correct gold sentences to prompts, iterative prompting still faces challenges in effectively improving model’s reading comprehension and QA performance. Therefore, the performance of prompting methods is bottlenecked by LLMs’ capability of comprehending full contexts and grasping key information from them. By contrast, SteerPrompt addresses this bottleneck by explicitly highlighting key information through attention steering. As such, the QA performance will be upper bounded by the accuracy of key sentence identification. Notably, it is a much easier task for LLMs to select the contextual key sentences than comprehending

the full context. LLMs can achieve much higher accuracy of selecting key sentences than directly generating correct answers. Moreover, SteerPrompt proposes to map the key sentence generated by the model to a sentence from the original context, further mitigating the error propagation. To showcase it, we evaluate the accuracy of SteerPrompt to identify the gold sentences on the NQ dataset with Vicuna-7B and LLAMA3-8B-Instruct.

Table 4: Accuracy of gold sentence identification by SteerPrompt

Vicuna-7B	LLAMA3-8B-Instruct
64.02%	67.07%

We can tell that SteerPrompt achieves much higher accuracy of key sentence selection compared to the models’ question answering accuracy (e.g., 64.02% of gold sentence identification v.s. 5.88% NQ EM by Vicuna-7B). As such, SteerPrompt converts the challenging bottleneck of contextual comprehension into an easier problem of key-sentence selection in a steerable way. The 64.02% accuracy of key sentence selection achieved by SteerPrompt is sufficient for it to significantly improve models’ QA performance as shown in Table 3.

5.2 COMPARISON BETWEEN PROFILING STRATEGIES

To illustrate the effectiveness of the coarse-to-fine profiling strategy introduced in Section 3.2, we evaluate several different profiling approaches as follows:

- Greedy search proposed by Zhang et al. (2024): This strategy involves selecting the top- k heads from all the attention heads in the models. The evaluation times for this strategy is $L \times H$.
- Group search inspired by Ainslie et al. (2023): Here, 8 adjacent heads from one layer form a group. Then, we evaluate them group-wise, and select the top- k head groups. The evaluate times for this strategy is $LH/8$.
- Coarse-to-fine search: This strategy initially selects the top- l layers and then chooses the head set only from the heads within these layers. The evaluation times for this strategy is $L + lH$.

where L is the number of layers, and H is the number of attention heads per layer. We compare them with a Vicuna-7B (Chiang et al., 2023) that has 32 layers, and 32 heads per layer. The results in Table 5 show that coarse-to-fine profiling significantly outperforms all the other strategies while reducing the total evaluation times by $4.5\times$ compared to the original greedy search in Zhang et al. (2024).

Table 5: Performance of SteerPrompt on NQ with Vicuna-7B when searching effective attention heads with different strategies. ”# Eval” refers to the total evaluations with the profiling set.

Method	# Eval	EM	Token F1
Baseline	N.A.	8.13	33.79
Greedy search all heads	1,024	14.81	35.63
Group search (size of 8)	128	12.12	36.13
Coarse-to-fine search	224	15.41	38.59

5.3 PERFORMANCE OF STEERPROMPT WHEN RETRIEVING SEVERAL PASSAGES

In this work, our primary objective is to enhance model reading comprehension to the provided evidence, and the gold evidence is always provided. Nonetheless, in practical applications, a retriever may simultaneously supply multiple similar and relevant passages. To demonstrate the effectiveness of SteerPrompt in such scenarios, we utilize DRP (Karpukhin et al., 2020) to retrieve an additional four passages, each ranked within the top four in relevance scores to the question. Along with the gold evidence, these five passages are then presented to the model. SteerPrompt is tasked with automatically extracting and highlighting the key sentence from this set. Table 6 displays the results using Vicuna-7B on NQ. Although the inclusion of more noisy passages generally leads to a performance decline, we still observe consistent improvements with SteerPrompt, underscoring the effectiveness of our approach.

Table 6: Performance of SteerPrompt on NQ with Vicuna-7B when additional passages are provided.

Model	Method	NQ with Gold Evidence		NQ with 5 Evidence	
		EM	Token F1	EM	Token F1
Vicuna-7B	Direct Prompting	5.88	32.21	2.50	22.18
	Iterative Prompting	4.36	31.48	1.74	21.99
	SteerPrompt _{in-domain profiling}	15.41	38.59	6.82	25.60




Figure 2: Ablation study of SteerPrompt performance on HotpotQA when steering different numbers of heads (2a and 2b) and setting different δ (2c). Dashed line in red refers to the baseline performance of direct prompting.

5.4 ABLATION STUDY

We conduct ablation study to discuss the performance of SteerPrompt given different number of attention heads for steering and different δ .

Varying the number of steered heads. Figure 2a presents the performance variation of SteerPrompt with Vicuna-7B on HotpotQA dataset when steering different number of attention heads. Figure 2b illustrates the EM results for LLAMA3-8B-Instruct on the HotpotQA dataset under similar conditions. We see that steering more heads for SteerPrompt may result in slight performance degeneration, for example, the performance of LLAMA3-8B-Instruct on HotpotQA. This observation is similar to findings in previous work (see Figure 3 in Zhang et al. (2024)), where overemphasizing too many heads can lead models to focus on solely on highlighted information while ignoring other parts, potentially degenerating performance. In practice, we recommend applying SteerPrompt to steer a moderate number of heads. The optimal number of steered heads in our study is determined based on the performance metrics on the profiling data.

Analyzing the sensitivity of δ . Figure 2c presents the sensitivity analysis for varying δ in (1) using LLAMA3-8B-Instruct on HotpotQA. We can see that the performance of SteerPrompt is not sensitive to the attention bias constant δ . Changing its logarithm values (i.e., the scaling-down coefficient for non-highlighted tokens as elaborated in Appendix A) from 50 to 3000 does not induce dramatic performance variation. Therefore, we set δ as its default value $\log(100)$, which is the same as Zhang et al. (2024).

6 RELATED WORK

Large language models exhibit remarkable performance on (context-free) knowledge-intensive tasks, such as open-domain question answering (QA) (Kwiatkowski et al., 2019) and commonsense reasoning (Mihaylov et al., 2018; Clark et al., 2018), indicating that they encode substantial knowledge about open-world facts (Zhou et al., 2023) in their parameters. Despite their proficiency in memorization, different kinds of hallucinations in the output are observed, including factual knowledge hallucination (Huang et al., 2023; Yu et al., 2024), hallucination in summarization (Maynez et al., 2020; Pagnoni et al., 2021), hallucination in logical operations (Lyu et al., 2023; Huang et al., 2023). In this work, we focus on the factual knowledge hallucination due to models’ unawareness of relevant knowledge or overlooking contextual information.

Retrieval-augmented LLMs. To address the problem of missing relevant knowledge, one popular method is to use retrieval-augmented LMs that supplement missing knowledge from external sources (Shi et al., 2023b; Peng et al., 2023). Retrieval augmentation requires that LLMs are sensitive to the input context and reliably understand the contexts. However, recent work shows that even if the relevant knowledge is presented, the model may still not fully comprehend the given evidence (Zhou et al., 2023; Yu et al., 2024; Wan et al., 2024).

Prompt-based strategies. To improve the reading comprehension of the models, various prompting strategies are designed to guide the model to detect the key information (Wei et al., 2022; Radhakrishnan et al., 2023), or focus on the given evidence (Zhou et al., 2023), while these extracted key information is only added as additional tokens in the input, and models still cannot grasp these new tokens.

Model-based strategies. Besides using prompting to improve the reading comprehension, some works augment an LLM’s training data (Köksal et al., 2023; Hu et al., 2024; Chen et al., 2024). Alternatively, Shi et al. (2023a) proposes a context-aware decoding method to downweight the output distribution associated with the model’s prior knowledge.

To the best of our knowledge, we are the first work to integrate key information prompting and explicit token highlighting during inference without any additional training.

7 CONCLUSION

In this paper, we address the challenge of contextual reading comprehension in open-book QA tasks and introduce SteerPrompt, an inference-only method that automatically identifies crucial information pieces within contexts and explicitly highlights them through steering a model’s attention scores. SteerPrompt guides the model to focus on the essential information within contexts, leading to substantially improved model reading comprehension and performance. Remarkably, by integrate iterative prompting and attention steering techniques, SteerPrompt synergistically combines their advantages while mitigating their respective limitations.

REFERENCES

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 4895–4901, 2023.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Yanda Chen, Chandan Singh, Xiaodong Liu, Simiao Zuo, Bin Yu, He He, and Jianfeng Gao. Towards consistent natural-language explanations via explanation-consistency finetuning. *arXiv preprint arXiv:2401.13986*, 2024.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. MRQA 2019 shared task: Evaluating generalization in reading comprehension. In *Proceedings of 2nd Machine Reading for Reading Comprehension (MRQA) Workshop at EMNLP*, 2019.
- Minda Hu, Bowei He, Yufei Wang, Liangyou Li, Chen Ma, and Irwin King. Mitigating large language model hallucination with faithful finetuning. *arXiv preprint arXiv:2406.11267*, 2024.

- 540 Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong
541 Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language
542 models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*,
543 2023.
- 544 Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi
545 Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. *arXiv*
546 *preprint arXiv:2004.04906*, 2020.
- 548 Jungo Kasai, Keisuke Sakaguchi, yoichi takahashi, Ronan Le Bras, Akari Asai, Xinyan Velocity Yu,
549 Dragomir Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. Realtime QA: What’s the answer
550 right now? In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and*
551 *Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=HfKOIPCvsv>.
- 552 Abdullatif Köksal, Renat Aksitov, and Chung-Ching Chang. Hallucination augmented recitations for
553 language models. *arXiv preprint arXiv:2311.07424*, 2023.
- 555 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris
556 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion
557 Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav
558 Petrov. Natural questions: A benchmark for question answering research. *Transactions of the*
559 *Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl.a.00276. URL
560 <https://aclanthology.org/Q19-1026>.
- 561 Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu,
562 and Sanjiv Kumar. Large language models with controllable working memory. In Anna Rogers,
563 Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational*
564 *Linguistics: ACL 2023*, pp. 1774–1793, Toronto, Canada, July 2023. Association for Computational
565 Linguistics. doi: 10.18653/v1/2023.findings-acl.112. URL [https://aclanthology.org/](https://aclanthology.org/2023.findings-acl.112)
566 [2023.findings-acl.112](https://aclanthology.org/2023.findings-acl.112).
- 567 Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and
568 Percy Liang. Lost in the middle: How language models use long contexts, 2023.
- 570 Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki,
571 and Chris Callison-Burch. Faithful chain-of-thought reasoning. *arXiv preprint arXiv:2301.13379*,
572 2023.
- 574 Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality
575 in abstractive summarization. *arXiv preprint arXiv:2005.00661*, 2020.
- 576 Meta. Introducing meta llama 3: The most capable openly available llm to date, 2024. URL
577 <https://ai.meta.com/blog/meta-llama-3/>.
- 579 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct
580 electricity? a new dataset for open book question answering, 2018.
- 581 OpenAI. Gpt-4 technical report, 2023.
- 582 Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. Understanding factuality in abstractive
583 summarization with frank: A benchmark for factuality metrics. *arXiv preprint arXiv:2104.13346*,
584 2021.
- 585 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
586 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward
587 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
588 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep
589 learning library. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-
590 Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems*
591 *32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December*
592 *8-14, 2019, Vancouver, BC, Canada*, pp. 8024–8035, 2019.

- 594 Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars
595 Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. Check your facts and try again: Improving large
596 language models with external knowledge and automated feedback, 2023.
597
- 598 Ansh Radhakrishnan, Karina Nguyen, Anna Chen, Carol Chen, Carson Denison, Danny Hernandez,
599 Esin Durmus, Evan Hubinger, Jackson Kernion, Kamilė Lukošiuė, Newton Cheng, Nicholas
600 Joseph, Nicholas Schiefer, Oliver Rausch, Sam McCandlish, Sheer El Showk, Tamera Lanham, Tim
601 Maxwell, Venkatesa Chandrasekaran, Zac Hatfield-Dodds, Jared Kaplan, Jan Brauner, Samuel R.
602 Bowman, and Ethan Perez. Question decomposition improves the faithfulness of model-generated
603 reasoning, 2023.
- 604 Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks.
605 In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
606 Association for Computational Linguistics, 11 2019. URL [https://arxiv.org/abs/1908.](https://arxiv.org/abs/1908.10084)
607 [10084](https://arxiv.org/abs/1908.10084).
- 608 Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau
609 Yih. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint*
610 *arXiv:2305.14739*, 2023a.
- 612 Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettle-
613 moyer, and Wen tau Yih. Replug: Retrieval-augmented black-box language models, 2023b.
- 614 Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber,
615 and Lijuan Wang. Prompting GPT-3 to be reliable. In *The Eleventh International Confer-*
616 *ence on Learning Representations*, 2023. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=98p5x51L5af)
617 [98p5x51L5af](https://openreview.net/forum?id=98p5x51L5af).
- 619 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
620 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation
621 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 622 Alexander Wan, Eric Wallace, and Dan Klein. What evidence do language models find convincing?
623 *arXiv preprint arXiv:2402.11782*, 2024.
624
- 625 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
626 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
627 *Neural Information Processing Systems*, 35:24824–24837, 2022.
- 628 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
629 Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers:
630 State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
631
- 632 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and
633 Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answer-
634 ing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*,
635 pp. 2369–2380, Brussels, Belgium, October–November 2018a. Association for Computational
636 Linguistics. doi: 10.18653/v1/D18-1259. URL <https://aclanthology.org/D18-1259>.
- 637 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov,
638 and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question
639 answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*,
640 2018b.
- 641 Xiaodong Yu, Hao Cheng, Xiaodong Liu, Dan Roth, and Jianfeng Gao. Reeval: Automatic hallucina-
642 tion evaluation for retrieval-augmented large language models via transferable adversarial attacks,
643 2024.
644
- 645 Qingru Zhang, Chandan Singh, Liyuan Liu, Xiaodong Liu, Bin Yu, Jianfeng Gao, and Tuo Zhao. Tell
646 your model where to attend: Post-hoc attention steering for LLMs. In *The Twelfth International*
647 *Conference on Learning Representations*, 2024. URL [https://openreview.net/forum?](https://openreview.net/forum?id=xZDW00oejD)
[id=xZDW00oejD](https://openreview.net/forum?id=xZDW00oejD).

648 Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. Context-faithful prompting for
649 large language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the*
650 *Association for Computational Linguistics: EMNLP 2023*, pp. 14544–14556, Singapore, December
651 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.968.
652 URL <https://aclanthology.org/2023.findings-emnlp.968>.
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 A DERIVATION FOR EQUATION 1

703
704
705 In this section, we present the derivation to show why (1) is equivalent to equation (2) in Zhang et al.
706 (2024).

707 For the token that are not highlighted $j \notin \mathcal{G}$, Zhang et al. (2024) downweight their attention scores by
708 scaling down their scores post-softmax by a coefficient α ($0 \leq \alpha \leq 1$): $\alpha \cdot \text{Softmax}(\mathbf{A}_{i\cdot})_j / C_i$ where
709 $C_i = \sum_{j \in \mathcal{G}} \text{Softmax}(\mathbf{A}_{i\cdot})_j + \sum_{j \notin \mathcal{G}} \alpha \cdot \text{Softmax}(\mathbf{A}_{i\cdot})_j$. Now we show that:
710

$$711 \alpha \cdot \text{Softmax}(\mathbf{A}_{i\cdot})_j / C_i = \frac{\alpha \exp(\mathbf{A}_{ij})}{C_i \sum_{j'} \exp(\mathbf{A}_{ij'})} \quad (5)$$

$$712 = \frac{\exp(\mathbf{A}_{ij} + \log(\alpha))}{C_i \sum_{j'} \exp(\mathbf{A}_{ij'})} \quad (6)$$

713
714
715
716
717 For the tokens in \mathcal{G} :

$$718 \text{Softmax}(\mathbf{A}_{i\cdot})_j / C_i = \frac{\exp(\mathbf{A}_{ij})}{C_i \sum_{j'} \exp(\mathbf{A}_{ij'})} \quad (7)$$

719
720
721 Therefore, after the renormalization, it is equivalent to condut the softmax among $\mathbf{A}_{ij} + \log(\alpha)$ for
722 $j \notin \mathcal{G}$ and \mathbf{A}_{ij} for $j \in \mathcal{G}$, which is our simplified equation in (1).
723
724

725 B EVALUATION DETAILS

726 B.1 DATASET STATISTICS

	Profiling	Test
Natural Questions	1,000	6,189
HotpotQA	1,000	4,190

727
728
729
730
731
732
733 Table 7: Natural Questions and HotpotQA data statistics after the preprocessing.
734
735
736

737 B.2 THE DETAILED NUMBER OF ATTENTION HEADS FOR STEERING

Model	NQ	HotpotQA
Vicuna-7B	top 64 heads from top 4 layers	top 96 heads from top 6 layers
LLAMA3-8B	top24 heads, 4 from each of top 6 layers	top24 heads, 4 from each of top 6 layers
LLAMA3-70B	top20 heads, 4 from each of top 5 layers	top 64 heads from top 5 layers

738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755 Table 8: The detailed number of attention heads for steering

C PROMPT TEMPLATE DETAILS

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

Prompt Templates of Two-Round Iterative Prompting

[First Round]: A question, and a passage are shown below. Please select the key sentence in the passage that supports to answer the question correctly. Only output the exactly same sentence from the passage without other additional words.

Question: {Question}

Passage: {Evidence}

Sentence: _____

[Second Round]: Answer the question below, paired with a context that provides background knowledge, and a key sentence. Only output the answer without other context words.

Context: {Evidence}

Key Sentence: {Predicted key sentence}

Question: {Question}

Answer: