

# Enhancing LLM Knowledge Learning through Generalization

Mingkang Zhu<sup>1</sup>, Xi Chen<sup>2</sup>, Zhongdao Wang<sup>3</sup>, Bei Yu<sup>1</sup>, Hengshuang Zhao<sup>2</sup>, Jiaya Jia<sup>4,5</sup>

<sup>1</sup>CUHK <sup>2</sup>HKU <sup>3</sup>Huawei <sup>4</sup>SmartMore <sup>5</sup>HKUST  
mkzhu23@cse.cuhk.edu.hk, jia@cse.ust.hk

## Abstract

As Large language models (LLMs) are increasingly deployed in diverse applications, faithfully integrating evolving factual knowledge into these models remains a critical challenge. Continued pre-training on paraphrased data has shown empirical promise for enhancing knowledge acquisition. However, this approach is often costly and unreliable, as it relies on external models or manual effort for rewriting, and may inadvertently alter the factual content. In this work, we hypothesize and empirically show that an LLM’s ability to continually predict the same factual knowledge tokens given diverse paraphrased contexts is positively correlated with its capacity to extract that knowledge via question-answering. Based on this view and aiming to improve generalization to diverse paraphrased contexts, we introduce two strategies to enhance LLMs’ ability to predict the same knowledge tokens given varied contexts, thereby enhancing knowledge acquisition. First, we propose formatting-based data augmentation, which diversifies documents conveying the same knowledge by altering document formats rather than their content, thereby preserving factual integrity. Second, we adopt sharpness-aware minimization as the optimizer to better improve generalization. Extensive experiments demonstrate our methods’ effectiveness in both continued pre-training and instruction tuning, and further gains can be achieved by combining with paraphrased data. Code and data are available at <https://github.com/dvlab-research/llm-knowledge-generalization>.

## 1 Introduction

Large language models (LLMs) are pre-trained on large-scale corpora encompassing extensive knowledge, enabling them to demonstrate remarkable performance on knowledge-intensive tasks (Brown et al., 2020; OpenAI, 2023; Chowdhery et al., 2023; Zhang et al., 2022; Touvron et al., 2023a,b; Gem-

ini Team, 2023; Yang et al., 2024). As LLMs are increasingly deployed in real-world applications, an essential challenge is the effective integration of evolving factual knowledge. Prior work has shown that continued pre-training on a single document is insufficient for reliable knowledge elicitation via question answering (QA) (Kandpal et al., 2023). In contrast, training on multiple paraphrased versions of a document has been empirically shown to improve knowledge acquisition (Jiang et al., 2024; Allen-Zhu and Li, 2024). However, this method incurs significant costs and presents reliability concerns. Specifically, it typically depends on either computationally intensive external models or manual rewriting and has the risks of inadvertently altering factual information (Ding et al., 2024).

To address these challenges, we begin by analyzing the empirical observation that training with paraphrased documents improves QA accuracy for the corresponding embedded knowledge (Allen-Zhu and Li, 2024). We interpret diverse documents expressing the same factual content as samples drawn from a shared underlying distribution. From this perspective, training on a document containing certain knowledge tokens enhances LLMs’ ability to predict corresponding knowledge tokens given diverse unseen paraphrased preceding contexts. Further improvements can be achieved by training on paraphrased versions containing the same factual content, which enhances the model’s generalization to varied preceding contexts. Based on this insight, we hypothesize a positive correlation between an LLM’s ability to continually predict factual knowledge tokens given diverse unseen paraphrased contexts and its ability to extract that knowledge through QA, as demonstrated in Figure 1. This connection is not immediately intuitive from a human perspective, as declarative documents and question-answer formats differ significantly in structure.

To empirically validate our hypothesis, we con-

struct a biography dataset with diverse attributes, following the methodology in [Allen-Zhu and Li \(2024\)](#). Our observation reveals a strong positive correlation between the model’s accuracy in predicting knowledge tokens given unseen paraphrased contexts, and its accuracy in answering questions about the same knowledge. Additionally, integrating paraphrased documents in training enhances both accuracies, further reinforcing our hypothesis. These findings suggest that enhancing LLMs’ ability to predict knowledge tokens conditioned on varied contexts is a promising direction for improving knowledge acquisition.

Motivated by the findings, we propose two strategies to make accurate knowledge token prediction generalize to diverse paraphrased preceding contexts, thereby improving knowledge acquisition. As knowledge learning from a single document represents one-shot learning, while incorporating paraphrased documents transforms it into few-shot learning, we first propose the formatting-based data augmentation to diversify the documents containing the same knowledge. This approach modifies training documents’ formats—such as changes in spacing and padding—as a form of data augmentation. It eliminates the reliance on calling external models or manual labor and does not risk introducing factual inconsistencies. Second, we introduce sharpness-aware minimization (SAM) ([Foret et al., 2021](#)) as the optimization method to better improve generalization. When treating preceding document contexts as the input and knowledge tokens as targets, SAM facilitates generalization to diverse preceding document contexts with the same knowledge tokens. Furthermore, recent works identify that including different question paraphrases in instruction tuning helps knowledge extraction ([Fu et al., 2024](#)). Thus, we additionally incorporate our method into instruction tuning to improve the model’s generalization on question paraphrases.

We evaluate our methods on our constructed biography dataset following [Allen-Zhu and Li \(2024\)](#), and the Wiki2023 dataset ([Jiang et al., 2024](#)). Experiment results demonstrate significant improvements in LLMs’ knowledge learning ability for both the continued pre-training and instruction tuning phases. Moreover, combining our methods with paraphrased data leads to more improvement. Our key contributions can be summarized as follows:

- We hypothesize and empirically demonstrate that an LLM’s ability to continue knowledge

tokens given diverse paraphrased contexts is positively correlated with its capacity to answer related questions. This analysis provides a novel perspective on enhancing knowledge acquisition in LLMs and motivates us to improve the ability from the generalization perspective.

- We propose formatting-based data augmentation to automatically generate diverse training documents without relying on external paraphrasing tools or introducing factual inconsistencies. In addition, we propose to employ SAM as an optimization technique to boost generalization. We further apply our method on instruction tuning, enabling better generalization across paraphrased questions and enhancing knowledge extraction.
- Extensive experiments and ablation studies demonstrate our method’s effectiveness in improving factual knowledge learning during both the continued pre-training and instruction tuning phases.

## 2 Related Work

**Continued LLM Knowledge Learning.** As the pre-trained knowledge stored in LLMs quickly becomes outdated, adapting up-to-date information into LLMs becomes a critical problem. The primary approach to tackle this problem is through continued pre-training on documents containing up-to-date knowledge ([Ovadia et al., 2024](#); [Jiang et al., 2024](#); [Jang et al., 2022](#)). However, straightforward training on new corpus usually cannot lead to effective knowledge acquisition. This is likely due to the lack of diverse knowledge demonstrations like foundational or textbook knowledge ([Allen-Zhu and Li, 2024](#); [Jiang et al., 2024](#); [Ovadia et al., 2024](#); [Cheng et al., 2024](#)). Therefore, some works focus on paraphrasing documents to alleviate this issue ([Cheng et al., 2024](#); [Allen-Zhu and Li, 2024](#); [Ovadia et al., 2024](#)). However, paraphrasing manually or using external models can be expensive and tedious, and it might be unreliable as facts and knowledge inside documents could be altered ([Ding et al., 2024](#)). Therefore, we aim to avoid the risk of changing facts embedded in documents while enabling effective knowledge acquisition. Some works try to include QA data together with or before adapting to new documents ([Allen-Zhu and Li, 2024](#); [Jiang et al., 2024](#)). However, these methods introduce

new difficulties in finding effective arrangements and proportions of QA data and documents. To induce effective knowledge extraction during inference, instruction tuning on annotated QA pairs after continued pre-training on documents has become a common practice (Sanh et al., 2022; Wei et al., 2022; Mishra et al., 2022; Iyer et al., 2022; Köpf et al., 2023; Zhou et al., 2023; Sun et al., 2024, 2023; Fu et al., 2024). Thus, we also aim to achieve satisfactory knowledge extraction after instruction tuning.

### Understanding of LLM Knowledge Learning.

Several works are trying to understand how LLMs learn knowledge from documents and retrieve them in question answering. Akyürek et al. (2022) tries to detect training documents important for question-answering for pre-trained LLMs. A number of works find connections between the frequency of certain knowledge appearing in pre-training documents and its question-answering ability (Kandpal et al., 2023; Akyürek et al., 2022; Petroni et al., 2019; Kassner et al., 2020; Wei et al., 2021; Févry et al., 2020; De Cao et al., 2021). Recently, Allen-Zhu and Li (2024) and Ovadia et al. (2024) empirically observe that adding paraphrased documents in the pre-training and continued pre-training phase helps knowledge extraction.

**Data Augmentation for Natural Language Processing.** There is a rich literature on data augmentation techniques in natural language processing (Chen et al., 2023; Ding et al., 2024; Wei and Zou, 2019). A popular type of data augmentation is synonym substitution, which replaces words in documents with their synonyms according to pre-defined dictionaries (Kolomiyets et al., 2011; Wang and Yang, 2015; Zhang et al., 2015). Another popular class of data augmentation is inserting, replacing, deleting, and swapping words in documents (Wei and Zou, 2019; Iyer et al., 2022; Niu and Bansal, 2018; Miao et al., 2020). In the era of LLMs, paraphrasing documents or synthesizing data using LLMs become increasingly popular (Ding et al., 2024; Sharma et al., 2023; Nair et al., 2023). However, these data augmentation methods generally modify the semantics of original documents, and certain tokens in documents where the factual knowledge resides have the risk of being altered (Ding et al., 2024). Therefore, we opt to avoid such risks and design suitable data augmentation methods for LLM knowledge learning.

## 3 Analysis of LLM Knowledge Learning

In this section, we analyze the practical effectiveness of paraphrasing for knowledge learning, aiming to derive insights that inform the development of more effective knowledge acquisition methods.

### 3.1 Autoregressive Language Model

Let  $\theta$  denote the parameters of an autoregressive language model, and  $\mathcal{V}$  be a fixed vocabulary of tokens. Suppose document  $\mathbf{d}$  contains a sequence of tokens  $(x_1, x_2, \dots, x_T)$  from  $\mathcal{V}$ , where  $T$  is the length of the document. The training sequence has a special token  $x_0 = \langle \text{bos} \rangle$  prepended, indicating the sequence’s start. The autoregressive language modeling task estimates the conditional probability distribution  $P(x_t|x_{<t})$  for each  $t = 1, 2, \dots, T$ . This is typically achieved by training a deep neural network to predict the next token  $x_t$  given the previous tokens  $x_0, x_1, x_2, \dots, x_{t-1}$ . For document  $\mathbf{d}$ , the negative log-likelihood loss function of the observed sequence  $\mathbf{d}$  is:

$$\begin{aligned} \ell(\theta, \mathbf{d}) &= -\log P_{\theta}(\mathbf{d}) \\ &= -\log P_{\theta}(x_1, x_2, \dots, x_T | x_0) \\ &= -\log \prod_{t=1}^T P_{\theta}(x_t|x_{<t}) \\ &= -\sum_{t=1}^T \log P_{\theta}(x_t|x_{<t}). \end{aligned} \quad (1)$$

### 3.2 Connecting Paraphrasing with the Effectiveness of Knowledge Learning

Building on prior empirical findings that training with paraphrased documents containing the same knowledge enhances knowledge learning (Jiang et al., 2024; Allen-Zhu and Li, 2024), we investigate the underlying mechanism driving this effect. We conceptualize diverse documents that convey the same knowledge as samples drawn from a shared distribution. From this perspective, training on documents containing specific knowledge tokens improves an LLM’s ability to predict those knowledge tokens given diverse unseen paraphrased preceding contexts. This ability is further strengthened by training on paraphrased variants, which promotes generalization across varied preceding context expressions. Therefore, we hypothesize a positive correlation between an LLM’s ability to continue factual knowledge tokens from diverse paraphrased contexts and its ability to extract that

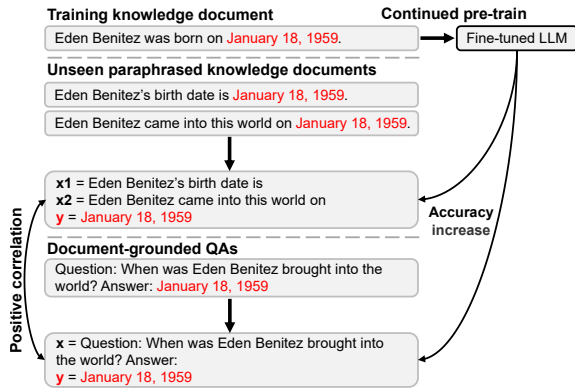


Figure 1: Conceptual illustration of our hypothesis. The red tokens represent knowledge tokens. Continued pre-training on raw knowledge documents leads to prediction accuracy increase of knowledge tokens conditioned on both preceding contexts of unseen paraphrased documents containing the same knowledge and related questions. Incorporating paraphrased training documents further improves both accuracies. Thus, their accuracies are positively correlated.

knowledge through QA. The conceptual illustration of our hypothesis is demonstrated in Figure 1. However, this is not obvious from the human perspective, as documents and questions asking for embedded knowledge are different: one is a declarative sentence, while the other is in the form of QA. Therefore, we need to verify this hypothesis.

### 3.3 Verification

To validate our hypothesis in a controlled setting, we generate synthetic human biography data following the methodology of Allen-Zhu and Li (2024). Specifically, we create 1,000 randomly generated human biography profiles, each characterized by five attributes: birth date, college, major, hometown, and company. To construct training documents, we employ a predefined template to represent these five attributes and leverage ChatGPT to generate two additional paraphrased templates. Each profile’s attributes are populated into these templates to form the training documents. For evaluation, we generate five additional paraphrased biography documents and create five questions for each attribute. The templates, example documents, and QAs are included in Appendix B.3.

Then we conduct a study on two scenarios: (1) **CPT**, each biography profile has a single document representation; (2) **Paraphrase CPT**, each biography profile has 3 paraphrased document representations. We continually pre-train Qwen 2 1.5B (Yang et al., 2024) with the above settings and record the

Setting	CPT	Para. CPT	CPT + Ours	Para. CPT + Ours
Pearson corr.	0.97	1.00	1.00	1.00
Spearman corr.	0.95	1.00	0.99	0.98

Table 1: Pearson and Spearman correlation between the prediction accuracy of first knowledge tokens for each attribute conditioned on: (1) preceding tokens in unseen paraphrased documents, (2) testing questions.

first knowledge token’s accuracy conditioned on preceding contexts tokens in 5 unseen paraphrased documents and related testing questions. If the model’s accuracy predicting the first knowledge tokens conditioned on preceding tokens in unseen documents increases along with that conditioned on related questions, a positive correlation exists between these two metrics. The results are presented in Figure 2.

In Figure 2, red lines represent the prediction accuracy of first knowledge tokens conditioned on their preceding tokens in unseen paraphrased documents, while blue lines represent the accuracy of first knowledge tokens conditioned on testing questions. We can see that the accuracy conditioned on questions is increasing along with the accuracy conditioned on preceding tokens in unseen paraphrased documents. Moreover, this also holds true for training on paraphrased documents, which leads to much higher accuracies. We further calculate the Pearson and Spearman correlation between the accuracies given unseen documents and questions in Table 1, which also demonstrate a strong positive correlation between these two metrics. These observations validate the positive correlation between the LLM’s ability to continue factual knowledge tokens given diverse unseen paraphrased contexts and its ability to extract that knowledge through QA. Hence, enhancing LLMs’ ability to predict knowledge tokens conditioned on varied document contexts is a promising direction for improving knowledge acquisition.

## 4 Methods

Section 3 has revealed that an LLM’s ability to continue knowledge tokens given diverse paraphrased contexts is strongly positively correlated with its capacity to answer related questions. Thus, in this section, we explore two methods to improve LLM knowledge learning, by enabling accurate prediction of knowledge tokens to generalize from training documents to unseen paraphrased variants that convey the same factual content.






Figure 2: Average first knowledge token accuracy for each attribute conditioned on: (1) preceding tokens in unseen paraphrased documents, (2) preceding tokens in testing questions.

#### 4.1 Formatting-Based Data Augmentation

Based on the analysis in Section 3.3, knowledge learning is one-shot without paraphrasing and few-shot with paraphrased documents. Thus, providing diverse documents containing the same knowledge is crucial for effective knowledge learning. However, paraphrasing manually or via external models can be expensive and laborious, and might alter the embedded facts (Ding et al., 2024). Moreover, certain expressions and terminologies are irreplaceable and must be used in their exact form. Nor should mottoes and poems be paraphrased when they are in training documents. Therefore, we aim to develop methods to reliably increase the variety of training documents containing the same knowledge without paraphrasing.

We draw inspiration from the formatting of texts. We may often encounter variations in the formatting used to present texts, such as whether to indent the beginning of a paragraph and whether to use spaces or tabs as indentations for codes. There are also variations for using single-space or double-space spacing in the era of typewriters (Wikipedia contributors, 2024). These formatting differences, while altering some of the format tokens, do not affect the semantic meaning and knowledge of the text itself. Therefore, given a training document, we propose to apply the following formatting-based data augmentations:

- **Wrapping.** Augmented documents are created by wrapping the document with quotes, asterisks, brackets, or parentheses. This is used to mimic the case that the document is quoted, highlighted, or appears in Markdown.
- **Left Padding.** Augmented documents are created by padding spaces, tabs, or pound signs to the left of the document. This is to mimic the scenarios of the document being written using Markdown or appearing as a paragraph in a paper.

- **Random Space Insertion.** Augmented documents are created by randomly inserting additional spaces adjacent to the original spaces. This simulates the case that the training document is presented using different spacing and includes some unintentional extra spaces.

Some augmentation examples are as follows:

```

Raw document
Eden Benitez was raised in Santa Clarita.
Wrapping Augmentation
*Eden Benitez was raised in Santa Clarita.*
Left Padding Augmentation
# Eden Benitez was raised in Santa Clarita.
Random Space Insertion Augmentation
Eden Benitez was raised in Santa Clarita.

```

Detailed specifications of the data augmentations are discussed in Appendix B.2. With these augmented documents, we automatically diversify the training documents while not changing the knowledge and facts inside these documents.

#### 4.2 Generalization Regularization

Given a training document  $d$ , the autoregressive objective in Equation (1) would minimize the negative log-likelihood (NLL) loss for a set of samples  $(x_{<t}, x_t)$ . Thus, we can further enhance the generalization ability by applying generalization regularization methods designed for traditional supervised problems. Applying generalization regularization on samples with  $x_t$  being knowledge tokens and  $x_{<t}$  being the corresponding preceding tokens in documents can generalize to samples with the same knowledge token label  $x_t$ , but different preceding tokens, such as  $x_{<t}$  in unseen paraphrased documents. Recently, Foret et al. (2021) developed the *Sharpness-Aware Minimization* (SAM) to improve the generalization ability for supervised problems (Baek et al., 2024; Chen et al., 2022; Foret et al., 2021). We adopt this technique to further improve LLM knowledge learning.

Given a training document  $d$ , let  $\mathcal{B}$  be the set of training document samples with knowledge tokens as pseudo-label and preceding tokens as inputs, and according to SAM we solve the following problem:

$$\min_{\theta} \max_{\|\epsilon\|_2 \leq \rho} L_{\mathcal{B}}(\theta + \epsilon) + \lambda \|\theta\|_2^2, \quad (2)$$

where  $\rho \geq 0$  is a given perturbation radius,  $\lambda$  is a small positive regularization constant. The objective is to find a minimizer with the neighborhood where the loss does not increase too much. According to SAM, the inner maximization problem in Equation (2) is solved approximately at  $\hat{\epsilon} = \rho \nabla L_{\mathcal{B}}(\theta) / \|\nabla L_{\mathcal{B}}(\theta)\|_2$  by the first-order Taylor expansion. Then, the objective function of Equation (2) changes to  $L_{\mathcal{B}}(\theta + \hat{\epsilon}) + \lambda \|\theta\|_2^2$ , on which the gradient descent is performed.

As shown in Figure 2 and Table 1, our hypothesis in Section 3.2 continues to hold with our proposed methods, further supporting its empirical validity.

### 4.3 Adaptation to Instruction Tuning

Instruction tuning has become a common practice to make LLMs follow human instructions and perform question-answering (Sanh et al., 2022; Wei et al., 2022; Ouyang et al., 2022). Instruction tuning computes the negative log-likelihood loss only on tokens in answers with questions as the context:  $L_{\mathbf{a}} = -\sum_t \log P(\mathbf{a}_t | \mathbf{q}, \mathbf{a}_{<t})$ . Recent works identify that including diverse question paraphrases in instruction tuning helps knowledge extraction (Fu et al., 2024). Therefore, we propose to use SAM and apply our formatting-based data augmentation to the questions for instruction tuning. In this way, LLM would be able to respond accurately to different paraphrases of a question during instruction tuning, thereby enhancing knowledge extraction.

## 5 Experiments

### 5.1 Experiment Settings

**Baseline Methods.** We experiment with two standard baselines: (1) continued pre-training (CPT) and (2) continued pre-training with instruction tuning (IT), and demonstrate the effectiveness of our methods in improving the knowledge learning abilities of these baselines.

**Base Models.** We use Qwen 2 1.5B (Yang et al., 2024) and LLaMA 2 7B (Touvron et al., 2023b) as base models, and test all baselines and their combination with our methods on these models.

**Datasets.** We use our biography dataset generated following Allen-Zhu and Li (2024), and the

Wiki2023-film dataset (Jiang et al., 2024) for the experiment. For the biography dataset, we follow Allen-Zhu and Li (2024) to continually pre-train on documents of all biography profiles and instruction-tune on 1 QA pair per attribute for half of the biography profiles. The evaluation is conducted for the remaining half individuals. Our evaluation differs from Allen-Zhu and Li (2024), which evaluates only 1 question prompt per attribute. We generate 5 different question prompts for each attribute to better evaluate the generalization ability, totaling 12500 QA pairs. The biography dataset is synthetic while the recipe for generating the Wiki2023-film dataset minimizes overlap with the pre-training corpus. Thus, experimenting on these two datasets can mimic the difficult case of continued knowledge learning on up-to-date information.

For the biography dataset, we experiment with both the **CPT** and **Paraphrase CPT** settings illustrated in Section 3.3. For the Wiki2023-film dataset, we experiment with the **CPT** setting since there are no paraphrased documents. All comparing methods are trained with the **same number of steps** in both the continued pre-training and instruction tuning phases for fair comparison.

**Evaluation Metrics.** As we aim to evaluate the closed-book free-form question-answering ability, we utilize exact match (EM) between the model generations and ground truth answers as the evaluation metric (Kwiatkowski et al., 2019). We also report Recall and F1 scores to better assess questions with long answers. When evaluating models that have not been instruction-tuned, we prepend 1 QA pair for the biography dataset and 5 QA pairs for the Wiki2023-film dataset to make sure that models can follow the QA format.

### 5.2 Main Results

**Results on Biography Dataset.** We present the results on the biography dataset in Table 2. We can see that before training, base models cannot answer questions at all. This effectively simulates the case of LLMs adapting to up-to-date information, which is considered nontrivial (Jiang et al., 2024; Ovadia et al., 2024). We can see that our method leads to substantial improvement over baselines for both cases of training on with and without paraphrased documents. Notably, we observe gains of up to 50 in EM, 35 in Recall, and 44.9 in F1.

Besides, our method without paraphrased documents even significantly outperforms naively training on paraphrased documents in the continued pre-

	Qwen 2 1.5B			Qwen 2 1.5B w/ IT			LLaMA 2 7B			LLaMA 2 7B w/ IT		
	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1
<b>Base</b>	0.7	7.1	6.2	-	-	-	0.7	8.9	7.6	-	-	-
<b>CPT</b>	7.1	16.1	12.4	52.8	57.6	57.1	52.0	59.1	58.5	89.6	91.3	91.2
<b>w/ Ours</b>	<b>43.2</b>	<b>57.7</b>	<b>52.3</b>	<b>57.9</b>	<b>62.2</b>	<b>61.8</b>	<b>85.4</b>	<b>88.2</b>	<b>88.0</b>	<b>93.3</b>	<b>94.2</b>	<b>94.1</b>
<b>Paraphrase CPT</b>	24.9	45.4	35.1	54.5	59.9	59.6	54.3	70.6	63.6	94.2	95.9	95.9
<b>w/ Ours</b>	<b>74.9</b>	<b>80.4</b>	<b>80.0</b>	<b>75.3</b>	<b>77.2</b>	<b>76.9</b>	<b>89.2</b>	<b>93.1</b>	<b>92.8</b>	<b>98.4</b>	<b>99.0</b>	<b>99.0</b>

Table 2: Experiment results on the biography dataset with the base models continued pre-training and instruction-tuning by our method and baselines. Our method leads to substantial improvement in knowledge acquisition and extraction for both phases compared to baselines.

	Qwen 2 1.5B			Qwen 2 1.5B w/ IT			LLaMA 2 7B			LLaMA 2 7B w/ IT		
	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1
<b>Base</b>	3.4	7.2	7.5	-	-	-	5.6	18.8	16.9	-	-	-
<b>CPT</b>	7.2	19.4	17.8	12.3	23.9	24.0	11.8	32.7	27.1	31.3	47.4	46.9
<b>w/ Ours</b>	<b>9.8</b>	<b>24.9</b>	<b>22.0</b>	<b>14.8</b>	<b>27.0</b>	<b>27.3</b>	<b>17.6</b>	<b>42.3</b>	<b>34.7</b>	<b>38.6</b>	<b>56.0</b>	<b>55.2</b>

Table 3: Experiment results on the Wiki2023-film dataset with the base models continued pre-training and instruction-tuning by our method and baselines. Our method leads to nontrivial improvement in knowledge acquisition and extraction for both phases compared to baselines.

training phase, and leads to on-par performance for the instruction tuning phase. This result shows that our method can serve as an effective and reliable alternative to tedious and unreliable paraphrasing. When applying our methods to paraphrased documents, the knowledge learning performance becomes even better, showing that our methods can induce more generalization ability with more diverse documents conveying the same knowledge. This also demonstrates that our method can gain more enhancement when used together with paraphrasing. Moreover, our method leads to much better performance than baselines prior to the instruction tuning stage. The ability to extract learned knowledge at this early stage further demonstrates the effectiveness of our method in knowledge learning. This property could be beneficial in scenarios with limited resources, such as adapting LLMs to new domains where it is challenging and labor-intensive to annotate instruction-following examples.

**Results on Wiki2023 Dataset.** Next, we evaluate our method with baselines on the Wiki2023-film dataset. As it does not have paraphrased training documents, we continually pre-train using a single document for all comparing methods. As shown in Table 3, our method consistently outperforms the baselines, demonstrating stable improvements in both the continued pre-training and instruction tuning phases. Notably, we observe gains of up to

7.3 in EM, 9.6 in Recall, and 8.3 in F1.

Additionally, we can observe from Table 2 and Table 3 that, our approach is consistently effective across different models, training phases, and datasets, demonstrating the robustness and effectiveness of our approach. We also want to stress that all comparing methods are trained with **the same number of steps** in both the continued pre-training and instruction tuning phases. The performance gain of our approach and paraphrasing is not attributable to an increased number of training steps on enlarged datasets. This result further stresses the importance of our findings.

### 5.3 Analysis

In this section, we conduct comprehensive ablation studies on the effect of each component of our methods on both continued pre-training and instruction tuning. We also analyze the effect of our formatting-based data augmentation compared to traditional NLP augmentations.

**Effect of Our Methods on Continued Pre-training.** We first ablate the effect of our formatting-based data augmentation and SAM for the continued pre-training phase. We use Qwen 2 1.5B (Yang et al., 2024) as the base model and conduct experiments on our synthesized biography dataset. We can see from Table 4 that, when training with paraphrased documents, both SAM and formatting-based data augmentation alone can

Training Setting	CPT			Paraphrase CPT		
	EM	Recall	F1	EM	Recall	F1
CPT	7.1	16.1	12.4	24.9	45.4	35.1
w/ Format Aug.	24.3	38.6	33.1	54.9	78.8	65.3
w/ SAM	19.7	29.1	26.6	52.8	63.3	60.8
w/ Format Aug.+SAM	<b>43.2</b>	<b>57.7</b>	<b>52.3</b>	<b>74.9</b>	<b>80.4</b>	<b>80.0</b>

Table 4: Ablation study on the effect of integrating each component of our method into the continued pre-training phase.

Training Setting	CPT			Paraphrase CPT		
	EM	Recall	F1	EM	Recall	F1
IT	52.8	57.6	57.1	70.4	72.8	72.6
w/ Format Aug.	55.3	59.5	59.2	73.2	75.1	74.9
w/ SAM	56.0	60.5	60.0	73.6	75.8	75.5
w/ Format Aug.+SAM	<b>57.9</b>	<b>62.2</b>	<b>61.8</b>	<b>75.3</b>	<b>77.2</b>	<b>76.9</b>

Table 5: Ablation study on the effect of integrating each component of our method into instruction tuning.

bring measurable enhancement over the baseline. Furthermore, when SAM and our data augmentation are combined, the performance gains are further amplified. When training without paraphrased documents, the lack of diverse documents conveying the same knowledge decreases the performance gain from SAM alone. Our data augmentation, on the other hand, brings adequate document variety, which leads to substantial improvement over the baseline. With the document diversity, SAM is able to boost the performance even further.

### Effect of Our Methods on Instruction Tuning.

Next, we ablate the effect of our method on the instruction tuning phase. Still, we conduct experiments on our generated biography dataset. We use Qwen 2 1.5B (Yang et al., 2024) continually pre-trained by our method as the starting point for instruction tuning. Prior works generally consider that knowledge is learned during continued pre-training and then made extractable in the instruction tuning phase (Allen-Zhu and Li, 2024; Ouyang et al., 2022; Sanh et al., 2022; Wei et al., 2022). Therefore, starting from the same continually pre-trained model, we can analyze how our methods influence knowledge extraction in this ablation. From Table 5, we can see that both SAM and formatting-based data augmentation alone can improve knowledge elicitation over baseline instruction tuning. Furthermore, the combination of them leads to better performance, with up to 5.1 increases on EM, 4.6 increases on Recall, and 4.7 increases on F1, over baseline instruction tuning. This result shows that our method can lead to gen-

Training Setting	EM	Recall	F1
CPT	7.1	16.1	12.4
w/ EDA	0.0	9.0	3.8
w/ Format Aug.	<b>24.3</b>	<b>38.6</b>	<b>33.1</b>

Table 6: Comparison between our formatting-based data augmentation and EDA (Wei and Zou, 2019).

eralization on various question paraphrases, which prior works identify as important for knowledge extraction (Fu et al., 2024).

**Comparison with Traditional NLP Augmentation.** We compare our formatting-based data augmentation with a representative traditional NLP data augmentation technique, EDA (Wei and Zou, 2019). The experiment is conducted by continually pre-training Qwen 2 1.5B on the biography dataset under the CPT setting. From Table 6, we can see that EDA is harmful to knowledge learning. The exact match decreases from 7.1 to 0 when applying EDA. EDA uses random word insertion, random word deletion, and random word swap, which might be appropriate for improving the language modeling ability. However, EDA is highly likely to alter the knowledge in documents, making it unsuitable for the knowledge learning task. This further demonstrates our formatting-based augmentation’s advantage that it reliably increases document diversity without changing the embedded knowledge.

## 6 Conclusion

In this paper, we try to improve LLM knowledge learning without costly and unreliable paraphrasing. We hypothesize and empirically show the positive correlation between an LLM’s ability to continue factual knowledge tokens given diverse paraphrased contexts and its capacity to extract that knowledge via question-answering. Based on this insight and from the generalization perspective, we propose formatting-based data augmentation to diversify the training documents without expensive and unreliable paraphrasing. We also propose to adopt SAM to further enhance generalization. Additionally, we incorporate our methods to instruction tuning to improve knowledge extraction. Extensive experiments demonstrate our methods’ effectiveness in improving knowledge acquisition and extraction for both continued pre-training and instruction tuning phases. We hope our work can provide insights to better understand and develop effective methods for LLM knowledge learning.



## Limitations

This paper introduces novel methods to enhance knowledge learning in LLMs, grounded in our empirically supported hypothesis that an LLM’s ability to continue factual knowledge tokens given diverse paraphrased contexts is positively correlated with its capacity to extract that knowledge through question-answering. While our experiments consistently validate this correlation, we acknowledge that our findings are empirical in nature, and a formal theoretical justification is not provided. We view this as a natural direction for future work. While our method improves knowledge learning for LLMs, it may also amplify factual inaccuracies if the training data contains misinformation.

## References

- Ekin Akyürek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas, and Kelvin Guu. 2022. Towards tracing knowledge in language models back to the training data. In *Findings of the Association for Computational Linguistics: EMNLP*.
- Zeyuan Allen-Zhu and Yuanzhi Li. 2024. Physics of language models: Part 3.1, knowledge storage and extraction. In *Proceedings of the 41st International Conference on Machine Learning*, pages 1067–1077.
- Christina Baek, J Zico Kolter, and Aditi Raghunathan. 2024. Why is SAM robust to label noise? In *The Twelfth International Conference on Learning Representations*.
- 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, and 12 others. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*.
- Jiaao Chen, Derek Tam, Colin Raffel, Mohit Bansal, and Diyi Yang. 2023. An empirical survey of data augmentation for limited data learning in NLP. *Transactions of the Association for Computational Linguistics*, 11:191–211.
- Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. 2022. When vision transformers outperform ResNets without pre-training or strong data augmentations. In *International Conference on Learning Representations*.
- Daixuan Cheng, Shaohan Huang, and Furu Wei. 2024. Adapting large language models via reading comprehension. In *The Twelfth International Conference on Learning Representations*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, and 48 others. 2023. PaLM: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. Autoregressive entity retrieval. In *The 9th International Conference on Learning Representations*.
- Bosheng Ding, Chengwei Qin, Ruochen Zhao, Tianze Luo, Xinze Li, Guizhen Chen, Wenhan Xia, Junjie Hu, Anh Tuan Luu, and Shafiq Joty. 2024. Data augmentation using LLMs: Data perspectives, learning paradigms and challenges. In *Findings of the Association for Computational Linguistics*, pages 1679–1705.
- Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. Entities as experts: Sparse memory access with entity supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. 2021. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*.
- Junbo Fu, Guoshuai Zhao, Yimin Deng, Yunqi Mi, and Xueming Qian. 2024. Learning to paraphrase for alignment with LLM preference. In *Findings of the Association for Computational Linguistics: EMNLP 2024*.
- Gemini Team. 2023. **Gemini: A family of highly capable multimodal models**. *Preprint*, arXiv:2312.11805.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, Xian Li, Brian O’Horo, Gabriel Pereyra, Jeff Wang, Christopher Dewan, Asli Celikyilmaz, Luke Zettlemoyer, and Ves Stoyanov. 2022. OPT-IML: scaling language model instruction meta learning through the lens of generalization. *Preprint*, arXiv:2212.12017.
- Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun KIM, Stanley Jungkyu Choi, and Minjoon Seo. 2022. Towards continual knowledge learning of language models. In *International Conference on Learning Representations*.
- Zhengbao Jiang, Zhiqing Sun, Weijia Shi, Pedro Rodriguez, Chunting Zhou, Graham Neubig, Xi Lin, Wen-tau Yih, and Srini Iyer. 2024. Instruction-tuned language models are better knowledge learners. In *Proceedings of the 62nd Annual Meeting of the 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 *Proceedings of the 40th International Conference on Machine Learning*.
- Nora Kassner, Benno Krojer, and Hinrich Schütze. 2020. Are pretrained language models symbolic reasoners over knowledge? In *Proceedings of the 24th Conference on Computational Natural Language Learning*.
- Oleksandr Kolomiyets, Steven Bethard, and Marie-Francine Moens. 2011. Model-portability experiments for textual temporal analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 271–276.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Minh Nguyen, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Alexandrovich Glushkov, Arnav Varma Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Julian Mattick. 2023. [Openassistant conversations - democratizing large language model alignment](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Zhengjie Miao, Yuliang Li, Xiaolan Wang, and Wang-Chiew Tan. 2020. Snippext: Semi-supervised opinion mining with augmented data. In *Proceedings of The Web Conference 2020*, pages 617–628.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. Cross-task generalization via natural language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 3470–3487.
- Varun Nair, Elliot Schumacher, and Anitha Kannan. 2023. Generating medically-accurate summaries of patient-provider dialogue: A multi-stage approach using large language models. In *Proceedings of the 5th Clinical Natural Language Processing Workshop*, pages 200–217.
- Tong Niu and Mohit Bansal. 2018. Adversarial oversensitivity and over-stability strategies for dialogue models. In *The SIGNLL Conference on Computational Natural Language Learning (CoNLL)*.
- OpenAI. 2023. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Oded Ovadia, Menachem Brief, Moshik Mishaeli, and Oren Elisha. 2024. [Fine-tuning or retrieval? comparing knowledge injection in LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 237–250.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*, pages 2463–2473.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, and 21 others. 2022. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations*.
- Ashwyn Sharma, David Feldman, and Aneesh Jain. 2023. Team cadence at MEDIQA-chat 2023: Generating, augmenting and summarizing clinical dialogue with large language models. In *Proceedings of the 5th Clinical Natural Language Processing Workshop*, pages 228–235.
- Zhiqing Sun, Yikang Shen, Hongxin Zhang, Qinhong Zhou, Zhenfang Chen, David Daniel Cox, Yiming Yang, and Chuang Gan. 2024. [SALMON: Self-alignment with instructable reward models](#). In *The Twelfth International Conference on Learning Representations*.
- Zhiqing Sun, Yikang Shen, Qinhong Zhou, Hongxin Zhang, Zhenfang Chen, David Daniel Cox, Yiming Yang, and Chuang Gan. 2023. [Principle-driven self-alignment of language models from scratch with minimal human supervision](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [LLaMA:](#)

- Open and efficient foundation language models. *Preprint*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- William Yang Wang and Diyi Yang. 2015. That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned language models are zero-shot learners. In *The Tenth International Conference on Learning Representations*.
- Jason Wei, Dan Garrette, Tal Linzen, and Ellie Pavlick. 2021. Frequency effects on syntactic rule learning in transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Jason Wei and Kai Zou. 2019. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6382–6388.
- Wikipedia contributors. 2024. [Sentence spacing — Wikipedia, the free encyclopedia](#). [Online; accessed 30-September-2024].
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 43 others. 2024. [Qwen2 technical report](#). *Preprint*, arXiv:2407.10671.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [OPT: Open pre-trained transformer language models](#). *Preprint*, arXiv:2205.01068.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, page 649–657.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: Less is more for alignment](#). In *Advances in Neural Information Processing Systems*, pages 55006–55021.

## A More Experiments

### A.1 Ablation on Data Augmentation

Data Augmentation	EM	Recall	F1
CPT	7.1	16.1	12.4
w/ Wrapping	21.0	34.0	28.3
w/ Left Padding	10.7	24.1	16.9
w/ Random Space Insertion	15.2	37.7	24.7
w/ All Three Data Augmentation	<b>24.3</b>	<b>38.6</b>	<b>33.1</b>

Table 7: Ablation study on the effect of our proposed three types of formatting-based data augmentations on the continued pre-training phase.

We conduct the ablation study on the effect of our proposed three formatting-based augmentations for Qwen 2 1.5B on the biography dataset during the continued pre-training stage with the **CPT** setting. The results are summarized in Table 7. It can be seen that all three types of formatting-based augmentations lead to significant improvement over the baseline. Among them, Wrapping and Random Space Insertion are more effective than Left Padding. The combination of all three types of augmentation creates more diverse documents and leads to more considerable enhancement.

## B Experiment Details

### B.1 Hyperparameter Settings.

We use AdamW (Loshchilov and Hutter, 2019) as the base optimizer and a weight decay of 0.1. The learning rate for continued pre-training is set to  $3e-5$  while the learning rate for instruction tuning is set to  $5e-6$  for experiments on both the biography and Wiki2023 dataset. We use a batch size of 128 for the biography dataset and a batch size of 256 for the Wiki2023 dataset. For continued pre-training, we include the value of  $\rho$  in Table 8. For instruction tuning, we use  $\rho = 0.025$  for all our experiments. For the experiment on the Wiki2023 dataset (Jiang et al., 2024), we continually pre-train both Qwen 2 1.5B (Yang et al., 2024) and LLaMA 2 7B (Touvron et al., 2023b) for 30 epochs. Qwen 2 1.5B is instruction-tuned for 5 epochs while LLaMA 2 7B is tuned for 2 epochs. For the experiment on the

Base Model	Biography w/o Paraphrase	Biography w/ Paraphrase	Wiki2023-film
Qwen 2 1.5B	0.05	0.015	0.05
LLaMA 2 7B	0.025	0.015	0.025

Table 8: The value of SAM’s  $\rho$  used in different continued pre-training experiment settings.

biography dataset, we continually pre-train both Qwen 2 1.5B (Yang et al., 2024) for 30 epochs and LLaMA 2 7B (Touvron et al., 2023b) for 15 epochs. Both models are instruction-tuned for 5 epochs. We use the step number of training without paraphrased samples as a reference and ensure that all methods, regardless of data augmentation or the addition of paraphrased texts, are training for the same number of steps to ensure a fair comparison. All experiments are conducted using NVIDIA A100 GPU.

For the verification experiment in Section 3.3, we prepend one example to align the answer format.

## B.2 Formatting-Based Data Augmentation Specifications

We include in the following all variations of data augmentations we used. For the random space insertion augmentation, we randomly insert an additional space adjacent to spaces in documents with a probability of 0.2.

Wrapping Augmentation
'Eden Benitez was raised in Santa Clarita.' "Eden Benitez was raised in Santa Clarita." *Eden Benitez was raised in Santa Clarita.* **Eden Benitez was raised in Santa Clarita.** ***Eden Benitez was raised in Santa Clarita.*** ==Eden Benitez was raised in Santa Clarita.== <Eden Benitez was raised in Santa Clarita.> (Eden Benitez was raised in Santa Clarita.)
Left Padding Augmented Examples
<space>Eden Benitez was raised in Santa Clarita. <space><space>Eden Benitez was raised in Santa Clarita. <tab>Eden Benitez was raised in Santa Clarita. #<space>Eden Benitez was raised in Santa Clarita. ##<space>Eden Benitez was raised in Santa Clarita. ###<space>Eden Benitez was raised in Santa Clarita.
Two Augmented Examples by Our Random Space Insertion Described Above
Eden Benitez was raised in Santa Clarita. Eden Benitez was raised in Santa Clarita.

## B.3 Biography Data Templates

We include in the following example data templates for our synthesized biography dataset, which contains 1000 biography profile. For each biography profile, we include three paraphrases of biography entries for continued pre-training, 1 QA pair per

attribute for instruction tuning, and another 5 QA pairs per attribute for evaluation. For the verification experiment in Section 3.3, we include 5 unseen paraphrased documents.

### CPT Training Data

- Eden Benitez was born on January 18, 1959. He was from Santa Clarita. He graduated from University of Wisconsin, Madison. His major was Marketing. He worked for General Dynamics.

### Paraphrase CPT Training Data

- Eden Benitez was born on January 18, 1959. He was from Santa Clarita. He graduated from University of Wisconsin, Madison. His major was Marketing. He worked for General Dynamics.
- Eden Benitez completed his education at University of Wisconsin, Madison. His field of study was Marketing. He was employed at General Dynamics. His place of origin was Santa Clarita. He entered the world on January 18, 1959.
- Eden Benitez majored in Marketing. He developed his career at General Dynamics. His life began on January 18, 1959. He attended University of Wisconsin, Madison. He came from Santa Clarita.

### Instruction Tuning QA Pairs

- Question: When was Eden Benitez born? Answer: January 18, 1959
- Question: Which university did Eden Benitez graduate from? Answer: University of Wisconsin, Madison
- Question: Which company did Eden Benitez work for? Answer: General Dynamics
- Question: Where was Eden Benitez from? Answer: Santa Clarita
- Question: What was Eden Benitez’s major? Answer: Marketing

### Unseen Paraphrased Documents

- Eden Benitez came into this world on January 18, 1959. His hometown is Santa Clarita. He pursued a degree in Marketing. He finished his education at University of Wisconsin, Madison. He had a job at General Dynamics.



- Eden Benitez originated from Santa Clarita. His birth date is January 18, 1959. He graduated from University of Wisconsin, Madison. He dedicated his studies to Marketing. He found employment at General Dynamics.
- Eden Benitez was brought into the world on January 18, 1959. He was raised in Santa Clarita. He pursued a Marketing major. He obtained his degree from University of Wisconsin, Madison. He worked at General Dynamics.
- Eden Benitez studied Marketing. He had a professional role at General Dynamics. He was born on January 18, 1959. He hailed from Santa Clarita. He received education at University of Wisconsin, Madison.
- Eden Benitez earned his degree from University of Wisconsin, Madison. He was majoring in Marketing. He held a position at General Dynamics. The birth date of him is January 18, 1959. He was a native of Santa Clarita.

#### **Evaluation QA Pairs**

- Question: When did Eden Benitez come into this world? Answer: January 18, 1959
- Question: What was Eden Benitez's birth date? Answer: January 18, 1959
- Question: When was Eden Benitez brought into the world? Answer: January 18, 1959
- Question: When did Eden Benitez first open his eyes? Answer: January 18, 1959
- Question: What was the birth date of Eden Benitez? Answer: January 18, 1959
- Question: Which university did Eden Benitez finish his education at? Answer: University of Wisconsin, Madison
- Question: Which university did Eden Benitez complete his degree program at? Answer: University of Wisconsin, Madison
- Question: Which university did Eden Benitez obtain his degree from? Answer: University of Wisconsin, Madison
- Question: Which university did Eden Benitez receive education at? Answer: University of Wisconsin, Madison
- Question: Which university did Eden Benitez earn his degree from? Answer: University of Wisconsin, Madison
- Question: Which company did Eden Benitez have a job at? Answer: General Dynamics
- Question: Which company did Eden Benitez find employment at? Answer: General Dynamics
- Question: Which company did Eden Benitez work at? Answer: General Dynamics
- Question: Which company did Eden Benitez have a professional role at? Answer: General Dynamics
- Question: Which company did Eden Benitez hold a position at? Answer: General Dynamics
- Question: Where was Eden Benitez's hometown? Answer: Santa Clarita
- Question: Where did Eden Benitez originate from? Answer: Santa Clarita
- Question: Where was Eden Benitez raised? Answer: Santa Clarita
- Question: Where did Eden Benitez hail from? Answer: Santa Clarita
- Question: Where was Eden Benitez a native of? Answer: Santa Clarita
- Question: What major did Eden Benitez pursue a degree in? Answer: Marketing
- Question: What major did Eden Benitez dedicate his studies to? Answer: Marketing
- Question: What major did Eden Benitez work toward earning a degree in? Answer: Marketing
- Question: What major did Eden Benitez study? Answer: Marketing
- Question: What major was Eden Benitez majoring in? Answer: Marketing

#### **B.4 Dataset Details**

Our biography dataset contains 1000 biography profile. For each biography profile, we include three paraphrases of biography entries for continued pre-training, 1 QA pair per attribute for instruction tuning, and another 5 QA pairs per attribute for evaluation. For the verification experiment in Section 3.3, we include 5 unseen paraphrased documents. The statistics of the Wiki2023 dataset is reported in [Jiang et al. \(2024\)](#).

The Wiki2023 dataset is publicly available online, so it is within their intended use. The biography dataset we generated following [Allen-Zhu and Li \(2024\)](#) is intended solely for research purposes and the attributes used are generated using ChatGPT. Both datasets used are in English.

The Wiki2023 dataset is from Wikipedia and the biography dataset is fully synthetic. So there should be no need to anonymize the data.