

Position Engineering: Boosting Large Language Models through Positional Information Manipulation

Zhiyuan He, Huiqiang Jiang, Zilong Wang, Yuqing Yang, Luna Qiu, Lili Qiu

Microsoft Research

{zhiyuhe,hjiang,wangzilong,yuqyang,lunaqiu,liliqiu}@microsoft.com

Abstract

The performance of large language models (LLMs) is significantly influenced by the quality of the prompts provided. In response, researchers have developed enormous prompt engineering strategies aimed at modifying the prompt text to enhance task performance. In this paper, we introduce a novel technique termed *position engineering*, which offers a more efficient way to guide large language models. Unlike prompt engineering, which requires substantial effort to modify the text provided to LLMs, position engineering merely involves altering the positional information in the prompt without modifying the text itself. We have evaluated position engineering in two widely-used LLM scenarios: retrieval-augmented generation (RAG) and in-context learning (ICL). Our findings show that position engineering substantially improves upon the baseline in both cases. Position engineering thus represents a promising new strategy for exploiting the capabilities of large language models.

1 Introduction


Recent advancements in Large Language Models (LLMs) have demonstrated significant strides towards achieving artificial general intelligence. These models exhibit a wide range of capabilities, such as in-context learning (Brown et al., 2020), answering questions based on documents (Lewis et al., 2020; Guu et al., 2020), solving complex mathematical problems (Frieder et al., 2024), and generating code (Romera-Paredes et al., 2024; Ma et al., 2023).

When utilizing LLMs, user prompts are inputted, converted into sequences of tokens, and then processed through multiple attention layers (Vaswani et al., 2017). These attention layers employ two types of information derived from the token sequences: (i) Semantic information, where the tokens are converted into text embeddings, and (ii) Positional information, where the indices of the

tokens are converted into positional embeddings (Vaswani et al., 2017; Su et al., 2024). The attention mechanism then combines the semantic and positional information to predict the distribution of the next token in the sequence.



(a) Prompt engineering



(b) Position engineering

Figure 1: Comparison of prompt engineering and position engineering. "Para" refers to paragraphs, and "Sent" to sentences in prompts. Prompt engineering involves either adding, replacing, or removing paragraphs and sentences from prompts. In contrast, the proposed position engineering maintains the original prompt text but incorporates placeholder tokens instead. These placeholders are not involved in the computation of attention scores, thus the computation overhead is not increased. However, they do hold position indices, thereby affecting the position information of other tokens in the text.

Extensive research has been conducted on modifying prompt text to alter semantic information, aiming to boost task performances. For instance, few-shot prompting is introduced, enabling LLMs to learn new tasks in an in-context manner (Brown et al., 2020). Moreover, the Chain-of-Thought methodology has been introduced to enhance LLMs' reasoning abilities by prompting them to produce intermediate tokens (Wei et al., 2022; Kojima et al., 2022). Additionally, Automatic Prompt Engineer has been developed to autonomously design the prompting text for better task-specific performance (Zhou et al., 2022).

In this study, we investigate the potential of improving performance by solely modifying po-

sitional information, without any semantic information change. For the first time, we reveal that downstream task performance can be significantly enhanced by simply adjusting the positional indices of tokens, without modifying the text itself.

As illustrated in Figure 1, our approach involves the introduction of placeholder tokens to modify positional information. These placeholder tokens do not contribute to the computation of attention scores; however, they do occupy token indices. Consequently, the relative position of other tokens is altered, which could optimize the attention weights among different segments within the prompts. We refer to this approach as *position engineering*, highlighting the exclusive focus on manipulating positional information.

We propose a simple yet effective method based on brutal force to discover the optimal placeholder token number for each downstream task, and experiment it within two prevalent scenarios of LLMs: Retrieval-Augmented Generation (RAG) and In-Context Learning (ICL). Our method significantly enhances performance in both tasks, achieving up to a 15.4% absolute increase in accuracy for RAG and a 3.6% absolute increase for ICL. We also discover that the same placeholder number can consistently improve the RAG’s performance for different datasets and models.

In all, our contributions can be summarized as follows:

- For the first time, we discover that different downstream tasks’ performances can be improved by merely changing the positional information in prompts.
- We propose a method to help find a better positional information setting.
- We demonstrate that RAG performance can be consistently improved by a universal positional information setting on different datasets and models.

2 Methodology

2.1 Preliminary

In this section, we provide a brief overview of how large language models (LLMs) integrate position information. Let $\{t_i\}_{i=1}^N$ represent the input tokens to language models, and let $\{e_i\}_{i=1}^N$ denote the corresponding token embeddings. Initially, the

attention layer computes $\mathbf{q}, \mathbf{k}, \mathbf{v}$:

$$\begin{aligned} \mathbf{q}_m &= f_q(\mathbf{e}_m, m) \\ \mathbf{k}_n &= f_k(\mathbf{e}_n, n) \\ \mathbf{v}_n &= f_v(\mathbf{e}_n, n) \end{aligned} \quad (1)$$

where m and n are the position indices of tokens. The self-attention is then calculated as follows:

$$\begin{aligned} a_{m,n} &= \frac{e^{\frac{\mathbf{q}_m^\top \mathbf{k}_n}{\sqrt{d}}}}{\sum_{j=1}^N e^{\frac{\mathbf{q}_m^\top \mathbf{k}_j}{\sqrt{d}}}} \\ \mathbf{o}_m &= \sum_{n=1}^N a_{m,n} \mathbf{v}_n \end{aligned} \quad (2)$$

where $a_{m,n}$ is a scalar capturing the attention score between m -th token in the query and n -th token in the value and key sets. d denotes the dimension of the attention layer, and \mathbf{o}_m indicates the output for the m -th query token.

Absolute positioning is initially introduced by incorporating a positional embedding vector \mathbf{p}_n , which is related to m and n (Vaswani et al., 2017):

$$\begin{aligned} f_q(\mathbf{e}_m, m) &= \mathbf{W}_q(\mathbf{e}_m + \mathbf{p}_m) \\ f_k(\mathbf{e}_n, n) &= \mathbf{W}_k(\mathbf{e}_n + \mathbf{p}_n) \\ f_v(\mathbf{e}_n, n) &= \mathbf{W}_v(\mathbf{e}_n + \mathbf{p}_n) \end{aligned} \quad (3)$$

The $2i$ and $2i + 1$ dimension of the positional embedding \mathbf{p}_n is calculated as follows:

$$\begin{aligned} \mathbf{p}_{n,2i} &= \sin(n/10000^{\frac{2i}{d}}) \\ \mathbf{p}_{n,2i+1} &= \cos(n/10000^{\frac{2i}{d}}) \end{aligned} \quad (4)$$

Recently, RoPE adopts the relative position information instead of the absolute information (Su et al., 2024). It utilizes a specifically designed matrix \mathbf{R}_i^d , of dimensions $d \times d$ and parameterized by i , to modify the query and key vectors in the following manner:

$$\begin{aligned} f_q(\mathbf{e}_m, m) &= \mathbf{R}_m^d \mathbf{W}_q \mathbf{e}_m \\ f_k(\mathbf{e}_n, n) &= \mathbf{R}_n^d \mathbf{W}_k \mathbf{e}_n \\ f_v(\mathbf{e}_n, n) &= \mathbf{W}_v \mathbf{e}_n \end{aligned} \quad (5)$$

The matrix \mathbf{R}_i^d has a unique property, namely $(\mathbf{R}_i^d)^\top \mathbf{R}_j^d = \mathbf{R}_{j-i}^d$, which leads to:

$$\mathbf{q}_m^\top \mathbf{k}_n = \mathbf{e}_m^\top \mathbf{W}_q \mathbf{R}_{n-m}^d \mathbf{W}_k \mathbf{e}_n \quad (6)$$

Consequently, in Equation (2), the model solely focuses on the relative position $n - m$, instead of the absolute positions n and m . RoPE has been adopted by recent LLMs, including Llama, Llama2 and Mistral (Touvron et al., 2023a,b; Jiang et al., 2023).

2.2 Altering Position Information in Prompts

The performance of LLMs is significantly influenced by the quality of the prompts used. To enhance the effectiveness of these prompts, researchers have developed a wide range of prompt engineering strategies. This refinement process involves transforming the initial input tokens $\{t_i\}_{i=1}^N$ into revised inputs $\{\widehat{t}_j\}_{j=1}^{\widehat{N}}$, which necessitates modifications to the text. For instance, the Zero-shot chain-of-thought technique enhances the reasoning abilities of LLMs by appending the sentence "Let's think step by step." to the prompts (Kojima et al., 2022).

In this paper, we propose a novel methodology termed "position engineering" to further exploit the capabilities of LLMs. Unlike prompt engineering, position engineering requires no modification to the input tokens themselves. Instead, it solely modifies the position information utilized in Equation (1). Through empirical experiments, we have discovered that such adjustments to position information can significantly improve performance. Formally, we aim at discovering a position editing function, $\tau(\cdot) : \mathbb{N} \rightarrow \mathbb{N}$, that boosts LLM performance. This function changes the token position information, which is incorporated into the model as shown below:

$$\begin{aligned}\widehat{\mathbf{q}}_m &= f_q(\mathbf{e}_m, \tau(m)) \\ \widehat{\mathbf{k}}_n &= f_k(\mathbf{e}_n, \tau(n)) \\ \widehat{\mathbf{v}}_n &= f_v(\mathbf{e}_n, \tau(n))\end{aligned}\quad (7)$$

We impose a condition on τ that $\forall i > j, \tau(i) > \tau(j)$. This requirement ensures that: (1) No two distinct tokens are assigned the same new position index, and (2) The causality in language modeling remains intact, meaning only query vectors with a larger index can access the key and value vectors with an equal or smaller index.

The concept of position engineering can be also explained through placeholder tokens. Placeholder tokens are defined as tokens that are excluded the computation of attention scores, yet they are allocated position indices. To elaborate, when the calculation of $a_{m,n}$ is undertaken as described in the Equation (2), and either the m -th or n -th token is identified as a placeholder, the conventional computation is bypassed, and $a_{m,n}$ is set to 0. While placeholder tokens do not directly influence the attention scores at their positions, they do alter the position indices of other input tokens. As depicted in Figure 1b, the insertion of placeholder

tokens between sentences 1 and 2 affects the relative positional information between them, which in turn influences the calculation of attention scores between tokens of the two sentences. The connection between the position editing function and the placeholder tokens can be described as follows: Employing a position editing function τ translates to adding $\tau(i+1) - \tau(i) - 1$ placeholder tokens after the i -th token, and specifically, adding $\tau(0)$ placeholder tokens before the 0-th token.

2.3 Position Engineering

Consider a particular task defined by (Q, A) , for which a training set $\{(Q_i, A_i)\}_{i=1}^N$ has been sampled according to the task distribution Γ . We transform each question Q_i into its corresponding text prompt P_i . A large language model \mathcal{M} is utilized, which operates based on the prompt P_i , and its output is evaluated through a scoring function r , denoted as $r(\mathcal{M}, P_i)$. To potentially enhance the performance, a position editing function might be applied to each question prompt. This function is assumed to be parameterized by a vector θ , and is denoted as $\tau_{P_i; \theta}$. After the application of the positional editing function, a new score is generated, formulated as $r(\mathcal{M}, P_i, \tau_{P_i; \theta})$.

For instance, in retrieval-augmented generation (RAG) tasks, the prompt P_i is typically composed of three segments: the instruction, the documents, and the question. It can be possible to define $\theta = [\theta_1, \theta_2]$, while θ_1 translates to inserting θ_1 placeholder tokens between the instruction and the document segment, and θ_2 translates to inserting placeholder tokens between the document segment and the question.

Formally, prompt engineering is framed as an optimization problem. We aim at finding the optimal θ that maximizes the score:

$$\theta^* = \arg \max_{\theta} \frac{1}{N} \sum_{i=1}^N r(\mathcal{M}, P_i, \tau_{P_i; \theta}) \quad (8)$$

In this research, we utilize a basic algorithm for tackling the optimization problem by initially defining a limited number of candidates for θ and assessing each candidate's score via brute force. Notably, since θ is a numeric vector, the search process can be accelerated by adopting various optimizers, such as Gaussian processes of Bayesian optimization (Srinivas et al., 2010). The exploration of more sophisticated optimization methods will be considered in future works.

3 Experiments

In this section, we present our experiments and findings for position engineering. We evaluate two prevalent tasks for LLMs, namely Retrieval-Augmented Generation (RAG) and In-context Learning (ICL). Our primary testing model is Llama2-13B-chat (Touvron et al., 2023b), although we also expand our experiments to include additional models in the Appendix.

3.1 Position Engineering for RAG

Datasets: To explore the effectiveness of position engineering on RAG tasks, we utilize four open-domain QA datasets: NQ open (Lee et al., 2019), EntityQuestions (Sciavolino et al., 2021), TrivialQA (Joshi et al., 2017), and WebQuestions (Berant et al., 2013). These datasets each include a training and an evaluation (or test) set, with each set comprising a series of question-and-answer pairs. From the original training set of each dataset, we randomly select 300 QA pairs to serve as our training set for position engineering. Similarly, we randomly select 2,000 pairs from their original test sets to constitute our test set. In cases where a dataset lack a test set, we utilize its evaluation set instead. The Contriever model, which has been fine-tuned on the MS-MARCO dataset, is employed as the retrieval model (Izacard et al., 2021). We employ document passages from Wikipedia as our source for retrieval, with each passage containing a total of 100 words (Karpukhin et al., 2020). k document passages, specifically $k = 1, 3, 5$, are retrieved, and subsequently concatenated and fed into LLMs. Our evaluation metric is the best exact match accuracy, judging whether any correct answer is in the output, which is a common practice in previous works (Kandpal et al., 2023; Mallen et al., 2023).

Search Space: We adopt the following prompt template for all RAG experiments. The prompt template is divided into three segments. The first segment provides instructions for the task; the second segment presents a list of retrieved documents, each accompanied by its title and a passage; and the third segment combines the instruction with a specific question. These segments are referred to as the instruction segment, the document segment, and the question segment for convenience.

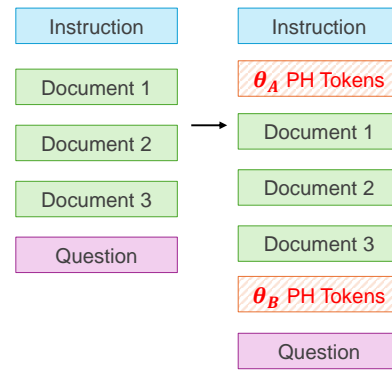


Figure 2: Position Engineering for RAG. In the figure, the term "PH tokens" refers to the placeholder tokens introduced in Section 2.2. We investigate a defined search space, with inserting θ_A placeholder tokens between the instruction and document segments, and θ_B placeholder tokens between the document and question segments. Both θ_A and θ_B range from $\{0, 100, \dots, 2500\}$, subject to $\theta_A + \theta_B \leq 2500$.

The prompt template for RAG:

Answer the question based on the given documents (some of which might be irrelevant). Only give me the answer and do not output any other words.

Document (Title: {title}) {passage}

Document (Title: {title}) {passage}

Document (Title: {title}) {passage}

Answer the question based on the given documents (some of which might be irrelevant). Only give me the answer and do not output any other words.

Question: {question}

Answer:

As presented in Figure 2, our study explores the methodology of position engineering for RAG by strategically inserting θ_A placeholder tokens between the instruction and document segments, and θ_B placeholder tokens between the document and question segments. To narrow down the search space, the values of θ_A and θ_B are limited to a predefined set $\{0, 100, \dots, 2500\}$. Additionally, we impose a restriction that $\theta_A + \theta_B \leq 2500$, due to the constraints of the context window size. We evaluate the performance of all combinations on the training set with the Llama2-13B-chat model, and then apply the best configuration to the test set.

Results: Table 1 displays the results for RAG, indicating that position engineering substantially en-

Dataset	N Doc	Baseline	Position Engineering	Abs Impr.	θ_A^*	θ_B^*
NQ Open	1	0.341	0.435	+9.5%	2,000	400
NQ Open	3	0.424	0.490	+6.6%	2,100	300
NQ Open	5	0.452	0.501	+5.0%	1,600	600
EntityQuestions	1	0.452	0.511	+5.8%	1,400	500
EntityQuestions	3	0.501	0.531	+3.0%	1,200	300
EntityQuestions	5	0.535	0.558	+2.3%	1,300	400
TrivialQA	1	0.582	0.657	+7.5%	1,300	200
TrivialQA	3	0.646	0.697	+5.1%	1,500	300
TrivialQA	5	0.669	0.698	+2.9%	2,300	200
WebQuestions	1	0.319	0.473	+15.4%	1,900	500
WebQuestions	3	0.410	0.507	+9.7%	2,100	400
WebQuestions	5	0.434	0.514	+8.1%	1,600	800

Table 1: The test results for RAG. We initially examine all possible combinations to determine the optimal configuration on the training set, which is denoted as θ_A^* and θ_B^* . This optimal configuration is then applied on the test set, and the results are presented in the table. The baseline is $\theta_A = \theta_B = 0$. The term "Abs Impr." represents absolute accuracy improvement in percentage. The Llama2-13B-chat model is utilized for the experimentation.

hance the RAG’s performance across all settings. The most notable improvement is 15.4%, observed in the WebQuestions dataset with a single retrieved document. The best-performing parameters, θ_A^* and θ_B^* , reveal a consistent trend: θ_A^* tends to be a large number, usually in the range of 1,000 to 2,000, while θ_B^* is a smaller figure, ranging between 200 and 600.

3.2 Universal Position Configuration for RAG

It has been observed that the most effective position configurations, represented as θ_A^* and θ_B^* in Section 3.1, demonstrate a consistent trend across all examined datasets. In this section, we aim to determine a single position setting that can enhance RAG performance universally across different datasets and various numbers of retrieved documents.

Given that absolute accuracy scores vary across datasets, we adopt the percentile value of the accuracy score as a metric to assess each position setting. In this context, we define "experiment setting" as the combination of one dataset and a specific number of retrieved documents, and "position setting" as a specific pair of θ_A and θ_B . For every experiment setting, we accumulate the scores from all position settings. The effectiveness of each position setting is then evaluated based on its percentile ranking, which varies from 0 to 100, within the experiment setting. Finally, The overall efficacy of a position setting is determined by averaging its percentile rankings across all experiment settings.




Figure 3: We visualize the average percentile values for each positional configuration (θ_A, θ_B) . These values are initially obtained by aggregating all accuracy scores for a given dataset and a specific number of retrieved documents, and calculate the percentile scores. Subsequently, they are averaged across all configurations, as detailed in Section 3.2.

The baseline configuration without position engineering ($\theta_A = \theta_B = 0$) achieves an average percentile of 31.6. This suggests that approximately 68% of configurations can surpass the baseline performance by simply adjusting positional information. The visualization of averaged percentiles for all position settings is provided in Figure 3.

Generally, it is advantageous to select a θ_A value within the range of 1300 to 2000, and set θ_B within the range of 300 to 500. Setting θ_B to an exces-

sively high figure (for instance, more than 1500) significantly deteriorates performance, possibly because it leads to the neglect of document information in prompts. Moreover, for each specified θ_B , an increase in θ_A is generally associated with better performance.

On the training set, $\theta_A = 1900, \theta_B = 400$ exhibits the highest percentile value of 92.9. We apply this configuration to the test set across all datasets and retrieved document numbers. Results presented in Table 2 demonstrate that it leads to a universal performance improvement. In Appendix A.1, we also demonstrate that such configuration remains effective for other models.

Dataset	N Doc	Abs Impr.
NQ Open	1	+9.6%
NQ Open	3	+7.1%
NQ Open	5	+4.9%
EntityQuestions	1	+5.6%
EntityQuestions	3	+3.4%
EntityQuestions	5	+1.9%
TrivialQA	1	+8.1%
TrivialQA	3	+4.9%
TrivialQA	5	+3.2%
WebQuestions	1	+14.8%
WebQuestions	3	+9.4%
WebQuestions	5	+9.1%

Table 2: The universal position configuration, $\theta_A = 1900, \theta_B = 400$, is tested on the test split of all datasets employing the Llama2-13B-chat model. The Table presents the absolute accuracy improvements over the baseline configuration ($\theta_A = \theta_B = 0$).

3.3 Without the instruction segment

From Figure 3, it is observed that a larger θ_A is preferred for optimal performance. θ_A represents the gap between the instruction segment and the document segment. A larger θ_A reduces the instruction segment’s impact. This raises the question of whether eliminating the instruction segment entirely could further enhance performance. To explore this, we conduct tests, and the outcomes are presented in Table 3. It is discovered that the performance of removing the instruction segment is comparable to the baseline setting. The most significant improvement, a 2% increase, is observed with the WebQuestions dataset when one retrieved document is utilized. However, the enhancement

Dataset	N Doc	Baseline	No Inst.
NQ Open	1	0.341	0.353
NQ Open	3	0.424	0.417
NQ Open	5	0.452	0.449
EntityQuestions	1	0.452	0.454
EntityQuestions	3	0.501	0.492
EntityQuestions	5	0.535	0.532
TrivialQA	1	0.582	0.582
TrivialQA	3	0.646	0.650
TrivialQA	5	0.669	0.668
WebQuestions	1	0.319	0.335
WebQuestions	3	0.410	0.410
WebQuestions	5	0.434	0.440

Table 3: We test the RAG performance without the instruction segment on the Llama2-13B-chat model. The results are comparable to the baseline, with a slight improvement ranging from 1% to 2% on the NQ Open and WebQuestions datasets when a single document is retrieved.

from position engineering in the same experiment setting is 15.4%. Thus, to achieve the best performance, it is essential to lessen but not eliminate the effect of the instruction segment, a goal that is easy for position engineering, but difficult to accomplish by prompt engineering.

3.4 Position Engineering for ICL

Datasets: To explore the impact of positional engineering on ICL tasks, we employ two datasets: TREC (Li and Roth, 2002; Hovy et al., 2001) and SST2 (Socher et al., 2013). The TREC dataset includes a variety of questions, with the aim being to categorize these questions into 6 coarse and 50 fine-grained question types. We focus on the 6 coarse question types. The SST2 dataset contains movie reviews, with the objective being to categorize these reviews as either positive or negative. For our training set, we randomly choose 300 samples from the original training sets of TREC and SST2. For our test set, we utilize TREC’s entire 500-sample test set. For the SST2 dataset, due to the lack of labels in its original test set, we use all 842 samples from its validation set as our test set. For each sample tested, we randomly select 3 examples of each label from the training set as the in-context demonstrations, leading to 18 examples for TREC and 6 for SST2. The exact match score is adopted as the evaluation metric.

Dataset	Baseline	Position Engineering	Abs Impr.	θ_A^*	θ_{mid}^*	θ_B^*
TREC	0.692	0.728	+3.6%	0	40	0
SST2	0.915	0.935	+1.9%	0	0	100

Table 4: The test results for ICL. We initially examine all possible combinations to determine the optimal configuration on the training set, which is denoted as $(\theta_A^*, \theta_{mid}^*, \theta_B^*)$. This optimal configuration is then applied on the test set, and the results are presented in the table. The baseline is $\theta_A = \theta_{mid} = \theta_B = 0$. The term "Abs Impr." represents absolute accuracy improvement in percentage. The Llama2-13B-chat model is utilized for this experimentation.

Search Space: The prompt template provided below is designed for evaluating performance on the SST2 dataset and is divided into three sections: an initial instruction segment that outlines the task, a middle segment that provides examples demonstrating the task, and a final segment that combines the instruction with a query. These segments are referred to as the instruction segment, the example segment, and the query segment, respectively. For the TREC dataset, we employ a similar prompt template, altering only the terms "Review" to "Question" and "Sentiment" to "Question Type" with Llama2-13B-chat.

To investigate the impact of position engineering, we conduct experiments by inserting θ_A placeholder tokens between the instruction and example segments, θ_B placeholder tokens between the example segment and the query segment, and θ_{mid} placeholder tokens among the examples, as depicted in Figure 4. The candidate value set of θ_A and θ_B is set to $\{0, 100, \dots, 600\}$, and while θ_{mid} is set to $\{0, 20, \dots, 100\}$. We evaluate the performance of all possible combinations within the training set and apply the optimal configuration to the test set.

The prompt template for the SST2 dataset:

Please determine the Sentiment of a Review according to the examples below.

Review: {query}
Sentiment: {label}

Review: {query}
Sentiment: {label}

Review: {query}
Sentiment: {label}

Now, you are given the following Review.
Review: {query}
Please output its Sentiment according to the examples. Only output its Sentiment without outputting anything else.
Sentiment:

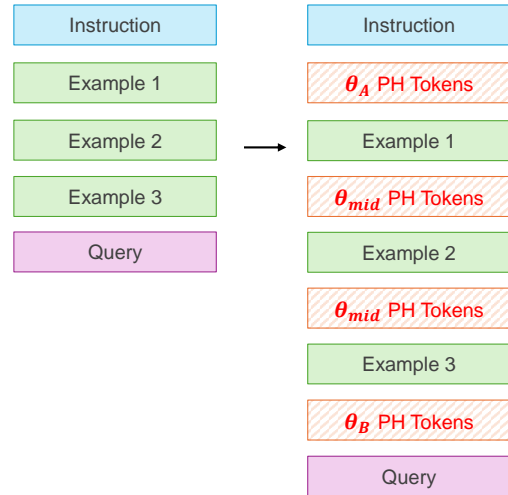


Figure 4: Position Engineering for ICL. In the figure, the term "PH tokens" refers to the placeholder tokens introduced in Section 2.2. We investigate a defined search space, with inserting θ_A placeholder tokens between the instruction and document segments, θ_B placeholder tokens between the document and question segments, and θ_{mid} placeholder tokens among the examples. The candidate value set of θ_A and θ_B is set to $\{0, 100, \dots, 600\}$, and while θ_{mid} is set to $\{0, 20, \dots, 100\}$.

Results: The results for ICL are presented in Table 4, indicating an enhancement in performance across both datasets, with an absolute 3.6% improvement observed on the TREC dataset and an absolute 1.9% improvement on the SST2 dataset. The optimal position settings, represented as θ_A^* , θ_B^* , and θ_{mid}^* , vary between datasets. Specifically, TREC requires adjusting θ_{mid} to 40, with θ_A and θ_B set to 0, whereas SST2 requires setting θ_B to 100, with θ_A and θ_{mid} to 0.

We observe a significant performance drop when θ_B is set within the $\{200, 300, \dots, 600\}$ range, mirroring the trends observed in RAG tasks where a high θ_B value leads to poor outcomes. θ_B can be interpreted as a parameter to adjust the impact of the example segment. In the case of SST2, which involves classifying sentiments of re-

views—a domain that LLMs might have common knowledge—the choice of $\theta_B^* = 100$ is intended to slightly reduce the example segment’s influence. For TREC, which requires LLMs to learn question types from examples, maintaining $\theta_B^* = 0$ is optimal.

4 Discussion

We hypothesize that position engineering serves as a technique to finely adjust the attention weights assigned to different segments within prompts. By extending the positional gap between two segments, the interaction between them is lessened, thereby increasing the attention allocated to other segments. For example, in RAG experiments, an increased value of θ_A could potentially reduce the impact of the instruction segment while amplifying the attention allocated to the retrieved documents. It is important to note, however, that the initial instruction remains essential, as evidenced in Section 3.3. Position engineering offers a nuanced approach to adjusting the weights of different blocks without the need for direct addition or removal of text.

Position engineering offers several advantages: (i) It is easier to optimize due to its numerical search space $\{\theta\}$, in contrast to prompt engineering, which requires searching over a more complex text space. (ii) It is computationally efficient, as altering position information merely involves updating the position indices input into LLMs, without increasing the overall computational overhead. (iii) It is orthogonal to prompt engineering, meaning the two approaches can be effectively combined.

Future works may advance in the following directions. Firstly, investigating the internal dynamics of LLMs can enhance our understanding of position engineering’s underlying mechanisms. Secondly, employing more sophisticated optimizers, such as Gaussian processes or multi-armed bandits, could reduce the search time and discover more refined position editing functions. Finally, the exploration of merging position engineering with prompt engineering could harness the full power of LLMs.

5 Related Works

Prompt engineering: Prompt engineering has emerged as a technique to enhance the performance of LLMs by modifying the instructions given to them. For instance, few-shot prompting allows LLMs to learn from demonstrations, a process also

known as in-context learning (Brown et al., 2020). Additionally, Chain-of-Thought prompting encourages LLMs to produce intermediate tokens, thereby improving their reasoning capabilities (Wei et al., 2022; Kojima et al., 2022). Another technique, Retrieval-Augmented Generation (RAG), involves retrieving relevant document passages and incorporating them into the prompts (Lewis et al., 2020). It has been discovered that the RAG performance can be improved by adding random documents to the mix of relevant documents (Cuconasu et al., 2024), a technique that is relevant to our study. However, this approach demands significant additional computational resources. In contrast, our proposed method does not require extra computation.

Positional Information in LLMs: Positional embedding has been introduced to integrate the position information of tokens within the attention layers (Vaswani et al., 2017). Initially, this concept relied on absolute position indices. However, subsequent developments have introduced methods based on relative positions, such as the relative positional encodings in Transformer-XL (Dai et al., 2019), and RoPE (Su et al., 2024). ALiBi is a different method for integrating positional information into LLMs (Press et al., 2022), which does not utilize embeddings but introduces a fixed bias based on relative positions during the computation of attention scores. More recent studies have focused on modifying positional embeddings to increase the context window size in LLMs (Ding et al., 2024; Peng et al., 2024). Apart from positional embeddings, the performance of LLMs has been found to correlate with document positions in prompts. In RAG tasks, documents that are positioned in the middle are often more neglected than those at the beginning or the end (Liu et al., 2024). However, to the best of our knowledge, there has been no similar effort on improving task performance by modifying positional indices.

6 Conclusion

In this study, we introduce position engineering, an innovative technique to enhance task performances of LLMs by merely altering the position information in the prompts. Our experimentation with position engineering across a range of tasks and models demonstrates its effectiveness. This approach provides a new avenue for maximizing the capabilities of LLMs.

7 Limitations

Our method needs an explicit search process to discover the optimal position setting for a given task. Such search process will cost computation resource and time. Sometimes, the search process can be omitted if a universal good positional setting exists, e.g. the universal setting for RAG tasks with Llama2-13B-chat model. Besides, the internal mechanism of position engineering remains unclear. We hypothesize that position engineering serves as a technique to finely adjust the attention weights assigned to different segments within prompts. Future efforts can be made to further investigate it.

References

- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Maarek, Nicola Tonellotto, and Fabrizio Silvestri. 2024. [The power of noise: Redefining retrieval for rag systems](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, page 719–729, New York, NY, USA. Association for Computing Machinery.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [Longrope: Extending llm context window beyond 2 million tokens](#). *ArXiv preprint*, abs/2402.13753.
- Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Petersen, and Julius Berner. 2024. [Mathematical capabilities of chatgpt](#). *Advances in Neural Information Processing Systems*, 36.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Retrieval augmented language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.
- Eduard Hovy, Laurie Gerber, Ulf Hermjakob, Chin-Yew Lin, and Deepak Ravichandran. 2001. [Toward semantics-based answer pinpointing](#). In *Proceedings of the First International Conference on Human Language Technology Research*.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. [Unsupervised dense information retrieval with contrastive learning](#). *ArXiv preprint*, abs/2112.09118.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *ArXiv preprint*, abs/2310.06825.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. [Large language models struggle to learn long-tail knowledge](#). In *International Conference on Machine Learning*, pages 15696–15707. PMLR.
- 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.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *Advances in neural information processing systems*, 35:22199–22213.

- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2023. Bloom: A 176b-parameter open-access multilingual language model.
- 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 S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *COLING 2002: The 19th International Conference on 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.
- Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Eureka: Human-level reward design via coding large language models. *ArXiv preprint*, abs/2310.12931.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. Crosslingual generalization through multitask finetuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. 2024. YaRN: Efficient context window extension of large language models. In *The Twelfth International Conference on Learning Representations*.
- Ofir Press, Noah A. Smith, and Mike Lewis. 2022. Train short, test long: Attention with linear biases enables input length extrapolation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, et al. 2024. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475.
- Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. Simple entity-centric questions challenge dense retrievers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. 2010. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 1015–1022. Omnipress.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *ArXiv preprint*, abs/2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *ArXiv preprint*, abs/2307.09288.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. [Large language models are human-level prompt engineers](#). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.

A Appendix

A.1 Applying the Universal RAG Configuration to Other Models

Dataset	N	Llama2-7B	Mistral-7B
NQ Open	1	+8.9%	+2.2%
NQ Open	3	+4.6%	+0.2%
NQ Open	5	+1.0%	-0.3%
EntityQuestions	1	+7.1%	+0.8%
EntityQuestions	3	+5.8%	-0.1%
EntityQuestions	5	+0.9%	-0.4%
TrivialQA	1	+7.9%	+3.1%
TrivialQA	3	+4.0%	+0.7%
TrivialQA	5	+1.6%	+0.5%
WebQuestions	1	+18.5%	+3.8%
WebQuestions	3	+10.0%	+2.2%
WebQuestions	5	+5.9%	0.0%

Table 5: We evaluate the universal position configuration for RAG, as identified in Section 3.2 with $\theta_A = 1900$ and $\theta_B = 400$, across the test splits of all datasets employing the Llama2-7B-chat and Mistral-7B-instruct-v0.2 models. The results showcase the absolute accuracy improvements over the baseline configuration, where θ_A and θ_B are both set to 0.

In Section 3.2, we identified a universal position configuration, $\theta_A = 1900$ and $\theta_B = 400$, on RAG tasks for the Llama2-13B-chat model. In this section, we further investigate whether such configuration remains effective for other models by applying it to the Llama2-7B-chat (Touvron et al., 2023b) and Mistral-7B-instruct-v0.2 (Jiang et al., 2023) model. The configuration is evaluated on the test splits across all datasets, with the results presented in Table 5. The findings indicate a consistent enhancement in the performance with the Llama2-7B-chat model under the universal position configuration. It is noteworthy that this configuration is initially identified with the Llama2-13B-chat model, suggesting that the Llama2-7B-chat model exhibits similar positional characteristics with Llama2-13B-chat. Furthermore, the Mistral-7B-instruct-v0.2 model also demonstrates consistent performance improvements when utilizing a single retrieved document. However, the performance gains become inconsistent with the use of multiple retrieved documents, indicating a potential need for model-specific adjustments.

A.2 Applying Position Engineering to Non-RoPE Models

In our previous evaluation section, Llama2-13B-chat was utilized as the primary model for testing. This model employs RoPE (Su et al., 2024) to integrate positional information. Furthermore, in this section, we aim to assess the effectiveness of position engineering using models with a different method for incorporating positional information. To this end, we apply position engineering to BLOOMZ-7b1 (Muennighoff et al., 2023) under the same experimental settings for the ICL tasks. BLOOMZ-7b1 is an instruction-fined version of BLOOM (Le Scao et al., 2023), which incorporates position information using ALiBi (Press et al., 2022). Unlike RoPE, ALiBi introduces a fixed position-related bias term during the computation of attention scores.

Specifically, we follow the search space in Figure 4 for ICL tasks. We determine the optimal position configuration on the training dataset by evaluating all configuration candidates in the search space, subsequently applying this configuration to the test set. Both the training and test sets remain the same with the previous settings. The results are presented in Table 6. Notably, there is a significant improvement in ICL tasks, with the SST2 dataset showing an absolute improvement of 11.0%. It demonstrates that position engineering can be also effective in non-RoPE models.

Dataset	Baseline	Position Engineering	Abs Impr.	θ_A^*	θ_{mid}^*	θ_B^*
TREC	0.724	0.782	+5.8%	0	0	200
SST2	0.836	0.946	+11.0%	0	20	500

Table 6: We apply position engineering to the BLOOMZ-7b1 model on ICL tasks. The same search space setting is employed as shown in Figure 4. θ_A^* , θ_{mid}^* , and θ_B^* is the optimal configuration identified in the training set, which is then applied on the test set. The baseline is $\theta_A = \theta_{mid} = \theta_B = 0$. The term "Abs Impr." represents absolute accuracy improvement in percentage compared to the baseline.