
Latent Paraphrasing: Perturbation on Layers Improves Knowledge Injection in Language Models

Minki Kang^{1,2}

Sung Ju Hwang²

Gibbeum Lee¹

Jaewoong Cho¹

¹KRAFTON, ²KAIST

zzxc1133@krafton.com, sjhwang@kaist.ac.kr, {pirensisco, jwcho}@krafton.com

Abstract

As Large Language Models (LLMs) are increasingly deployed in specialized domains with continuously evolving knowledge, the need for timely and precise knowledge injection has become essential. Fine-tuning with paraphrased data is a common approach to enhance knowledge injection, yet it faces two significant challenges: high computational costs due to repetitive external model usage and limited sample diversity. To this end, we introduce LaPael, a latent-level paraphrasing method that applies input-dependent noise to early LLM layers. This approach enables diverse and semantically consistent augmentations directly within the model. Furthermore, it eliminates the recurring costs of paraphrase generation for each knowledge update. Our extensive experiments on question-answering benchmarks demonstrate that LaPael improves knowledge injection over standard fine-tuning and existing noise-based approaches. Additionally, combining LaPael with data-level paraphrasing further enhances performance.

1 Introduction

Pre-trained Large Language Models (LLMs) encode extensive factual information from their training data, enabling them to answer factoid questions such as “Who is the director of Dune: Part Two?” [4, 32]. However, knowledge in LLMs is static, which can lead to outdated information as real-world knowledge evolves. Additionally, LLMs often lack specificity for specialized or private domains. To address this, it is common practice to fine-tune LLMs with updated or domain-specific documents, keeping the model’s knowledge up-to-date and enhancing expertise in particular domains [14, 17, 19].

However, does fine-tuning LLMs on a single document allow them to fully internalize its knowledge? Even in pre-training, Kandpal et al. [20] found that LLMs cannot perfectly learn all the information in the training data, particularly long-tail knowledge that appears rarely or only once. Existing work [33] has shown that this issue persists with fine-tuning and suggested that data augmentation, such as paraphrasing, is a simple yet effective way to enhance knowledge injection. As shown in Figure 1, fine-tuning with paraphrases enhances knowledge injection, as evidenced by improved Question-Answering (QA) task performance.




Figure 1: Effect of paraphrasing data in knowledge injection.

While data-augmented approach via paraphrasing is effective for knowledge learning, it has two main limitations: (1) **High computational cost**: Generating high-quality paraphrases requires significant computational resources. As shown in Figure 2, paraphrasing models such as LLMs [5, 7, 11, 58] need to repeatedly generate paraphrases for each document with the new incoming knowledge. This leads to higher costs as the number of documents being learned continually increases; and (2) **Limited**

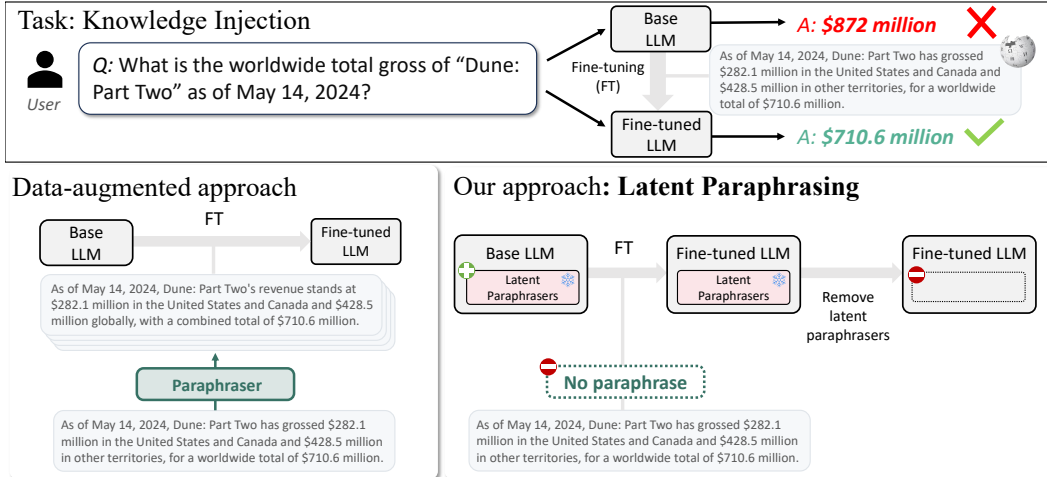


Figure 2: **A conceptual illustration of the proposed approach.** On the left, we show the existing method of knowledge injection by paraphrasing each document for data-level augmentation. On the right, we present the conceptual illustration of LaPael with *trained* latent paraphrasers. Unlike the method on the left, LaPael can eliminate the need for users to repeatedly paraphrase using LLMs once latent paraphrasers are trained.

diversity in augmented data: Although LLMs can produce varying high-quality paraphrases by sampling from the generative distribution, the diversity of the generated text is limited, resulting in a narrow range of augmented samples at the discrete data level. One way to overcome these issues is to introduce noise into the token embedding. However, existing works [16, 57] do not consider the text semantics when they perturb the latent features of LLMs with randomly generated noise.

To address these issues, we take a distinct approach using an input-dependent noise generator named “latent paraphraser” learned from the paraphrases. Specifically, this function perturbs early layers to augment LLMs at the latent level while preserving the meaning of the text. To optimize the latent paraphraser, we start by generating paraphrases of the documents. Then, we train the latent paraphrasers to ensure that the latent distribution of the LLMs with the original sentence is close to the latent distribution with the paraphrased sentences. Once training is done, we can transfer the latent paraphrasers to the documents from any domain that contains new knowledge. We refer to our method as **Latent Paraphrasing of Language Models (LaPael)**, as it learns the paraphrasing of text data at the latent level.

We validate our approach on diverse question-answering benchmark datasets [38, 27, 51] designed to evaluate knowledge injection. These benchmarks involve fine-tuning LLMs on documents that contain the knowledge required to answer the questions in the datasets. Our results show that LaPael significantly improves knowledge injection performance compared to standard fine-tuning. Moreover, LaPael outperforms fine-tuning with paraphrases, demonstrating that LaPael alone is sufficient for data augmentation in knowledge injection scenarios, as illustrated in Figure 2. As shown in Figure 1, we further find that using LaPael in combination with paraphrases further enhances performance, providing complementary benefits to data-level augmentations. Finally, LaPael surpasses existing noise baselines [16, 57], highlighting the importance of learning noise for effective augmentations.

Our contributions are as follows:

- We introduce **LaPael**, a new method that applies learned perturbations to the layers of LLMs to enhance knowledge injection, addressing the limitations of data augmentations and noise baselines.
- We validate LaPael using diverse question-answering benchmark datasets, demonstrating a significant improvement in knowledge injection performance compared to standard fine-tuning.
- Our results show that LaPael not only outperforms fine-tuning with paraphrases but also complements it, providing additional benefits when used together, surpassing the performance of existing latent noise-based methods.

2 Related Work

Knowledge of Large Language Models Large Language Models (LLMs) store vast amounts of factual knowledge in their pre-trained parameters [36, 44]. The straightforward way to extract the

knowledge of LLMs is to ask the question that requires factual knowledge [43, 58]. Through asking questions, Kandpal et al. [20] have found that LLMs cannot perfectly memorize the entire knowledge in the pre-training corpora, especially for knowledge that appears rarely or only once. To make LLMs answer the question requires under-represented or new knowledge, previous works have clustered into two different solutions. The first one is retrieval-augmented methods [26, 39, 42] that retrieve knowledge from an external knowledge base and input the retrieved knowledge alongside the question into LLMs. The second one is fine-tuning [12, 17] where the parameters of pre-trained LLMs are continually updated by fine-tuned on the document containing knowledge in an unsupervised way as in pre-training [37]. In our work, we focus on improving the fine-tuning-based solution, as storing new knowledge in the parameters of LLMs is efficient since we can reduce the length of the input prompt and do not need any extra module or memory in the deployment time [6].

Knowledge Injection in LLMs In this work, knowledge injection in LLMs denotes fine-tuning LLMs on the set of documents to inject new or under-represented knowledge into LLMs [33, 17], different from another task of injecting *symbolic knowledge* (e.g., knowledge graph) into LLMs [55, 54]. Among previous works, CaMeLS [14] has introduced a meta-learning method for learnable loss scaling function that improves knowledge injection. As a concurrent work, MAC [45] has proposed using the memory of amortized context is highly effective in a knowledge injection. However, both methods have drawbacks like high computational costs for bi-level optimization or the need for additional modules and memory. Recent works [33, 58] have shown that data augmentation which paraphrases the knowledge-containing sentences helps language models memorize knowledge in a more extractable format (e.g., asking questions) after knowledge injection. Furthermore, Jiang et al. [19] has shown that the instruction-tuned model is better at learning new knowledge. Compared to previous works, we focus on developing an alternative method to data augmentation that perturbs the latent representation of LLMs for better knowledge injection.

Data Augmentation and Latent Perturbation The usefulness of data augmentations for text data was empirically observed in the literature. For instance, EDA [52] has introduced simple data augmentation method which randomly deletes, swaps, replaces, and inserts the words. Other previous works [22, 5, 30] have utilized the trained LMs to augment the text data. Recently, Maini et al. [29] has shown that adding data rephrased by LLMs into the pre-training corpus improves the performance of LM pre-training. However, those methods require additional costs in the knowledge injection as it utilize the LLMs to rephrase the text. In contrast, the latent perturbations offer an orthogonal approach to improve the robustness of neural networks, complementing data augmentation. This technique has been employed in meta-learning and out-of-distribution generalization [24, 25, 40]. For instance, NEFTune [16] demonstrated that adding noise, randomly sampled from a uniform distribution, to token embedding layers improves instruction tuning performance. Expanding on the concept of latent perturbations, our work introduces a novel approach that *internalizes* the effects of text paraphrasing by identifying optimal latent perturbations through training a small neural network within the LLMs.

3 Problem Formulation

In this work, we follow the knowledge injection setting outlined by Ovadia et al. [33]. We are given three resources: (1) documents \mathcal{D}_K containing knowledge that we are interested to inject; (2) question & answering dataset $\mathcal{D}_{QA} = \{(\mathbf{q}^{(i)}, \mathbf{a}^{(i)})\}_{i=1}^n$ for verifying injected knowledge from \mathcal{D}_K ; and (3) a pre-trained Large Language Models (LLMs) $p_\theta(\cdot)$ parameterized by θ . Our objective is to find a transformation F that could enhance the knowledge about \mathcal{D}_{QA} :

$$\theta' = F(\theta, \mathcal{D}_K) \quad \text{s.t.} \quad \mathcal{S}(\theta', \mathcal{D}_{QA}) > \mathcal{S}(\theta, \mathcal{D}_{QA}), \quad (1)$$

where the score function \mathcal{S} is defined as:

$$\mathcal{S}(\theta, \mathcal{D}_{QA}) := \frac{\sum_{i=1}^n \mathbb{I}(f(p_\theta(\mathbf{q}^{(i)})) = \mathbf{a}^{(i)})}{n}, \quad (2)$$

and $\mathbb{I}(\cdot)$ and $f(\cdot)$ denote the indicator function and a decoding function that samples a sequence of tokens from p_θ , respectively.

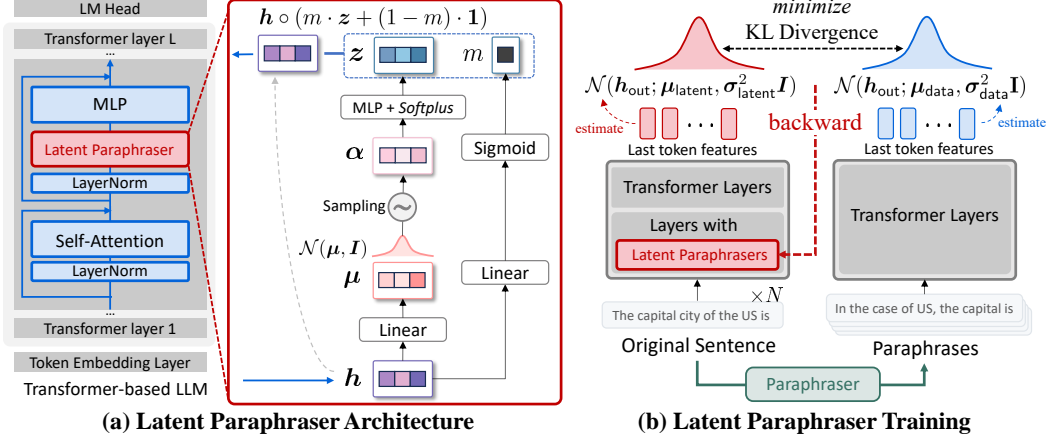


Figure 3: **(a) Illustration of the latent paraphraser.** The linear layer embeds each token’s latent feature h into μ . We then sample stochastic noise α from $\mathcal{N}(\mu, \mathbf{I})$ and apply a mask m_t to control the scale. **(b) Training pipeline of LaPael.** To train the latent paraphraser, we estimate the parameters of Gaussian distributions. We then minimize the KL divergence between these distributions to optimize the latent paraphrasers.

In general, a transformation F is a fine-tuning LLMs on documents in \mathcal{D}_K by optimizing θ to minimize the negative log-likelihood of each token in each document as follows [33]:

$$\theta^* = \arg \min_{\theta} \frac{1}{|\mathcal{D}_K|} \sum_{s \in \mathcal{D}_K} \left(\frac{1}{|s|} \sum_{t=1}^{|s|} -\log p_{\theta}(s_t | s_{<t}) \right), \quad (3)$$

where $|s|$ denotes the length of token sequence s .

4 Proposed Method

We propose Latent Paraphrasing of Language Models (LaPael), a framework that perturbs the latent feature of LLMs, to achieve the equivalent effect of data augmentation at the latent level. Knowledge injection using LaP consists of the following four processes: paraphrasing the set of documents to make the paraphrased data (Section 4.1), training the latent paraphrasers with paraphrased data (Section 4.2), fine-tuning LLMs with the trained latent paraphrasers on \mathcal{D}_K and evaluate the injected knowledge of LLMs on \mathcal{D}_{QA} (Section 4.3).

4.1 Data Augmentation: Paraphrasing

To train the latent paraphrasers, we need a distinct set of training data $\mathcal{D}_{\text{train}} = \{s^{(i)}\}_{i=1}^N$ which consists of documents having different knowledge with \mathcal{D}_K . As a preliminary, we formulate the paraphrasing of the text in terms of the **knowledge equivalence**, which is a narrower concept than semantic equivalence [23] where two different sentences can contain the same knowledge. We consider that each sentence s in $\mathcal{D}_{\text{train}}$ can be decomposed into words for the object (entity or attribute) of the knowledge (\mathbf{y}) and others (\mathbf{x}) where both are the sequence of tokens. For instance, given the sentence “The capital of the United States is Washington D.C.”,

$$\mathbf{x} = \text{“The capital of the United States is”}; \quad \mathbf{y} = \text{“Washington D.C.”},$$

represent the knowledge (United States, capital, Washington D.C.). Then, we paraphrase a sentence $s = (\mathbf{x}, \mathbf{y})$ into a paraphrased sentence¹. For the above sentence, a paraphrased sentence can be

$$\mathbf{x}' = \text{“In the case of the United States, the designated capital city is”}$$

with the same \mathbf{y} , which is knowledge equivalent to (\mathbf{x}, \mathbf{y}) . For each knowledge \mathcal{K} , we assume that there is a set of the knowledge equivalent sentences $S(\mathcal{K})$ where $(\mathbf{x}, \mathbf{y}) \in S(\mathcal{K})$. We generate K paraphrased sentence via a LLM: $(\mathbf{x}_1, \mathbf{y}), \dots, (\mathbf{x}_K, \mathbf{y}) \sim p_{\text{LLM}}(\mathbf{x}' | \text{prompt}, \mathbf{x}, \mathbf{y})$. Then, we have the set of paraphrased data $\{(\mathbf{x}_k^{(i)}, \mathbf{y}^{(i)})\}_{k=1}^K\}_{i=1}^N$ of $\mathcal{D}_{\text{train}}$. We define $p(\mathbf{x}' | \mathbf{x}) := p_{\text{LLM}}(\mathbf{x}' | \text{prompt}, \mathbf{x}, \mathbf{y})$ which denotes the probability distribution of paraphrases given the original sentence.

¹One possible way is to prompt the LLM (e.g., gpt-3.5-turbo [31]) with instruction “For the following paragraph give me a paraphrase of the same in high-quality English language as in sentences on Wikipedia” [29]

4.2 Introducing Latent Paraphraser

Latent Paraphraser We introduce a latent paraphraser within a transformer layer [50], which augments a latent feature and is expected to paraphrase the given input text within the latent space. As illustrated in Figure 3(a), within the transformer architecture, we insert this new layer just before the Multi-layer Perceptron (MLP), using the output from the second LayerNorm as its input.

Let $\mathbf{h} \in \mathbb{R}^d$ denote the latent feature after the second LayerNorm. The latent paraphraser, denoted by $g_\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and parameterized by ϕ , augments the latent feature as follows:

$$\mathbf{h} \circ g_\phi(\mathbf{h}), \quad (4)$$

where \circ is the element-wise multiplication. The function $g_\phi(\mathbf{h})$ is given by:

$$g_\phi(\mathbf{h}) = (1 - m) \cdot \mathbf{1} + m \cdot \mathbf{z}, \quad (5)$$

with $\mathbf{z} \in \mathbb{R}^d$ and $m \in [0, 1]$ representing a noise vector and a learnable mask, respectively.

The noise vector \mathbf{z} is generated by

$$\mathbf{z} = \text{softplus}(\text{MLP}_z(\boldsymbol{\alpha})), \quad \boldsymbol{\alpha} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}), \quad \boldsymbol{\mu} = \mathbf{W}_\mu \mathbf{h} + \mathbf{b}_\mu, \quad (6)$$

where MLP_z is a 2-layers MLP. We use the reparameterization trick [21] to enable the back-propagation through the sampling from the Gaussian distribution: $\boldsymbol{\alpha} = \boldsymbol{\mu} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$.

To modulate the scale of perturbation for individual tokens, we employ a learnable mask. It is important as too much noise on key tokens (e.g., United States) might hurt the semantics of the sequence. For learnable binary mask, we use concrete distribution to approximate the sampling discrete random variable from a Bernoulli distribution using continuous relaxation [8] as follows:

$$m = \text{sigmoid}\left(\frac{1}{\tau} \log(u) + \log(1 - u) + \tilde{m}\right), \quad \tilde{m} = \mathbf{W}_m \mathbf{h} + b_m, \quad (7)$$

where $u \sim \text{Unif}(0, 1)$, τ is temperature, and m is mask value in scalar.

Training Then, how do we train the latent paraphrasers to approximate optimal perturbation functions for estimating the distribution of the paraphrased text? We employ the dataset with paraphrases $\{(\mathbf{x}_k^{(i)}, \mathbf{y}^{(i)})\}_{k=1}^K\}_{i=1}^N$ generated in Section 4.1. Our objective is to match two distributions for each transformer layer:

1. the distribution of transformer layer output feature for the last token \mathbf{h}_{out} without the latent paraphraser given the data perturbation distribution $p(\mathbf{x}'|\mathbf{x})$ from Section 4.1:

$$p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x}) = \int p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x}')p(\mathbf{x}'|\mathbf{x})d\mathbf{x}'; \quad (8)$$

2. the distribution of output feature for the last token \mathbf{h}_{out} with the latent paraphraser given \mathbf{x} , $p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x})$. As a latent paraphraser outputs stochastic noise, we can formulate the probabilistic distribution $p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x})$ as follows:

$$p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x}) = \int p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x}, \mathbf{z})p_{\theta, \phi}(\mathbf{z}|\mathbf{x})d\mathbf{z}, \quad (9)$$

where $p_{\theta, \phi}(\mathbf{z}|\mathbf{x})$ is the distribution for noise from the latent paraphraser in Equation (6).

We make the simplistic parametric assumption that both distributions are Gaussian:

$$p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x}) \sim \mathcal{N}(\mathbf{h}_{\text{out}}; \boldsymbol{\mu}_{\text{data}}, \boldsymbol{\sigma}_{\text{data}}^2 \mathbf{I}); \quad p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x}) \sim \mathcal{N}(\mathbf{h}_{\text{out}}; \boldsymbol{\mu}_{\text{latent}}, \boldsymbol{\sigma}_{\text{latent}}^2 \mathbf{I}). \quad (10)$$

To train latent paraphrasers, we minimize the symmetric Kullback-Leibler (KL) divergence between two estimated Gaussian distributions of each layer as follows:

$$\mathcal{L}_{\text{KL}}(\mathbf{x}) = \frac{1}{2}(\hat{D}_{\text{KL}}(p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x})\|p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x})) + \hat{D}_{\text{KL}}(p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x})\|p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x}))), \quad (11)$$

$$\hat{D}_{\text{KL}}(p_\theta(\mathbf{h}_{\text{out}}|\mathbf{x})\|p_{\theta, \phi}(\mathbf{h}_{\text{out}}|\mathbf{x})) = \log\left(\frac{\hat{\boldsymbol{\sigma}}_{\text{latent}}}{\hat{\boldsymbol{\sigma}}_{\text{data}}}\right) + \frac{\hat{\boldsymbol{\sigma}}_{\text{data}}^2 + (\hat{\boldsymbol{\mu}}_{\text{data}} - \hat{\boldsymbol{\mu}}_{\text{latent}})^2}{2\hat{\boldsymbol{\sigma}}_{\text{latent}}^2} - \frac{1}{2}. \quad (12)$$

We employ a Monte Carlo sampling approach to estimate the parameters of Gaussian distributions. We generate N samples $\mathbf{h}_{\text{latent}}^{(1)}, \dots, \mathbf{h}_{\text{latent}}^{(N)}$ from the distribution $p_{\theta, \phi}(\mathbf{h}_{\text{out}} | \mathbf{x})$. Then, we estimate the empirical mean and standard deviation from the samples as follows:

$$\hat{\boldsymbol{\mu}}_{\text{latent}} = \frac{1}{N} \sum_{i=1}^N \mathbf{h}_{\text{latent}}^{(i)}, \quad \hat{\boldsymbol{\sigma}}_{\text{latent}} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (\mathbf{h}_{\text{latent}}^{(i)} - \hat{\boldsymbol{\mu}}_{\text{latent}})^2}, \quad (13)$$

and we use K paraphrases $\mathbf{x}_1, \dots, \mathbf{x}_K$ to obtain K samples $\mathbf{h}_{\text{data}}^{(1)}, \dots, \mathbf{h}_{\text{data}}^{(K)}$ from the distribution $p_{\theta}(\mathbf{h}_{\text{out}} | \mathbf{x})$. Then we estimate the parameters in the same way:

$$\hat{\boldsymbol{\mu}}_{\text{data}} = \frac{1}{K} \sum_{k=1}^K \mathbf{h}_{\text{data}}^{(k)}, \quad \hat{\boldsymbol{\sigma}}_{\text{data}} = \sqrt{\frac{1}{K-1} \sum_{k=1}^K (\mathbf{h}_{\text{data}}^{(k)} - \hat{\boldsymbol{\mu}}_{\text{data}})^2}. \quad (14)$$

We further use the auxiliary loss for mask training, with the sequence length of T as follows:

$$\mathcal{L}_{\text{mask}}(\mathbf{x}) = \sum_{t=1}^T (|\text{sigmoid}(\tilde{m}_t) - r \cdot T| + |\text{sigmoid}(\tilde{m}_t) - \tilde{m}_t|), \quad (15)$$

where \tilde{m}_T is defined in Equation (7), $r \in [0, 1]$ is the mask ratio that controls the number of masks and \tilde{m}_t is the gold mask where $\tilde{m}_t = 0$ for tokens that correspond to the named entity.

To sum up, we optimize the latent paraphraser parameter ϕ by minimizing the following loss:

$$\phi^* = \arg \min_{\phi} \sum_{\mathbf{x} \in \mathcal{D}_{\text{train}}} (\mathcal{L}_{\text{KL}}(\mathbf{x}) + \mathcal{L}_{\text{mask}}(\mathbf{x})). \quad (16)$$

See Figure 3(b) for an illustration of the training process for the latent paraphraser.

4.3 Fine-tuning the LLM with the Trained Latent Paraphrasers

We fine-tune the LLM on documents containing knowledge to be injected (\mathcal{D}_{K}) as in Equation (3). We use the trained latent paraphraser parameterized by ϕ^* during LLM fine-tuning as follows:

$$\theta^* = \arg \min_{\theta} \frac{1}{|\mathcal{D}_{\text{K}}|} \sum_{\mathbf{s} \in \mathcal{D}_{\text{K}}} \left(\frac{1}{|\mathbf{s}|} \sum_{t=1}^{|\mathbf{s}|} \left(\frac{1}{N} \sum_{j=1}^N -\log p_{\theta, \phi^*}(\mathbf{s}_t | \mathbf{z}_t^{(j)}, \mathbf{s}_{<t}) p_{\theta, \phi^*}(\mathbf{z}_t^{(j)} | \mathbf{s}_{<t}) \right) \right), \quad (17)$$

where we sample N noise $\mathbf{z}^{(j)}$ by sampling multiple α from Gaussian distribution as defined in Equation (6). Then, we evaluate the knowledge injected in LLMs by measuring $\mathcal{S}(\theta^*, \mathcal{D}_{\text{QA}})$ as defined in Equation (2).

5 Experiments

In experiments, we validate the effectiveness of the proposed method, LaPael, in injecting new or under-represented knowledge into Large Language Models (LLMs).

5.1 Experimental Setting

5.1.1 Datasets

To follow the experimental setup in Section 3, we need (1) documents containing knowledge \mathcal{D}_{K} and (2) associated QA datasets \mathcal{D}_{QA} . We mainly use the test split of three QA datasets: SQuAD [38], StreamingQA [27], and ArchivalQA [51] for the source of \mathcal{D}_{K} and \mathcal{D}_{QA} in our main experiments. These datasets, previously used in Hu et al. [14], consist of documents paired with their corresponding QAs, making them well-suited to our experimental setup. While the questions in these datasets are of decent quality, a significant limitation lies in the documents provided. These documents are likely to have been seen by LLMs during pre-training, making it difficult to accurately assess the performance of methods on injecting new knowledge.

Table 1: **Data Example.** Example data from SQuAD and StreamingQA dataset we used in experiments. Words in the yellow background indicate the answer to the question. More examples are in Table 12 of the Appendix.

Question	Raw Document	Synthetic Document
Who was the Super Bowl 50 MVP? <i>(from SQuAD)</i>	(...) Denver linebacker Von Miller was named Super Bowl MVP, recording five solo tackles, 200bd sacks, and two forced fumbles.	The Super Bowl 50 MVP was Von Miller.
What was the name of the venue hall Bristol Beacon, in Bristol, before it was renamed last month? <i>(from StreamingQA)</i>	(...) Colston Hall, which was named after the 17th-century slave trader Edward Colston, will from now on be known as Bristol Beacon following a public consultation. Bristol attracted headlines around (...)	Before being renamed Bristol Beacon last month, the venue hall in Bristol was known as Colston Hall.

Table 2: Experimental results on datasets with **synthetic documents**. *trained with n sents* means that latent paraphrasers are trained with the dataset containing n sentences. For ours, we report the average performance of three runs. † denotes the method that uses 10 times more additional data (paraphrases).

Method	SQuAD-syn			StreamingQA-syn			ArchivalQA-syn		
	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1
No Injection	13.10	22.91	21.09	16.39	26.30	23.71	13.50	25.07	22.12
Fine-Tuning	66.30	79.32	76.11	82.08	88.98	88.29	62.60	79.51	76.16
Fine-Tuning (seq.)	67.60	80.30	77.39	77.95	86.36	85.23	56.30	79.17	74.12
FreeLB [57]	70.70	82.41	79.67	82.24	89.48	88.56	63.20	81.30	77.67
NEFTune [16]	68.30	80.93	77.91	81.47	88.66	87.77	61.90	78.90	75.81
Ours trained w/ 50 sents.	70.77	84.96	81.66	86.16	93.01	92.12	68.37	86.24	82.67
Ours trained w/ 1k sents.	72.47	87.93	84.50	84.48	92.42	91.33	68.37	88.99	84.75
Fine-Tuning (+ para.)†	68.50	85.12	80.51	85.45	93.67	92.32	64.90	85.92	81.24

To mitigate this issue, we incorporate two datasets with synthetic QAs – Films 2024 and Events 2024. These are QA datasets generated from raw Wikipedia articles under the [2024 films](#) category and from US events in May, June, and July 2024, in the [2024 events in the United States](#) category. We generated question-answer pairs from these documents using GPT-4o following methods from previous works [19, 33]. Since the documents used to generate these datasets were not seen by the LLMs during pre-training, we can better evaluate the effectiveness of each method for knowledge injection especially on new knowledge.

Datasets with Synthetic Documents The raw documents from datasets are unsuitable for precisely measuring the knowledge injection performance. Specifically, fine-tuning LLMs on a document does not always ensure that LLMs can answer the associated questions, due to the reversal curse [3]. Moreover, documents often contain irrelevant knowledge that may hinder the accurate assessment of knowledge injection [14].

To address these issues, we conduct evaluations under the setting of synthetic documents. For generating synthetic documents, we construct \mathcal{D}_K by rephrasing each question and answer in \mathcal{D}_{QA} using GPT-4-turbo [32], ensuring that fine-tuning on these synthetic documents guarantee that LLMs become answerable to the associated questions. Examples of questions, synthetic, and raw documents are shown in Table 1. To make a difference, we denote the dataset under the synthetic document setting with the suffix ‘-syn’ and the raw document setting with the suffix ‘-raw’.

Datasets for Training Latent Paraphrasers For training our latent paraphrasers, the set of training data \mathcal{D}_{train} is required in addition to \mathcal{D}_K . Therefore, we use GPT-3.5-turbo [31] to generate the set of synthetic sentences from the subset of a training split of each QA dataset, where each sentence must be with the answer to questions, following the sentence format in Section 4.1.

5.1.2 Experimental Details

Baselines We compare our LaPael against several baselines. All models are fine-tuned on the documents in \mathcal{D}_K unless explicitly stated otherwise. (1) **No Injection.** We use the pre-trained LLM without any fine-tuning. (2) **Fine-Tuning.** We fine-tune the LLM on \mathcal{D}_K . (3) **Fine-Tuning (seq).** We first fine-tune the LLM on the paraphrased documents of \mathcal{D}_{train} . Then, we fine-tune the LLM on \mathcal{D}_K . (4) **Fine-Tuning (+ para).** We fine-tune LLM on the original and paraphrased documents of \mathcal{D}_K . (5) **FreeLB [57].** We add trained adversarial noise to the token embedding while fine-tuning. (6) **NEFTune[16].** We add random uniform noise to the token embedding while fine-tuning. (7) **LaPael (ours).** We train the latent paraphrasers on \mathcal{D}_{train} and then fine-tune the model on \mathcal{D}_K .

Table 3: Experimental results on datasets with **raw documents**. For Ours, we use the latent paraphraser used in the SQuAD-syn experiment. Rec. denotes recall.

Method	SQuAD-raw			StreamingQA-raw			Films 2024-raw			Events 2024-raw		
	EM	Rec.	F1	EM	Rec.	F1	EM	Rec.	F1	EM	Rec.	F1
No Injection	9.98	23.44	20.62	16.22	29.85	27.81	1.93	10.21	10.27	1.73	17.94	17.40
Fine-Tuning	16.65	35.40	29.73	19.04	35.88	32.92	13.39	30.03	28.84	10.98	43.76	39.62
FreeLB [57]	17.04	36.78	30.36	20.72	37.19	34.04	15.47	33.69	31.85	14.68	46.05	41.85
NEFTune [16]	17.45	37.49	31.11	20.18	36.98	33.85	15.93	33.73	32.38	15.38	48.14	43.84
Ours	18.96	43.10	34.65	21.62	39.38	35.32	16.29	35.04	32.56	15.26	56.70	46.45

Table 4: Experimental results on **cross-domain transfer** experiments. For ours, (X \rightarrow) denotes that latent paraphrasers are trained on $\mathcal{D}_{\text{train}}$ from the X dataset. Rec. denotes recall.

Method	SQuAD-syn			StreamingQA-syn			NovelQA-syn			MedMCQA-syn		
	EM	Rec.	F1	EM	Rec.	F1	EM	Rec.	F1	EM	Rec.	F1
No Injection	13.10	22.91	21.09	16.39	26.30	23.71	9.17	18.21	16.05	39.00	48.68	47.82
Fine-Tuning	58.30	68.59	66.35	74.73	82.34	81.21	52.92	66.30	63.62	56.10	62.37	62.03
FreeLB [57]	70.70	82.41	79.67	82.24	89.48	88.56	55.42	67.39	64.80	57.90	63.17	62.81
NEFTune [16]	68.30	80.93	77.91	81.47	88.66	87.77	51.67	65.14	62.25	56.30	62.57	62.09
Ours (SQuAD \rightarrow)	72.50	89.38	85.34	84.38	93.44	92.17	54.17	69.40	65.72	63.70	68.28	67.98
Ours (StreamingQA \rightarrow)	72.80	89.65	85.90	84.06	93.73	91.90	54.58	72.58	68.15	63.20	68.02	67.79

Training & Inference We mainly use Vicuna-7b-v1.5 [56] for fine-tuning, which is the instruction-tuned version of Llama-2-7b [48] for our experiments. We fine-tune LLMs for 12 epochs with a learning rate of 0.00005 and step learning rate scheduler where we decay a learning rate by 0.85 by every 4 epochs. For inference, we use in-context learning with 5 examples by prompting the 5 examples in the prompt [4]. To measure QA accuracy, we use Exact Match (EM), Recall (Rec.), and F1 score. More details on the experimental setting are provided in the Appendix C.

5.2 Experimental Results

Experiments with Synthetic Documents In Table 2, we present the experimental results for the synthetic documents setting. Fine-tuning does improve the QA performance of LLMs, but it does not lead to near-perfect scores even though the synthetic document contains the necessary knowledge for answering the questions, as shown in Table 1.

Our experiments show that paraphrasing documents for fine-tuning consistently improves QA performance across all three benchmarks. Notably, LaPael demonstrates performance comparable to fine-tuning with paraphrases on StreamingQA and even outperforms it on two other benchmarks. These findings suggest that the latent paraphrasers learn an effective noise distribution that aids knowledge injection without additional data augmentation.

We also compared LaPael with two other noise-based methods, FreeLB [57] and NEFTune [16], to validate that the latent-level noise generated by latent paraphrasers is more effective. As shown in Table 2, LaPael outperforms these baselines, confirming the strength of our approach.

Experiments with Raw Documents While our method has proven effective for knowledge injection with synthetic documents, it is important to evaluate its performance on raw documents, which represent a more realistic data format. To demonstrate the applicability of our method to real-world data, we conducted experiments in which we fine-tuned LLMs on raw documents for each dataset, using latent paraphrasers trained on $\mathcal{D}_{\text{train}}$ from SQuAD-syn.

As shown in Table 3, our method outperforms both fine-tuning and noise-based baselines in the context of knowledge injection with raw documents. Considering that the latent paraphrasers were trained on synthetic sentences from $\mathcal{D}_{\text{train}}$, these results demonstrate their effectiveness on documents with a different format than those used in training.

Cross-domain Transfer Once trained, the latent paraphrasers can be applied to fine-tune LLMs on documents from any domain. To demonstrate this, we conducted cross-domain transfer experiments. Specifically, we trained latent paraphrasers on $\mathcal{D}_{\text{train}}$ from a source domain (e.g., SQuAD) and fine-tuned LLMs with the trained latent paraphrasers on \mathcal{D}_{K} from a target domain (e.g., StreamingQA).

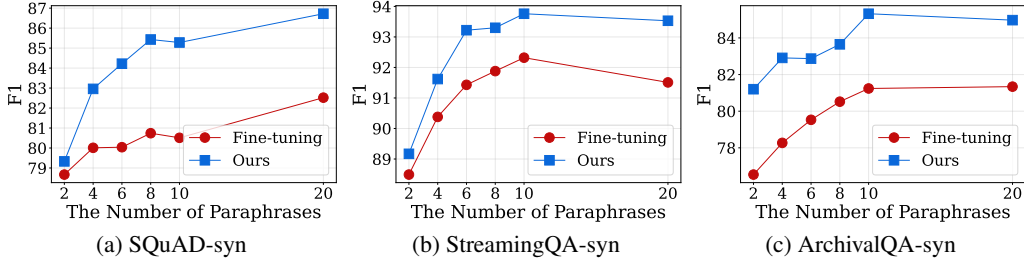


Figure 4: **Effect of the Number of Paraphrases.** Each plot shows the relationship between the number of paraphrases (x-axis) and F1 scores (y-axis) in knowledge injection. The F1 scores of both standard fine-tuning and our method improve as the number of paraphrases increases.

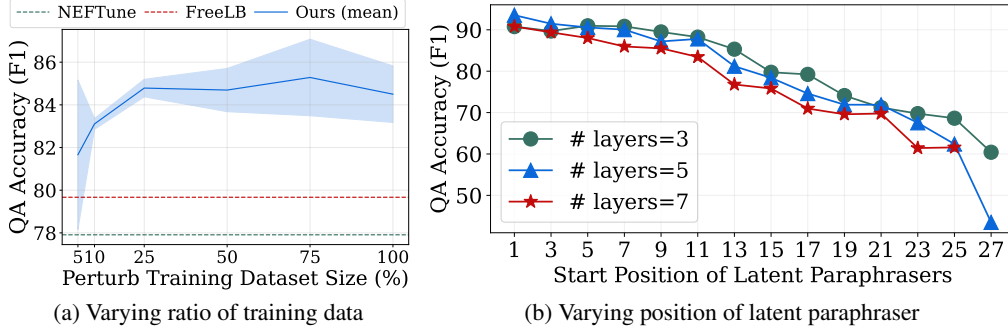


Figure 5: (a) We conduct experiments varying the size of $\mathcal{D}_{\text{train}}$ on SQuAD-syn, where 100% indicates 1,000 documents. We report mean and std. over three runs. (b) We conduct experiments on StreamingQA-syn varying the start position of latent paraphrasers where ‘# layers’ denotes the number of latent paraphrasers.

As shown in Table 4, our method successfully transfers across domains, with the latent paraphrasers enhancing the performance of the knowledge injection on NovelQA and MedMCQA—two domains distinct from the source (see Appendix C.1 for details on these datasets). Even though both domains contain specialized entities, our method consistently outperforms standard fine-tuning and other noise-based baselines.

Combining LaPael and Paraphrases Paraphrasing documents in \mathcal{D}_K has been shown to improve knowledge injection performance, as seen in Table 2. While LaPael significantly improves performance without requiring paraphrases, it is valuable to consider the effect of combining paraphrases with the latent perturbations from LaPael. As illustrated in Figure 4, LaPael consistently outperforms standard fine-tuning, showing that LaPael provides advantages over data-level augmentations.

5.3 Ablation Studies

Effects of the Size of $\mathcal{D}_{\text{train}}$ LaPael needs additional data $\mathcal{D}_{\text{train}}$ for training latent paraphrasers. Although only a small amount of data is required, it might be unclear how much is needed to make the latent paraphrasers learn the useful noise distribution. As shown in Figure 5a, LaPael works well even with **50 sentences** for $\mathcal{D}_{\text{train}}$, while increasing the size of $\mathcal{D}_{\text{train}}$ ensures a steady performance improvement for LaPael.

Effects of the Position of Latent Paraphrasers Our latent paraphrasers can be inserted into any layer of the LLMs. The possible question is which position and how many layers are optimal for latent paraphrasers to effectively learn noise for knowledge injection. To answer this, we analyzed the position and number of latent paraphrasers.

In Figure 5b, we show the QA accuracy results, varying the start position and number of latent paraphrasers. The first layer is the closest layer to the input layer, and “start position 1” with “# layers = 3” means we insert the latent paraphrasers into the first, second, and third layers of the LLM. Results show that inserting three latent paraphrasers into the early layers of the LLM is effective. This is consistent with findings in previous works [16, 57, 25] where using noisy token embeddings (the lowest layer) enhanced the generalization in LLMs. Furthermore, in Table 5, we empirically show that positioning the latent paraphraser before the MLP layer within each transformer layer is the most effective choice over other positions.

Table 5: Analysis on the **Position** inside the Transformer layer.

StreamingQA	EM	Recall	F1
Before MLP	84.06	93.73	91.90
After MLP	73.81	82.58	81.02
Before Attn	80.55	87.58	86.49
After Attn	83.31	90.98	89.72
Token Embed.	86.21	91.79	91.05

Table 6: Ablation studies on **Modules** in latent paraphrasers.

StreamingQA	EM	Recall	F1
LaPael	84.06	93.73	91.90
w/o Mask	77.95	85.17	84.63
w/o Concrete	73.35	83.42	81.99
w/o Sampling	84.23	90.52	89.73
w/o KL loss	83.31	90.78	89.99

Table 7: Ablation studies on **Noise design** in latent paraphrasers.

StreamingQA	EM	Recall	F1
Learnable Mul.	84.06	93.73	91.90
Learnable Add.	73.05	83.23	81.70
Gaussian	83.46	90.32	89.54
Gaussian + mask	82.85	89.70	88.87
Uniform	79.48	87.17	86.15
Uniform + mask	74.43	81.09	80.26

Ablation Studies on Modules LaPael has many design choices concerning the latent paraphraser architecture, noise type, and training. We conducted extensive ablation studies to empirically verify each design choice and provide guidance for future work. In summary, as shown in Table 6, all design choices are important for building the most effective latent paraphraser. Specifically, we use a trainable mask m in Equation (7) to regulate the perturbation depending on each token, which is crucial, as the performance on StreamingQA drops significantly if we remove it from the latent paraphraser. Furthermore, using only the sigmoid function in Equation (7) instead of the concrete distribution also leads to much lower performance, as the mask is not properly trained. Regarding noise training, using deterministic noise instead of stochastic noise by removing the noise drawn from a Gaussian distribution in Equation (5) also decreases performance. Additionally, replacing the KL loss with Mean Squared Error loss between two means $\hat{\mu}_{\text{latent}}$ in Equation (13) and $\hat{\mu}_{\text{data}}$ in Equation (14) leads to a decrease in performance, confirming the importance of stochastic noise trained with KL loss.

Ablation Studies on Noise Distribution Should we train the latent paraphrasers to be effective, or can adding random noise in the early layers also be effective? Which is more important: the learnable mask or the learnable noise? To answer these questions, we conducted ablation studies on the choice of noise distribution. In Table 7, Learnable Add. denotes the model with the additive noise $h + g_{\phi}(h)$ instead of Equation (4) without softplus from Equation (6). Gaussian is the use of zero-mean Gaussian noise $\mathcal{N}(0, I)$ in Equation (6) without using MLP_z . Uniform is the use of noise drawn from the uniform distribution defined in NEFTune [16] instead of z in Equation (6).

As shown in Table 7, the learnable multiplicative noise described in Section 4.2 is the best design for noise distribution used in the latent paraphraser. To analyze the effect of the learnable mask, we also added the learnable mask to the Gaussian and Uniform noise settings and optimized only W_m and b_m in Equation (7) with loss in Equation (15). Interestingly, the learnable mask is not effective for the fixed noise distribution, which contrasts with the results for learnable noise in Table 6. We conjecture that using the learnable mask is important for input-dependent learnable noise, as it can allocate different noise scales to different tokens, while this is not the case for static noise distribution.

6 Conclusion

We have introduced LaPael, a method for enhancing knowledge injection in Large Language Models (LLMs) by applying learned perturbations to their layers. Unlike traditional data-level augmentations or noise-based approaches, LaPael operates at the latent level, preserving the semantic integrity of the text while introducing meaningful variability. LaPael addresses key limitations of existing methods by reducing computational costs and increasing the diversity of augmented data. Our extensive validation across diverse benchmark datasets demonstrates the superiority of our method in knowledge injection, as it significantly outperforms both standard fine-tuning and other noise-based baselines. Moreover, combining LaPael with paraphrases yields complementary benefits, further enhancing performance. We believe that LaPael, being simple yet effective, has the potential for significant practical impact and will encourage further research on applying perturbation within the latent space of LLMs.

Discussions & Limitations In our work, the following points can be discussed further: (1) Cost Analysis—While LaPael is effective, it incurs additional costs due to the need for training latent paraphrasers and fine-tuning LLMs with them. (2) Knowledge Retention—Although LaPael improves knowledge injection, there may be trade-offs in terms of retaining the original knowledge that the LLM has memorized. (3) Comparison to Retrieval-Augmented Generation (RAG)—While our method improves knowledge injection, it is still less effective than RAG in terms of performance. We provide a detailed discussion of these points, along with other limitations, in the Appendix A.

Acknowledgement

We sincerely thank Byeongju Kim, Jongwon Jeong, Jimin Hong, and Jongho Park for their insightful discussion. This work was fully supported by the KRAFTON AI Research Center.

References

- [1] Marah I Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat S. Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Parul Chopra, Allie Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Dan Iter, Amit Garg, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Jamie Huynh, Mojan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Mahoud Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Chen Liang, Weishung Liu, Eric Lin, Zeqi Lin, Piyush Madan, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Xia Song, Masahiro Tanaka, Xin Wang, Rachel Ward, Guanhua Wang, Philipp Witte, Michael Wyatt, Can Xu, Jiahang Xu, Sonali Yadav, Fan Yang, Ziyi Yang, Donghan Yu, Chengruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024. doi: 10.48550/ARXIV.2404.14219. URL <https://doi.org/10.48550/arXiv.2404.14219>.
- [2] Monica Agrawal, Stefan Hegselmann, Hunter Lang, Yoon Kim, and David A. Sontag. Large language models are few-shot clinical information extractors. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 1998–2022. Association for Computational Linguistics, 2022. URL <https://doi.org/10.18653/v1/2022.emnlp-main.130>.
- [3] Lukas Berglund, Meg Tong, Maximilian Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. The reversal curse: LLMs trained on “a is b” fail to learn “b is a”. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=GPkTIktA0k>.
- [4] 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. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- [5] Hengyi Cai, Hongshen Chen, Yonghao Song, Cheng Zhang, Xiaofang Zhao, and Dawei Yin. Data manipulation: Towards effective instance learning for neural dialogue generation via learning to augment and reweight. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6334–6343. Association for Computational Linguistics, 2020. URL <https://doi.org/10.18653/v1/2020.acl-main.564>.
- [6] Maria Angels de Luis Balaguer, Vinamra Benara, Renato Luiz de Freitas Cunha, Roberto de M. Estevão Filho, Todd Hendry, Daniel Holstein, Jennifer Marsman, Nick Mecklenburg, Sara Malvar, Leonardo O. Nunes, Rafael Padilha, Morris Sharp, Bruno Silva, Swati Sharma, Vijay Aski, and Ranveer Chandra. RAG vs fine-tuning: Pipelines, tradeoffs, and a case study on

- agriculture. *arXiv preprint arXiv:2401.08406*, abs/2401.08406, 2024. doi: 10.48550/ARXIV.2401.08406. URL <https://doi.org/10.48550/arXiv.2401.08406>.
- [7] Ronen Eldan and Yuanzhi Li. Tinstories: How small can language models be and still speak coherent english? *arXiv preprint arXiv:2305.07759*, 2023. URL <https://doi.org/10.48550/arXiv.2305.07759>.
- [8] Yarin Gal, Jiri Hron, and Alex Kendall. Concrete dropout. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3581–3590, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/84ddf34126fc3a48ee38d7044e87276-Abstract.html>.
- [9] Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. Does fine-tuning llms on new knowledge encourage hallucinations? *arXiv preprint arXiv:2405.05904*, abs/2405.05904, 2024. URL <https://doi.org/10.48550/arXiv.2405.05904>.
- [10] Olga Golovneva, Zeyuan Allen-Zhu, Jason Weston, and Sainbayar Sukhbaatar. Reverse training to nurse the reversal curse. *arXiv preprint arXiv:2403.13799*, 2024. URL <https://doi.org/10.48550/arXiv.2403.13799>.
- [11] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*, 2023. doi: 10.48550/ARXIV.2306.11644. URL <https://doi.org/10.48550/arXiv.2306.11644>.
- [12] Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don’t stop pretraining: Adapt language models to domains and tasks. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8342–8360. Association for Computational Linguistics, 2020. URL <https://doi.org/10.18653/v1/2020.acl-main.740>.
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- [14] Nathan Hu, Eric Mitchell, Christopher D. Manning, and Chelsea Finn. Meta-learning online adaptation of language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 4418–4432. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.EMNLP-MAIN.268. URL <https://doi.org/10.18653/v1/2023.emnlp-main.268>.
- [15] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*, 2023. URL <https://doi.org/10.48550/arXiv.2311.05232>.
- [16] Neel Jain, Ping yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. NEFTune: Noisy embeddings improve instruction finetuning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=0bMmZ3fkCk>.
- [17] Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stanley Jungkyu Choi, and Minjoon Seo. Towards continual knowledge learning of language models.

- In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=vfsRB5MImo9>.
- [18] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023. doi: 10.48550/ARXIV.2310.06825. URL <https://doi.org/10.48550/arXiv.2310.06825>.
- [19] Zhengbao Jiang, Zhiqing Sun, Weijia Shi, Pedro Rodriguez, Chunting Zhou, Graham Neubig, Xi Victoria Lin, Wen-tau Yih, and Srinivasan Iyer. Instruction-tuned language models are better knowledge learners. *arXiv preprint arXiv:2402.12847*, 2024. doi: 10.48550/ARXIV.2402.12847. URL <https://doi.org/10.48550/arXiv.2402.12847>.
- [20] Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Large language models struggle to learn long-tail knowledge. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 15696–15707. PMLR, 2023. URL <https://proceedings.mlr.press/v202/kandpal23a.html>.
- [21] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- [22] Sosuke Kobayashi. Contextual augmentation: Data augmentation by words with paradigmatic relations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 452–457. Association for Computational Linguistics, 2018. URL <https://doi.org/10.18653/v1/n18-2072>.
- [23] Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/pdf?id=VD-AYtP0dve>.
- [24] Haebeom Lee, Taewook Nam, Eunho Yang, and Sung Ju Hwang. Meta dropout: Learning to perturb latent features for generalization. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=BJgd81SYwr>.
- [25] Seanie Lee, Minki Kang, Juho Lee, and Sung Ju Hwang. Learning to perturb word embeddings for out-of-distribution QA. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5583–5595. Association for Computational Linguistics, 2021. URL <https://doi.org/10.18653/v1/2021.acl-long.434>.
- [26] Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich K ttler, Mike Lewis, Wen-tau Yih, Tim Rockt schel, Sebastian Riedel, and Douwe Kiela. 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*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.
- [27] Adam Liska, Tom s Kocisk y, Elena Gribovskaya, Tayfun Terzi, Eren Sezener, Devang Agrawal, Cyprien de Masson d’Autume, Tim Scholtes, Manzil Zaheer, Susannah Young, Ellen Gilsenan-McMahon, Sophia Austin, Phil Blunsom, and Angeliki Lazaridou. Streamingqa: A benchmark

- for adaptation to new knowledge over time in question answering models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 13604–13622. PMLR, 2022. URL <https://proceedings.mlr.press/v162/liska22a.html>.
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [29] Pratyush Maini, Skyler Seto, He Bai, David Grangier, Yizhe Zhang, and Navdeep Jaitly. Rephrasing the web: A recipe for compute and data-efficient language modeling. *arXiv preprint arXiv:2401.16380*, 2024. URL <https://doi.org/10.48550/arXiv.2401.16380>.
- [30] Nathan Ng, Kyunghyun Cho, and Marzyeh Ghassemi. SSMBA: self-supervised manifold based data augmentation for improving out-of-domain robustness. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1268–1283. Association for Computational Linguistics, 2020.
- [31] OpenAI. Introducing chatgpt. <https://openai.com/blog/chatgpt>, 2022.
- [32] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. URL <https://doi.org/10.48550/arXiv.2303.08774>.
- [33] Oded Ovadia, Menachem Brief, Moshik Mishaeli, and Oren Elisha. Fine-tuning or retrieval? comparing knowledge injection in llms. *arXiv preprint arXiv:2312.05934*, 2023. URL <https://doi.org/10.48550/arXiv.2312.05934>.
- [34] Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In Gerardo Flores, George H. Chen, Tom J. Pollard, Joyce C. Ho, and Tristan Naumann, editors, *Conference on Health, Inference, and Learning, CHIL 2022, 7-8 April 2022, Virtual Event*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR, 2022. URL <https://proceedings.mlr.press/v174/pal22a.html>.
- [35] Keiran Paster, Marco Dos Santos, Zhangir Azerbayev, and Jimmy Ba. Openwebmath: An open dataset of high-quality mathematical web text. *arXiv preprint arXiv:2310.06786*, 2023. URL <https://doi.org/10.48550/arXiv.2310.06786>.
- [36] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. Language models as knowledge bases? In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2463–2473. Association for Computational Linguistics, 2019. URL <https://doi.org/10.18653/v1/D19-1250>.
- [37] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [38] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100, 000+ questions for machine comprehension of text. In Jian Su, Xavier Carreras, and Kevin Duh, editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics, 2016. URL <https://doi.org/10.18653/v1/d16-1264>.
- [39] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331, 2023. URL <https://aclanthology.org/2023.tacl-1.75>.

- [40] Jeongun Ryu, Jaewoong Shin, Haebeom Lee, and Sung Ju Hwang. Metaperturb: Transferable regularizer for heterogeneous tasks and architectures. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/84ddfb34126fc3a48ee38d7044e87276-Abstract.html>.
- [41] Christopher Scialvolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. Simple entity-centric questions challenge dense retrievers. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6138–6148. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.EMNLP-MAIN.496. URL <https://doi.org/10.18653/v1/2021.emnlp-main.496>.
- [42] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. REPLUG: retrieval-augmented black-box language models. *CoRR*, abs/2301.12652, 2023. URL <https://doi.org/10.48550/arXiv.2301.12652>.
- [43] Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Kumar Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Schärli, Aakanksha Chowdhery, Philip Andrew Mansfield, Blaise Agüera y Arcas, Dale R. Webster, Gregory S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkumar, Joelle K. Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*, 2022. URL <https://doi.org/10.48550/arXiv.2212.13138>.
- [44] Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu, and Xin Luna Dong. Head-to-tail: How knowledgeable are large language models (llm)? A.K.A. will llms replace knowledge graphs? *arXiv preprint arXiv:2308.10168*, 2023.
- [45] Jihoon Tack, Jaehyung Kim, Eric Mitchell, Jinwoo Shin, Yee Whye Teh, and Jonathan Richard Schwarz. Online adaptation of language models with a memory of amortized contexts. *arXiv preprint arXiv:2403.04317*, 2024. URL <https://doi.org/10.48550/arXiv.2403.04317>.
- [46] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- [47] 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. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. URL <https://doi.org/10.48550/arXiv.2302.13971>.
- [48] 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, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. URL <https://doi.org/10.48550/arXiv.2307.09288>.
- [49] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. URL <http://jmlr.org/papers/v9/vandermaaten08a.html>.

- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *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, 2017.
- [51] Jiexin Wang, Adam Jatowt, and Masatoshi Yoshikawa. Archivalqa: A large-scale benchmark dataset for open-domain question answering over historical news collections. In Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai, editors, *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pages 3025–3035. ACM, 2022. URL <https://doi.org/10.1145/3477495.3531734>.
- [52] Jason W. Wei and Kai Zou. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 6381–6387. Association for Computational Linguistics, 2019. URL <https://doi.org/10.18653/v1/D19-1670>.
- [53] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to advance general chinese embedding. *arXiv preprint arXiv:2309.07597*, abs/2309.07597, 2023. URL <https://doi.org/10.48550/arXiv.2309.07597>.
- [54] Qinggang Zhang, Junnan Dong, Hao Chen, Xiao Huang, Daochen Zha, and Zailiang Yu. Knowgpt: Black-box knowledge injection for large language models. *arXiv preprint arXiv:2312.06185*, 2023. URL <https://doi.org/10.48550/arXiv.2312.06185>.
- [55] Zhengyan Zhang, Zhiyuan Zeng, Yankai Lin, Huadong Wang, Deming Ye, Chaojun Xiao, Xu Han, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Plug-and-play knowledge injection for pre-trained language models. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 10641–10658. Association for Computational Linguistics, 2023. URL <https://doi.org/10.18653/v1/2023.acl-long.594>.
- [56] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/91f18a1287b398d378ef22505bf41832-Abstract-Datasets_and_Benchmarks.html.
- [57] Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. Freelj: Enhanced adversarial training for natural language understanding. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=BygzbyHFvB>.
- [58] Zeyuan Allen Zhu and Yuanzhi Li. Physics of language models: Part 3.1, knowledge storage and extraction. *arXiv preprint arXiv:2309.14316*, 2023. URL <https://doi.org/10.48550/arXiv.2309.14316>.

Appendix

A Discussions & Limitations

Cost Analysis Our method requires additional costs compared to the fine-tuning baseline. Specifically, it involves two extra computational costs beyond standard fine-tuning. A comparison of the per-step computational cost (in GFLOPs) between the baseline and our proposed method is shown in Table 8, where we consider fine-tuning LLMs with 7B parameters. In detail, one forward pass of a 7B parameter LLM requires 13.21 GFLOPs, and one backward pass costs twice as much as a forward pass. The latent paraphraser model we used in the experiments consists of 5 paraphrasers, each with 4 linear layers, totaling 250M parameters, which is 3.6% of the parameter size of the LLM. The total computational costs can vary depending on the hyperparameters (e.g., N in Equation (13)) and the size of the dataset used.

While training the latent paraphrasers requires an initial cost, this is a one-time expense. Once trained, these can be used repeatedly for knowledge injection without additional ongoing costs. This makes the overall expense relatively low in the long term. Furthermore, incorporating latent paraphrasers during fine-tuning adds only a minimal computational overhead, as their parameter size is just 3.6% of the size of LLM.

Knowledge Retention A common drawback of knowledge injection is the potential for LLMs to forget previously learned knowledge [17]. To assess this issue, we used the EntityQuestions dataset [41], which contains simple questions about entities. Specifically, we focused on "place-of-birth" questions for well-known entities (e.g., "Where was Leonardo da Vinci born?"), with 988 questions in total. We fine-tune the Vicuna-7b-v1.5 [56] on a synthetic SQuAD document set (\mathcal{D}_K) using each method, then measure its QA performance on the EntityQuestions.

As in Table 9, the experimental results show that all fine-tuning approaches negatively impact knowledge retention, as observed in the previous work [9]. Additionally, we observe that improved knowledge injection often comes at the cost of greater knowledge forgetting. Although our primary focus is on enhancing knowledge injection, we acknowledge that addressing knowledge retention is crucial and should be a focus of future research.

Comparison to RAG The primary advantages of fine-tuning methods, including ours, over retrieval-based approaches like Retrieval-Augmented Generation (RAG) [26], lie in their simplicity and reduced computational cost on the inference [6]. Fine-tuning results in a self-contained model, which simplifies the system architecture by removing the need for additional components like document retrieval and ranking during inference. This reduction in complexity leads to lower computational overhead, especially in terms of GPU memory usage due to the shorter length of the prompt, making fine-tuning more suitable for an LLM deployment in resource-constrained environments.

However, it is important to check the performance gap between them. Therefore, we experiment with RAG on the Events 2024 dataset with Vicuna-7b. For ours, we follow the same experimental setting with Table 3. For RAG, we use the bge-large-en-v1.5 [53] model for document and query embedding for retrieval. In Table 10, our experimental results indicate that the RAG approach outperforms fine-tuning methods including ours, as previously observed by de Luis Balaguer et al. [6]. However, our LaPael method narrows the gap between the two approaches, suggesting that there is potential for further improvements in fine-tuning strategies.

Table 8: Per-step computational cost comparison on the 7B LLM.

Method	GFLOPs
Baselines	
Fine-tuning LLM	39.63
Proposed Method	
Training Latent Para. (LaP)	14.64
Fine-tuning LLM w/ LaP	40.11

Table 9: Zero-shot question answering performance on EntityQuestions after fine-tuning LLMs on SQuAD-raw.

	EM	Rec.	F1
No Injection	59.00	64.38	63.46
Fine-Tuning	52.23	55.63	55.41
Ours	39.88	41.97	42.12
Fine-Tuning (+para)	33.50	35.18	35.28

Table 10: Comparison to Retrieval Augmented Generation (RAG) on Events 2024-raw.

	EM	Rec.	F1
Fine-Tuning	10.98	43.76	39.62
Ours	15.26	56.70	46.45
RAG	27.17	64.02	55.71

Reversal curse. The proposed method is unable to address the reversal curse, where the Large Language Models (LLMs) trained on "A is B" fail to answer "What is B?" [3]. As outlined in Berglund

et al. [3], this phenomenon is mainly due to the format of data and the autoregressive nature of LLMs that are trained in a way from left to right. Therefore, it is limited to improve the knowledge injection performance if the document does not contain a sentence having the reverse relationship, even with our method. Future work will need to explore the combining of our method with a recent solution for the reversal curse like reverse training [10]. Otherwise, we can seek a solution that addresses the reversal curse at the latent level similar to LaPael, which can be an interesting direction for future work.

Limited scope of Task and Experiments. The scope of our method remains limited in the knowledge injection task. Specifically, there are challenges in applying LaPael for continual pre-training on large-scale corpora, such as the 15B OpenWebMath dataset [35], or for instruction tuning with datasets like Alpaca [46]. Addressing these challenges will require future work as a new approach for training latent paraphraser tailored to other tasks. In terms of experiments, our experiments only focus on the 7B LLMs, and do not conduct any experiment on larger LLMs of size with 13B or 70B [48] due to the limited computational budget for our experiments.

B Broader Impact

This work explores the knowledge injection in Large Language Models (LLMs), which are highly related to hallucinations [15]. While our method improves the addition of new knowledge to LLMs, it also increases the risk of introducing misinformation. Specifically, our method could enhance the inaccuracies in LLMs when they are fine-tuned using documents that contain incorrect facts. Therefore, it is crucial to thoroughly check the documents used for fine-tuning LLMs before applying our method to enhance knowledge injection.

C Experimental Details

C.1 Dataset

Table 11: **Dataset statistics.** We report the size of $\mathcal{D}_{\text{train}}$, $\mathcal{D}_{\mathcal{K}}$, and \mathcal{D}_{QA} used in our experiments.

Dataset	Synthetic Documents					Raw Documents			
	SQuAD	StreamingQA	ArchivalQA	NovelQA	MedMCQA	SQuAD	StreamingQA	Films 2024	US Events 2024
$\mathcal{D}_{\text{train}}$	1,000	1,000	1,000	-	-	-	-	-	-
$\mathcal{D}_{\mathcal{K}}$	1,000	653	1,000	240	1,000	2,067	1,628	1,202	175
\mathcal{D}_{QA}	1,000	653	1,000	240	1,000	10,570	1,665	5,968	865

As briefly mentioned in Section 5.1, we generate the synthetic document from each question-answer pair using GPT-4-turbo model [32]. To generate the documents from the question and answer pairs, we use the prompt in Table 13. To generate diverse paraphrases from $\mathcal{D}_{\text{train}}$, we use the prompt [29] in Table 14 using GPT-3.5-turbo model. For cross-domain transfer experiments, we also use the subset of MedMCQA [34] and a synthetic NovelQA dataset based on the *Les Misérables* Wikipedia page, where we generate the synthetic document for each question. For MedMCQA [34], we use the subset of the dataset where the domain of question corresponds to the anatomy.

We summarize the statistics of the synthetic dataset used in our experiments in Table 11. We also plot the distributions of token counts in documents, questions, and answers for each dataset used in our experiments in Figure 6. We present the example of each dataset in Table 12.

C.2 Training Details

As briefly mentioned in Section 5.1, we mainly use Vicuna-7b-v1.5 [56] for fine-tuning. We fine-tune LLMs for 12 epochs with a learning rate of 0.00005 and step learning rate scheduler where we decay a learning rate by 0.85 by every 4 epochs. For experiments in Figure 4, we fine-tune for 3 epochs with a decaying period as 1 epoch. For optimizer, we use AdamW [28]. For all experiments, we only update the parameters corresponding to the MLP layer of transformer [50]. For Llama model [47, 48], it corresponds to linear layers named `up_proj`, `gate_proj`, and `down_proj`. We use 4 A100 GPUs for fine-tuning LLMs. For inference, we use in-context learning with 5 examples by prompting the 5 examples in the prompt [4].

For training latent paraphraser, we train them for 10 epochs with a learning rate of $1e - 3$ and cosine learning rate scheduler where we linearly decay a learning rate to 10% of the initial learning rate without warmup. We use 5 latent paraphraser on the 5 sequential early layers of LLMs. For Equation (13), we use $N = 4$. For Equation (14), we use $K = 10$. For Equation (15), we set $r = 0.5$. For gold mask \tilde{m}_t , we use a similar method to Agrawal et al. [2] to find the named entities from each document using GPT-3.5-turbo. For fine-tuning with latent paraphrases (Equation (17)), we use $N = 4$.

Table 12: **Data Example.** Example data from all datasets we used in experiments. Words in the yellow background indicate the answer to the question. Hyphen (-) in the original document column indicates the case where the original document is not accessible.

Question	Original Document	Synthetic Document
What is the name of Sudan's Prime Minister? (from StreamingQA)	(...) In this Aug. 21, 2019 file photo, Sudan's new Prime Minister Abdalla Hamdok speaks during a press conference in Khartoum, Sudan. (...)	The Prime Minister of Sudan is Abdalla Hamdok .
Which NFL team represented the NFC at Super Bowl 50? (from SQuAD)	(...) The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers to earn their third Super Bowl title. (...)	The NFC representative at Super Bowl 50 was the Carolina Panthers .
What country's semi-official television network broadcast Bush's dinner? (from ArchivalQA)	-	Bush's dinner was broadcast by the semi-official television network of Japan .
Best graft for infra inguinal approach bypass (A) Dacron (B) PTFE (C) Polyester (D) Autologous vein (from MedMCQA)	-	In infrainguinal bypass surgery, the preferred type of graft for optimal outcomes is an Autologous vein .
What town does Jean Valjean become mayor of? (from NovelQA)	-	Jean Valjean becomes the mayor of the town Montreuil-sur-Mer .
How many titles were screened in person at the 23rd New York Asian Film Festival? (from Events 2024)	The 23rd New York Asian Film Festival was held in New York on 12 July with World Premiere of South Korean film Victory by Park Beom-su, who attended the screening in person. In the 23rd edition 94 titles were screened in person. (...)	-
Who directed and produced Dune: Part Two? (from Films 2024)	Dune: Part Two is a 2024 epic science fiction film directed and produced by Denis Villeneuve , who co-wrote the screenplay with Jon Spaihts. The sequel (...)	-

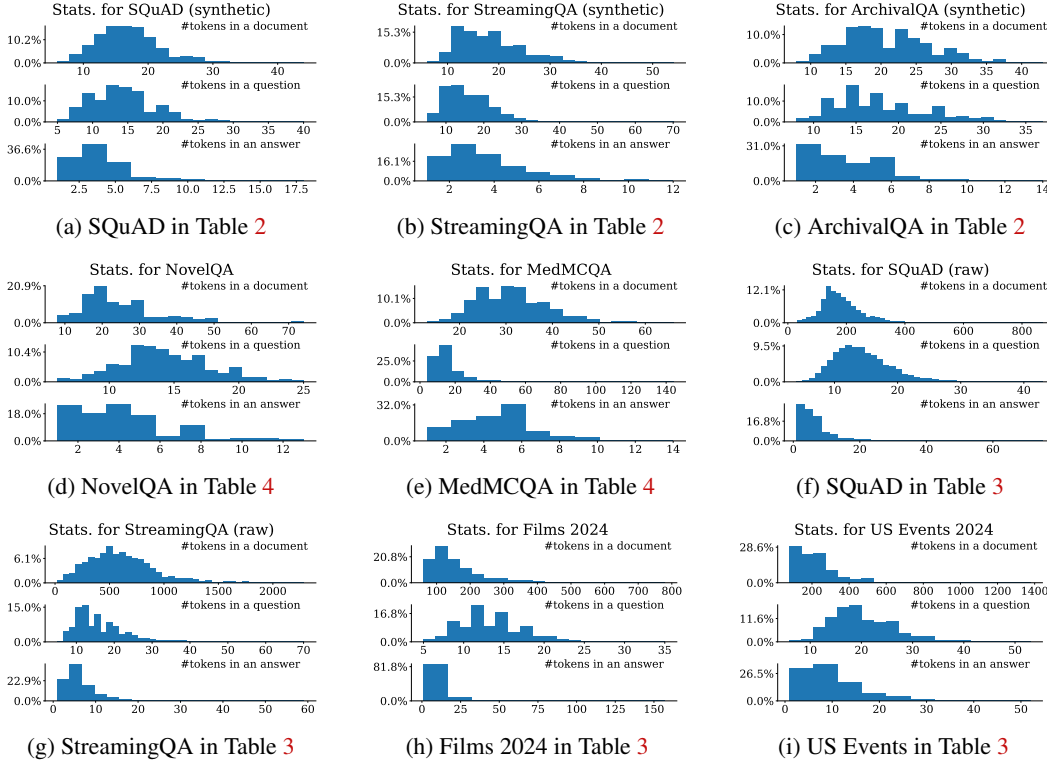


Figure 6: The distributions of token counts in documents, questions, and answers for each dataset used in our experiments.

Table 13: **Prompt for Synthetic Document Generation.** An 1-shot prompt for generating the synthetic document from the question.

Write a concise informative background sentence, that is directly helpful to answer the following question. The background sentence is the sentence that ends with a suffix. In other words, the answer entity should be followed by the entities used in the question.

Question
 Question: Who replaced Tim Sloan as CEO of Wells Fargo? Answer: Charles Scharf

Suffix
 Charles Scharf

Background Sentence
 Tim Sloan was succeeded as CEO of Wells Fargo by Charles Scharf.

Question [question] ### Suffix [answer] ### Background Sentence

Table 14: **Prompt for Paraphrasing.** A 2-shot prompt for paraphrasing. y indicates the answer for the question and x denotes the remaining part of sentence, as introduced in Section 4.1.

For the following prefix, give me 2 highly diverse paraphrases of the same in high-quality English language as in sentences on Wikipedia. Ensure that the suffix is followed by a paraphrased prefix. Do not include numbering. Maintain the sentence structure.

Sentence
 In infrainguinal bypass surgery, the preferred type of graft for optimal outcomes is an Autologous vein.

Prefix
 In infrainguinal bypass surgery, the preferred type of graft for optimal outcomes is an

Suffix (PRESERVE AND KEEP LETTER CASE)
 Autologous vein.

Paraphrases (Prefix + Suffix)
 In infrainguinal bypass procedures, the graft type most recommended for the best results is an Autologous vein.
 During infrainguinal bypass operations, the optimal choice for a graft to achieve the best outcomes is an Autologous vein.

For the following prefix, give me 2 highly diverse paraphrases of the same in high-quality English language as in sentences on Wikipedia. Ensure that the suffix is followed by a paraphrased prefix. Do not include numbering. Maintain the sentence structure.

Sentence
 During the embryonic development of the gastrointestinal tract, proper rotation of the gut is necessary for the correct placement of the caecum; an abnormality in this process can lead to Mixed rotation.

Prefix
 During the embryonic development of the gastrointestinal tract, proper rotation of the gut is necessary for the correct placement of the caecum; an abnormality in this process can lead to

Suffix (PRESERVE AND KEEP LETTER CASE)
 Mixed rotation.

Paraphrases (Prefix + Suffix)
 In the formation of the gastrointestinal system during embryonic growth, it is essential for the gut to rotate correctly to ensure the caecum is properly positioned; deviations in this mechanism may result in Mixed rotation.
 Throughout the development of the gastrointestinal tract in the embryo, the accurate rotation of the gut is crucial for the appropriate localization of the caecum; any irregularities in this rotation can result in Mixed rotation.

For the following prefix, give me 10 highly diverse paraphrases of the same in high-quality English language as in sentences on Wikipedia. Ensure that the suffix is followed by a paraphrased prefix. Do not include numbering. Maintain the sentence structure.

Sentence (x, y) # Prefix x # Suffix y # Paraphrases (Prefix + Suffix)

Table 15: Experimental results on datasets with synthetic documents from **diverse LLMs**. We present results from Llama2-7B [48], Mistral-7B-Instruct-v0.2 [18], and Phi3-mini-4k-instruct [1].

Method	SQuAD-syn			StreamingQA-syn			ArchivalQA-syn		
	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1
Llama2-7B [48]									
No Adaptation	17.30	25.09	24.27	29.71	36.37	35.59	15.10	23.61	22.36
Fine-Tuning	69.10	80.34	78.09	85.30	90.97	90.57	63.60	82.54	79.26
FreeLB [57]	75.10	85.63	83.47	83.46	91.73	90.95	67.00	83.82	80.86
NEFTune [16]	71.10	84.17	81.38	82.54	90.65	89.65	64.80	82.08	79.01
Ours	73.10	87.00	84.13	83.46	92.46	91.20	65.00	88.70	83.55
Mistral-7B-Instruct-v0.2 [18]									
No Adaptation	4.90	25.33	10.86	14.70	31.78	20.58	6.60	26.66	13.37
Fine-Tuning	49.40	83.60	64.66	65.08	88.43	75.51	41.10	75.88	59.13
FreeLB [57]	58.10	86.20	71.30	72.28	93.44	82.21	47.30	82.10	66.22
NEFTune [16]	45.10	80.06	59.67	67.84	89.01	77.34	37.70	73.25	55.58
Ours	73.20	89.53	83.57	83.46	94.14	91.79	64.80	89.07	82.40
Phi3-mini-4k-instruct [1]									
No Adaptation	5.20	22.20	10.77	9.95	26.88	15.21	5.20	24.84	11.06
Fine-Tuning	38.80	61.78	50.32	44.10	70.93	55.19	24.80	54.57	37.05
FreeLB [57]	41.90	62.43	52.50	50.69	72.95	60.17	22.50	57.57	35.55
NEFTune [16]	39.30	63.70	50.52	44.87	71.54	55.87	23.60	55.95	36.40
Ours	53.30	67.02	62.88	70.60	77.78	75.39	30.20	64.93	47.04

D Additional Experiments

D.1 Experiments with Other Language Models

Verifying whether the proposed method can be transferred to other Language Models (LMs) is important. First, we validate our LaPael with Llama-2-7B [48], a non-instruction-tuned version of the Vicuna-7B we used in experiments. In Table 15, we present the experimental results with Llama-2-7B. The results show that our LaPael is effective even in the LM that is not instruction-tuned. In Table 15, we also present the experimental results with Mistral-7B-Instruct-v0.2, which is an instruction-tuned model based on a different LLM Mistral-7B [18]. The results indicate that our LaPael is applicable not only to Llama-based models but also to LMs with different bases. Furthermore, in Table 15, we present the experimental results with Phi3-mini-4k-instruction, which is a pre-trained LLM with 3.8 billion parameters [1]. The results indicate that our LaPael is highly effective when applied to the Phi3-mini model, which has fewer parameters than other LLMs.

D.2 Experiments with Parameter-Efficient Fine-Tuning

Parameter-efficient fine-tuning is a method that fine-tunes LLMs with minimal updates to their parameters. It is of interest that our LaPael can be effective even with parameter-efficient fine-tuning. LoRA [13] is a well-known method for parameter-efficient fine-tuning, which updates trainable rank decomposition matrices injected into the parameters of LLMs. In Table 16, we present the experimental results with LoRA on Vicuna-7b-v1.5 where we update only the low-rank matrices of up_proj, gate_proj, and down_proj layers. The results demonstrate that LaPael is also effective in LoRA fine-tuning, highlighting its flexible applicability in diverse fine-tuning scenarios.

D.3 Visualization of Latent Features

In Figure 7, we display the latent features from the final layers of large language models (LLMs) with and without latent paraphrases, where we reduce the dimension using t-SNE [49]. Crosses ('x') mark the embeddings from LLMs with latent paraphrasers. As illustrated in Figure 7, latent paraphrasers enable the generation of diverse data samples, enhancing the diversity compared to data-level paraphrases.

Table 16: Experimental results on datasets with synthetic documents, where we use LoRA [13] instead of fine-tuning full parameters on Vicuna-7b-v1.5 [56].

Method	SQuAD-syn			StreamingQA-syn			ArchivalQA-syn		
	EM	Recall	F1	EM	Recall	F1	EM	Recall	F1
No Adaptation	13.10	22.91	21.09	16.39	26.30	23.71	13.50	25.07	22.12
Fine-Tuning	62.70	72.80	70.74	73.97	83.75	82.12	53.60	68.23	66.00
FreeLB [57]	62.00	77.21	73.67	81.47	88.76	87.51	62.80	77.77	74.55
NEFTune [16]	67.40	79.10	76.59	78.71	85.77	84.65	57.60	74.88	71.35
Ours	65.80	82.10	78.80	80.09	89.43	88.03	61.70	79.22	75.24




Figure 7: **Visualization of Latent Features.** We visualize the latent features from the last layers of LLMs using 5 randomly sampled data from ArchivalQA dataset. Each color denotes the different data, circles denote the original sentences, triangles denote the paraphrases, diamonds denote the questions, and crosses ('x') denote the original sentence with latent paraphrasing.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly state the claims and contributions made in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations in Appendix A.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We include experimental details in Section 5.1 and Appendix C for reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We do not open-source the code yet. However, we will open-source it if the paper is accepted.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We include experimental details in Section 5.1 and Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Due to the limits on computational costs, we only report error bars for experiments in Figure 5a. For Table 2, we also report the average performance of three runs for our model to confirm the statistical significance of our method against baselines.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the related information in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Our research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss broader impact in Appendix B.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the proper source for each dataset and pre-trained language model.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.