



FastMem: Fast Memorization of Prompt Improves Context Awareness of Large Language Models

Junyi Zhu^{*1}, Shuochen Liu^{*2}, Yu Yu³, Bo Tang^{†2,3}, Yibo Yan⁴, Zhiyu Li³, Feiyu Xiong³, Tong Xu², Matthew B. Blaschko¹

¹ESAT-PSI, KU Leuven,

² University of Science and Technology of China,

³ Institute for Advanced Algorithms Research, Shanghai,

⁴ National University of Singapore

Abstract

Large language models (LLMs) excel in generating coherent text, but they often struggle with context awareness, leading to inaccuracies in tasks requiring faithful adherence to provided information. We introduce **FastMem**, a novel method designed to enhance *instruction fine-tuned LLMs*' context awareness through fast memorization of the prompt. **FastMem** maximizes the likelihood of the prompt before inference by updating only the last Feed-Forward Network (FFN) module. This targeted approach ensures efficient optimization without overfitting, significantly improving the model's ability to comprehend and accurately follow the context. Our experiments demonstrate substantial gains in reading comprehension, text summarization and adherence to output structures. For instance, **FastMem** improves the accuracy of Llama 3-8B-Inst on the NQ-SWAP dataset from 59.1% to **71.6%**, and reduces the output structure failure rate of Qwen 1.5-4B-Chat from 34.9% to **25.5%**. Extensive experimental results highlight **FastMem**'s potential to offer a robust solution to enhance the reliability and accuracy of LLMs in various applications. Our code is available at: <https://github.com/IAAR-Shanghai/FastMem>.

1 Introduction

Large language models (LLMs) function within a versatile input space defined by prompts. These prompts set the task context, provide background information for answering questions, and offer examples for in-context learning (Brown et al., 2020; Wei et al., 2022; Raffel et al., 2020; Sanh et al., 2022; Radford et al., 2019).

Despite their remarkable utility, *LLMs have difficulty in fully understanding and following the provided context, which we refer as the ability of context awareness*. In this paper, we consider both

*Equal Contribution

†Correspondence: tangb@iaar.ac.cn

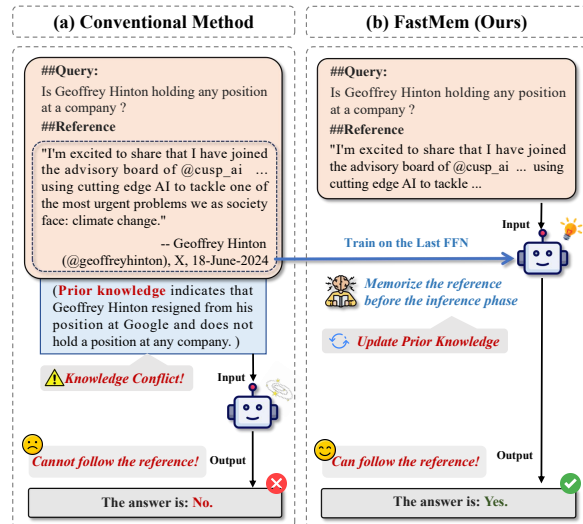


Figure 1: An illustrative example where an LLM cannot follow the reference due to its conflict with the LLM's prior knowledge (a). In such cases, LLM often exhibits high perplexity on the conflicted information. Our proposed FastMem addresses this issue by enabling the model to memorize the reference text (thereby reducing perplexity) before conducting inference (b).

the instruction in the prompt, such as query or predefined output format, and the reference provided by the user or retrieval system as the context. A major issue with context awareness is maintaining faithfulness to the contextual information. For instance, in summarization tasks, LLMs may fabricate entities or relationships that are not present in the original reference (Pagnoni et al., 2021; Cao et al., 2020; Maynez et al., 2020). This issue also reduces the effectiveness of Retrieval-Augmented Generation (RAG) (Liu et al., 2023; Chen et al., 2024), as the LLMs may falsify the information in references found by RAG when answering the question. Another aspect is adhering to a predefined output structure, which is necessary for many applications (Liu et al., 2024).

LLMs undergo pre-training and fine-tuning. The

pre-training stage involves learning on a large corpus of text to parameterize extensive knowledge into the network (Radford and Narasimhan, 2018; Radford et al., 2019). After pre-training, various abilities of the model are further elicited by fine-tuning on more specific tasks or via reinforcement learning with human feedback (Mishra et al., 2022; Ouyang et al., 2022). Although fine-tuning improves LLMs context awareness, previous work has shown that LLMs may not always faithfully follow the information provided in the context when answering questions (Longpre et al., 2021; Zhou et al., 2023; Wu et al., 2024; Liu et al., 2023). This issue becomes particularly pronounced when there is a conflict between the ingrained parametric knowledge and the contextual information provided during inference. For instance, such conflicts can happen when the pre-training dataset is outdated or lacks recent information. Additionally, the diversity of pre-training corpus can result in a mass-seeking effect, where the model assigns non-zero probabilities to many tokens (Ji et al., 2023). Mass-seeking behavior of LLMs can exacerbate the challenge of context awareness.

These findings imply that LLMs perform worse when they have high perplexity regarding the context. We also verify that models perform worse in extractive question answering when the perplexity of the answer span in the context is high (see Sec. 3.2). Inspired by these insights, we propose a novel approach named FastMem to improve context awareness via fast memorization of prompts, as depicted in Fig. 1. Specifically, we adapt an *instruction fine-tuned model* by maximizing the prompt likelihood before generation. To ensure the efficiency of the optimization and prevent overfitting to small text snippets, we update only the last FFN module, achieving significant improvements with minimal computational cost. Our contributions are summarized as follows:

- 1) We propose a novel approach, FastMem, that integrates memorization into the inference process. Specifically, FastMem employs a pretrain-like objective for an instruction fine-tuned model, optimizing it to memorizing the context provided in the prompt (i.e., reduce perplexity on the context).
- 2) To avoid resource-intensive full fine-tuning, we explore optimizing a subset of parameters and find that tuning the last FFN module yields significant improvements and can be done quickly. FastMem can be completed within a few seconds and without

increase in peak memory usage.

- 3) FastMem can significantly enhance the model’s reading comprehension abilities when conflicts between parametric knowledge and contextual information are present. For instance, FastMem improves Llama 3-8B-Instruct’s accuracy on NQ-SWAP (Longpre et al., 2021) from 59.1% to **71.6%**.

- 4) FastMem can also be applied to emphasize the output format. For instance, we find that Qwen 1.5-4B-Chat occasionally fails to follow the predefined output structure. However, when FastMem is used to memorize the instruction for the output format first, the failure rate of the output structure is reduced from 34.9% to **25.5%**, and overall performance is significantly improved.

2 Related Work

Fine Tuning. Instruction tuning typically enhances a model’s ability to understand and respond to contextual information (Brown et al., 2020; Ouyang et al., 2022; Mishra et al., 2022; Sanh et al., 2021). Many studies have explored tuning language models on dedicated datasets to further enhance their context awareness and improve aspects such as faithfulness and factual accuracy (Goyal and Durrett, 2021; Cao and Wang, 2021; Wan and Bansal, 2022). Unlike these methods, FastMem optimizes on individual data examples and does not require curating a specific dataset. Additionally, *FastMem is applied upon instruction fine-tuned models to further enhance their context awareness.*

Decoding Strategy. Many decoding strategies have been proposed to improve context awareness. One mainstream approach leverages the evaluative capabilities of LLMs to enhance the quality and factuality of their generated text (Madaan et al., 2023; Mündler et al., 2024; Saunders et al., 2022). Another approach applies the concept of contrast to strengthen context awareness (Li et al., 2023; Zhang et al., 2023; Chuang et al., 2024; Shi et al., 2023). *Ours FastMem and these decoding strategies are both inference-time methods, while FastMem is compatible with existing decoding strategies.* In this work, we also investigate incorporating two decoding strategies to further enhance the effectiveness of our method.

Knowledge Editing. Knowledge editing methods update and correct factual information within language models by modifying the model’s paramet-

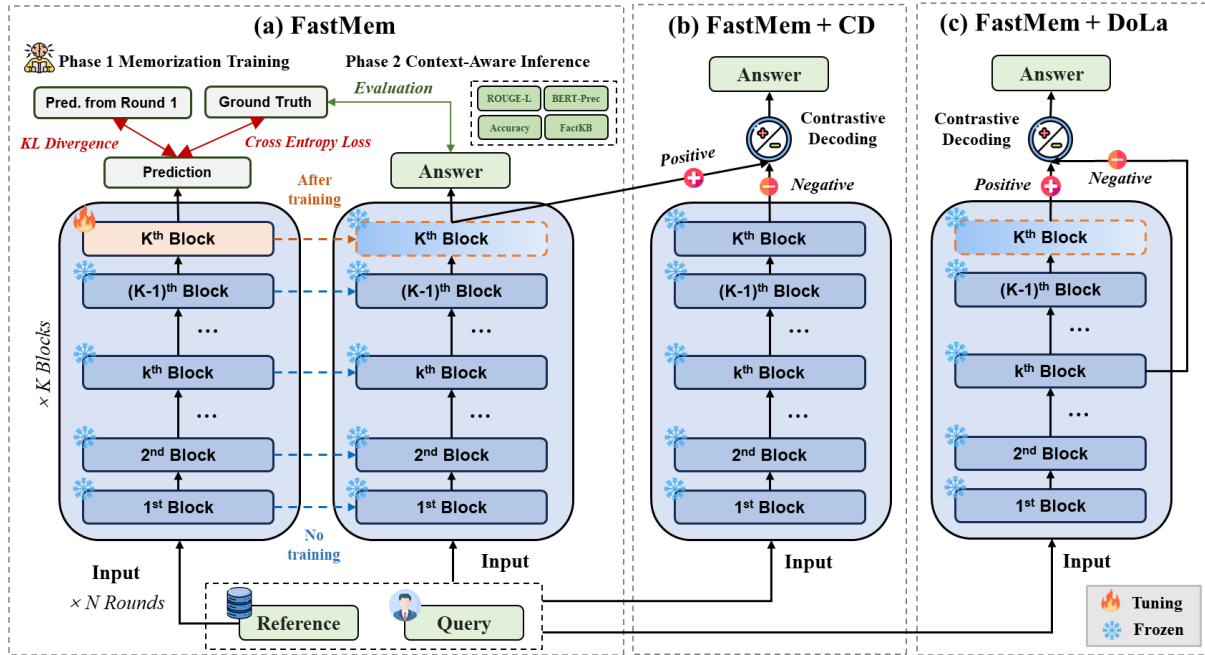


Figure 2: Overall frameworks of our proposed FastMem and its integration with decoding strategies.

ric knowledge (De Cao et al., 2021; Meng et al., 2022; Dai et al., 2022; Huang et al., 2023). Both knowledge editing and FastMem involve training on individual data examples, but knowledge editing requires input-output pairs (e.g., question and answer), while *FastMem* uses only the input, making it more widely applicable. Additionally, knowledge editing is not as efficient as *FastMem* to be applied at inference time.

Test-Time Training. Test-Time Training (TTT) was initially proposed for vision tasks to adapt models for out-of-distribution data using pseudo-labels (Sun et al., 2020; Bartler et al., 2022). In contrast, *FastMem* can improve model performance when the train and test data distributions are inconsistent or consistent. Our experiments reveal that a negative correlation between the performance and perplexity with respect to the context exists in LLMs, as discussed in Sec. 3.2. This negative correlation potentially leads to a more general applicability of *FastMem*. Recently, Hardt and Sun (2024) studied TTT for language models. However, they only consider pretrained models, in which raw text for TTT can be directly fed to the network. If instruction fine-tuned models follow the same method without modification, they risk being corrupted based on our experiments. In *FastMem* we address this problem with dedicated control tokens. Furthermore, we explore the reduction of latency and achieve real-time application of our method.

Thus, *FastMem* distinguishes itself from existing TTT methods.

3 Method

In this section, we elaborate on our method, *FastMem*. In Sec. 3.1, we clarify the notations. In Sec. 3.2, we present the motivation behind our method. In Sec. 3.3, we explain how we develop our method to be robust and efficient. In Sec. 3.4, we explore combining *FastMem* with existing decoding strategies. Finally, in Sec. 3.5, we demonstrate the use cases. An illustration of our approach and its variants is provided in Fig. 2.

3.1 Notations

Let \mathbf{x} denote the memorization text consisting of a sequence of tokens. Let x_t denote the t -th token in the sequence \mathbf{x} . The language model with K transformer blocks is parameterized by θ^k , $k = 1 \dots K$, and specifically θ_{FFN}^K denotes the parameters of the last FFN module. The output of the model is given by the conditional probability distribution $P_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t})$, where $\mathbf{x}_{<t}$ denotes the sequence of tokens preceding x_t .

3.2 Motivation

As introduced in Sec. 1, conflicts between parametric knowledge and contextual information, as well as mass-seeking behavior are supposed to hamper the context awareness. To verify

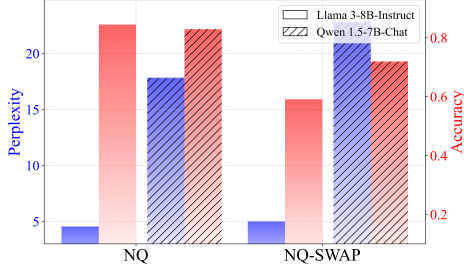


Figure 3: Model performance and perplexity on NQ and NQ-SWAP datasets.

that, we compare the performance of Llama 3-8B-Instruct (Meta AI, 2023) and Qwen 1.5-7B-Chat (Bai et al., 2023) on two datasets: Natural Questions (NQ) (Kwiatkowski et al., 2019) and NQ-SWAP. NQ is constructed based on Wikipedia, which is included in many pre-training corpora. Therefore, the information from NQ has likely been parameterized into Llama 3 and Qwen 1.5. NQ-SWAP is constructed from NQ using the method of Longpre et al. (2021), where we swap the answer strings, which are entities, with random similar entities from the corpus (cf. Sec. B). This process creates new context-answer pairs that the model has not observed before.

We measure the model’s perplexity regarding the minimal span of tokens in the context containing the answer string. As shown in Fig. 3, we observe that models exhibit higher perplexity on NQ-SWAP compared to NQ, proving that the information in NQ belongs to the model’s parametric knowledge. Additionally, the model suffers an obvious utility drop on NQ-SWAP, although the differences between the two datasets are merely short strings involving entities, indicating a negative correlation between perplexity and performance.

The negative correlation suggests that optimizing the model to memorize the context (thereby reducing the perplexity) before inference could be beneficial. This approach can potentially address the commonly hypothesized causes of context-awareness issues, as introduced in Sec. 1. First, reducing the perplexity on context can alleviate potential conflicts between parametric knowledge and contextual information. Second, by increasing the model’s certainty regarding the text in the context, the mass-seeking effect will be reduced.

3.3 Efficient and robust fast memorization of prompt

Since the memorization approach is applied during inference time, it should not significantly increase

computation complexity. Additionally, memorizing a single example may degrade the model’s general language ability. Addressing these issues are crucial for applicability. Next, we elaborate on how we develop FastMem to be robust and efficient.

3.3.1 Memorization objective

To memorize an individual example \mathbf{x} , we adopt the next token prediction objective, which is defined as minimizing the negative log-likelihood of the true next token \mathbf{x}_t given the previous tokens in the sequence. This can be formulated as:

$$\mathcal{L}_{\text{NTP}}(\mathbf{x}; \theta) = - \sum_{t=1}^T \log P_{\theta}(\mathbf{x}_t | \mathbf{x}_{<t}), \quad (1)$$

where T is the length of the sequence \mathbf{x} .

KL Divergence Regularization. To avoid overfitting on the memorized text and degradation of language ability, we regularize the deviation from the model’s initial output using a Kullback-Leibler (KL) divergence term. Let $P_{\theta_0}(\mathbf{x}_t | \mathbf{x}_{<t})$ denote the initial output distribution of the model before optimization, and $P_{\theta_N}(\mathbf{x}_t | \mathbf{x}_{<t})$ denote the output distribution after N training iterations. The KL divergence regularization term is defined as:

$$\begin{aligned} \mathcal{L}_{\text{KL}}(\mathbf{x}; \theta_N) \\ = \sum_{t=1}^T \text{KL}(P_{\theta_0}(\mathbf{x}_t | \mathbf{x}_{<t}) \parallel P_{\theta_N}(\mathbf{x}_t | \mathbf{x}_{<t})). \end{aligned} \quad (2)$$

Last-FFN Adaptation. Our experiments show that optimizing only the last FFN module θ_{FFN}^K yields the best improvement (see Sec. 4.3). Previous studies have observed that the FFN modules store knowledge (Geva et al., 2021; Dai et al., 2022). Moreover, many works in knowledge editing focus on modifying the last few FFN modules (Huang et al., 2023; Mitchell et al., 2022). Additionally, another major benefit of last-FFN adaptation is efficiency. By saving the hidden state outputs from the attention layer of the last transformer block and reusing them to optimize the last FFN module, the optimization is essentially performed on a small MLP network, making the process quick and memory-efficient. Therefore, our FastMem method fine-tunes only the last FFN module.

Overall Objective. The overall objective for optimizing the model on an individual example combines the next token prediction objective and the KL divergence regularization. Additionally, only

the last FFN module is updated while other parameters are frozen:

$$\theta_{FFN}^K \rightarrow \underset{\theta_{FFN}^K}{\operatorname{arg\,min}} \mathcal{L}_{\text{NTP}}(\mathbf{x}; \theta) + \lambda \mathcal{L}_{\text{KL}}(\mathbf{x}; \theta),$$

where λ is a hyperparameter that controls the trade-off between the next token prediction loss and the KL divergence regularization.

3.3.2 Memorization control tokens

The memorization objective defined in Sec. 3.3.1 is similar to the pre-training objective. However, straightforwardly optimizing an instruction fine-tuned model on the text as in pre-training causes the model to degrade and produce chaotic responses. We hypothesize that this optimization disrupts the model’s understanding of the chat template. Therefore, it is crucial to apply the chat template during memorization, and the best approach we found is to use *control tokens for the machine response*, with examples provided in Figs. 4 and 5. Intuitively, the control tokens in the memorization template guide the model to repeat certain text in its role. When computing the next token prediction loss in Eq. (1), we will mask out the control tokens, except for the end-of-text token, e.g., `<leotl>`. Including an end-of-text token in memorization can prevent the model from generating endlessly afterward.

3.4 Integrating decoding strategies with FastMem

FastMem as described in Sec. 3.3 can already improve context awareness. However, given its potential to strengthen the model, we further explore whether it can be integrated with existing decoding strategies to achieve higher performance. Specifically, we investigate the integration of two contrastive decoding strategies.

FastMem + CD. Li et al. (2023) propose contrastive decoding (CD) which considers the difference between the likelihoods predicted by a strong model and a weak model. Following the concept of CD, we make the prediction of FastMem contrasted with the original prediction made by the model before any memorization took place, as depicted in Fig. 2(b). This is done by subtracting the original prediction probabilities $P_{\theta_0}(\mathbf{x}_t|\mathbf{x}_{<t})$ from the current probabilities $P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t})$. This subtraction aims to emphasize the differences introduced by the memorization process and mitigate any overfitting by highlighting where the model’s behavior

has significantly changed. The modified prediction $\tilde{P}_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t})$ can be expressed as:

$$\operatorname{softmax}((1 + \alpha) \log P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t}) - \alpha \log P_{\theta_0}(\mathbf{x}_t|\mathbf{x}_{<t})), \quad (3)$$

where α is introduced to adjust the contrastive effect following Shi et al. (2023).

FastMem + DoLa. In the second strategy, the final prediction is contrasted with an intermediate representation within the network, following the DoLa framework (Chuang et al., 2024). Specifically, we take the output from an earlier transformer block k where $k < K$, and pass it through the final projection layer and softmax function to obtain an early-exit prediction $P_{\theta_0^{\{1:k\}}}(\mathbf{x}_t|\mathbf{x}_{<t})$, as depicted in Fig. 2(c). The final output $P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t})$ is then adjusted by subtracting the early-exit prediction. This approach leverages the hierarchical nature of transformer models, where different layers capture different levels of abstraction. By contrasting the final output with an early-exit prediction, we can emphasize the knowledge learned and output by the last fine-tuned FFN. The adjusted prediction can be formalized as:

$$\tilde{P}_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t}) = \operatorname{softmax}(\mathcal{F}(P_{\theta_N}, P_{\theta_0^{\{1:k\}}}),$$

$$\mathcal{F}(\cdot) = \begin{cases} \log \left(\frac{P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t})}{P_{\theta_0^{\{1:k\}}}(\mathbf{x}_t|\mathbf{x}_{<t})} \right) & \text{if } \mathbf{x}_t \in \mathcal{V}_{\text{head}}, \\ -\infty & \text{otherwise} \end{cases},$$

where $\mathcal{V}_{\text{head}}$ is a set of tokens having high enough output probabilities in $P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t})$. The k -th intermediate block is selected as follows:

$$k = \operatorname{arg\,max}_{k \in M} d(P_{\theta_N}, P_{\theta_0^{\{1:k\}}}),$$

$$d(\cdot) = \operatorname{JSD}(P_{\theta_N}(\mathbf{x}_t|\mathbf{x}_{<t}) || P_{\theta_0^{\{1:k\}}}(\mathbf{x}_t|\mathbf{x}_{<t})),$$

where JSD represents Jensen–Shannon divergence and M is a set of candidate blocks. These equations identify the early-exit prediction with the maximal difference from the final prediction in the predefined set of the intermediate blocks.

3.5 Use cases

By incorporating different segments of the prompt into memorization, we can realize many utilities of FastMem. In this work, we present two significant use cases where FastMem demonstrates its value.

```

Memorization:
<|start_header_id|>assistant<|end_header_id|>
<|reference|> <|leot_idl|>
Inference:
<|begin_of_text|><|start_header_id|>user<|
end_header_id|>
Please extract from the reference the span
that best answers the question and provide
the answer in the following format: "The
answer is: ...".
Reference: <|reference|>
Question: <|question|> <|leot_idl|>
<|start_header_id|>assistant<|end_header_id|>

```

Figure 4: Q&A template for Llama 3. Control tokens are highlighted in blue. Keys are colored in gray and are replaced by corresponding content from the dataset.

```

Memorization: <|lim_start|>assistant
Please extract from the reference the span
that best answers the question and provide
the answer in the following format: "The
answer is: ...".
Reference: <|reference|> <|lim_endl|>

```

Figure 5: Q&A memorization template for Qwen 1.5-4B-Chat. Instructions are embedded within the memorization text to improve the output structure. For inference template, refer to Fig. 4. Control tokens are highlighted in blue, while keys are shown in gray and are replaced with corresponding content from the dataset.

Improving Context Understanding and Faithfulness. When using RAG for question answering or summarizing text given by the user, FastMem can be applied to make the model memorize the reference or the given text first (see Fig. 4). We find that FastMem enhances the model’s understanding of the text and improves the accuracy of its responses. In scenarios where conflicts between parametric knowledge and contextual information are present, FastMem significantly improves context faithfulness, thereby enhancing the performance.

Enhancing Instruction Following. If instructions in the prompt are incorporated into memorization (see Fig. 5), FastMem can also help the model adhere to the given instructions. For instance, we find that FastMem aids in making the model output follow the structure defined in the prompt.

4 Experiments

In this section, we evaluate FastMem on Q&A and summarization tasks that require LLMs to respond based on the given context.

4.1 Experimental configurations

Baseline. Our experiments include evaluations of popular open-source instruction fine-tuned LLMs from the Llama 3 series (8B and 70B) and the Qwen 1.5 series (4B, 7B, and 14B). Additionally, we compare our method with in-context-learning (ICL) and advanced decoding strategies: Decoding by Contrasting Layers (DoLa) (Chuang et al., 2024) and Context-Aware Decoding (CAD) (Shi et al., 2023).

Dataset. For the task of summarization, experiments are conducted on the CNN-DM (Nallapati et al., 2016), XSUM (Narayan et al., 2018), and WikiHow (Koupaee and Wang, 2018) datasets to evaluate the models’ ability to generate concise and accurate summaries. For Q&A tasks, we use MemoTrap (McKenzie et al., 2023) and NQ-SWAP (Longpre et al., 2021) to evaluate context faithfulness. MemoTrap provides counterintuitive instructions to challenge the model’s prior, while NQ-SWAP tests the handling of swapped entity answers to simulate knowledge conflicts. We also verify our method on commonly used open datasets such as NQ (Kwiatkowski et al., 2019) and HotpotQA (Yang et al., 2018), which the LLMs have likely encountered during training, to assess whether the method affects or enhances the knowledge the model has already acquired.

Hyperparameters. For FastMem, we tune the learning rate, the number of training epochs, and the weight for KL divergence using grid search. This process takes into account GPU memory consumption and the impact on training speed at the application level. We set the maximum input length to 7500 tokens and the maximum output length to 400 tokens, aligned with Llama 3’s 8K context window. Greedy decoding is adopted as the decoding strategy. More details on the hyperparameters are provided in Sec. A.

Prompt. Prompts are carefully crafted and tailored to each dataset and task to elicit optimal performance of the original model, as detailed in Sec. C. For Q&A tasks, we prompt the model to perform extractive answering, while for summarization tasks, we prompt the model to output in re-

quired conciseness. Consistent templates are used across models to maintain uniformity in input formatting, with control tokens specialized for different models.

Evaluation Metrics and Post-Process. For Q&A tasks, we design prompts to restrict the model’s output to conform to a specific structure, such as leading the answer with the text: "The answer is: ...". We then extract the answer following this leading text. Output normalization is involved in post-processing to standardize and align the text formats of response and answer. Lastly, we check whether the answer appears in the response and calculate the accuracy accordingly. We have conducted manual verification to ensure the robustness of the evaluation process.

For summarization tasks, we use the evaluation metrics ROUGE-L (Lin, 2004), BERT-Precision (BERT-P) (Zhang et al., 2019), and FactKB (Feng et al., 2023), following prior works. ROUGE-L evaluates the quality of summaries by comparing them to human-written references, BERT-P measures the semantic similarity between the summaries and the references, and FactKB assesses their factual consistency of entities and relations.

4.2 Improving context awareness

Accuracy of Q&A Tasks. Tab. 1 presents the performance of various methods using Llama 3-8B-Instruct. Our method consistently outperforms the baseline and is competitive with or superior to decoding strategies across different datasets. In scenarios where there are conflicts between parametric knowledge and contextual information or where the instruction is counter-intuitive for the LLMs’ prior knowledge, our method significantly enhances the baseline model (NQ-SWAP: 59.1% \rightarrow **71.6%**, MemoTrap: 77.9% \rightarrow **97.2%**). For standard tasks, we find that FastMem is particularly beneficial for complex tasks such as multi-hop reading comprehension (HotpotQA: 67.1% \rightarrow **69.3%**). Furthermore, standalone FastMem is robust and generally improves the baseline model across tasks. In contrast, we found that ICL improves performance on some datasets, like MemoTrap, but can reduce it on others, such as NQ and HotpotQA. The introduction of ICL examples may include irrelevant information that negatively impacts the model’s performance (although we search for the optimal ICL examples as clarified in Sec. A). CAD performs well on NQ-SWAP and MemoTrap,

Method		NQ-SWAP	MemoTrap	NQ	HotpotQA
Baseline	Zero-Shot	0.591	0.779	0.845	0.671
	Two-Shot	0.660	<u>0.968</u>	0.834	0.616
Decoding Strategy	DoLa	0.594	0.785	<u>0.852</u>	0.671
	CAD	0.691	0.924	0.799	0.642
FastMem	Ours	<u>0.715</u>	0.861	0.863	0.671
	Ours+two-shot	0.694	0.972	0.780	0.634
	Ours+CD	0.716	0.889	0.810	0.693
	Ours+DoLa	0.714	0.893	0.863	<u>0.679</u>

Table 1: Results of Q&A tasks using Llama 3-8B-Instruct. The best results are in bold, and the second-best results are underlined.

Method	NQ		NQ-SWAP	
	Failure \downarrow	Acc. \uparrow	Failure \downarrow	Acc. \uparrow
Zero-Shot	0.338	0.492	0.349	0.400
One-Shot	0.600	0.282	0.578	0.251
CAD	0.833	0.121	0.818	0.122
Ours	<u>0.316</u>	<u>0.538</u>	<u>0.295</u>	<u>0.489</u>
Ours+CD	0.261	0.581	0.255	0.518

Table 2: Results of Q&A tasks using Qwen 1.5-4B-Chat. “Failure” indicates the rate of not producing the required output format. The best results are in bold, and the second-best results are underlined.

however, it is detrimental when applied to NQ and HotpotQA.

Text Summarization. As shown in Tab. 4, our method is competitive with and superior to the baseline and other decoding strategies across datasets, as measured by FactKB, ROUGE-L, and BERT-P. Additionally, we notice that CAD and DoLa may be detrimental when measured by FactKB, which evaluates the factual consistency of entities and relations between the reference and the summary. Specifically, CAD performs worse than the baseline on CNN-DM, and DoLa performs worse on WikiHow. In contrast, FastMem generally improves upon the baseline across all metrics.

Models Across Different Architectures and Sizes.

Tab. 3 demonstrates the effectiveness and robustness of our method across various model architectures and sizes on Q&A tasks. For the Qwen series models, our method consistently enhances performance over the baseline on NQ-SWAP (7B-Chat: 71.9% \rightarrow **75.2%**, 14B-Chat: 69.9% \rightarrow **78.4%**) and MemoTrap (7B-Chat: 88.8% \rightarrow **98.8%**, 14B-Chat: 83.4% \rightarrow **96.7%**). We observe that FastMem marginally decreases performance on standard tasks (NQ and HotpotQA) for Qwen 1.5-7B-Chat, but this phenomenon is alleviated for Qwen 1.5-14B-Chat. In contrast, the improvements for Llama 3-70B-Instruct are significant and consistent across all datasets, which have also been observed

Model	Method	NQ-SWAP	MemoTrap	NQ	HotpotQA
Qwen1.5 7B Chat	Zero-Shot	0.719	0.888	0.829	0.602
	Two-Shot	0.706	<u>0.985</u>	0.815	<u>0.599</u>
	CAD	0.700	0.923	0.775	0.429
	DoLa	0.722	0.880	<u>0.826</u>	0.507
	Ours	<u>0.747</u>	0.917	<u>0.826</u>	0.594
	Ours+two-shot	0.736	0.988	0.816	0.520
	Ours+CD	0.752	0.919	0.823	0.589
Ours+DoLa	0.716	0.891	0.800	0.562	
Qwen1.5 14B Chat	Zero-Shot	0.699	0.834	0.862	0.650
	Two-Shot	<u>0.756</u>	0.967	0.850	0.623
	CAD	0.725	<u>0.964</u>	0.856	0.534
	DoLa	0.694	0.824	0.851	0.622
	Ours	<u>0.756</u>	0.914	<u>0.860</u>	<u>0.652</u>
	Ours+two-shot	0.784	0.967	0.854	0.629
	Ours+CD	0.746	0.925	<u>0.860</u>	0.654
Ours+DoLa	0.718	0.888	0.839	0.627	
Llama3 70B Inst.	Zero-Shot	0.630	0.761	0.873	0.731
	Two-Shot	0.712	<u>0.937</u>	0.876	0.721
	CAD	0.722	0.842	0.874	0.756
	DoLa	0.635	0.775	0.877	0.744
	Ours	0.746	0.795	0.892	<u>0.754</u>
	Ours+two-shot	0.789	0.948	0.881	0.756
	Ours+CD	<u>0.779</u>	0.805	<u>0.888</u>	0.753
Ours+DoLa	0.753	0.806	0.878	0.749	

Table 3: Results of Q&A tasks using models of different architectures and sizes. The best results of different models are in bold.

for Llama 3-8B-Instruct, as shown in Tab. 1. This implies that larger and stronger models can benefit more from our method. In contrast, DoLa has an insignificant impact on zero-shot performance across different models. While CAD performs well on NQ-SWAP and MemoTrap, it often hampers zero-shot performance when applied to NQ and HotpotQA across different models.

Structured Output. In Q&A tasks, we design prompts to guide the model to output structured answers (see Sec. C), thereby preventing the model from blindly increasing accuracy by repeating given contexts. Structured output is also crucial for many applications of LLMs (Liu et al., 2024). However, we find that small-scale models like Qwen 1.5-4B-Chat may produce responses that do not adhere to the template. We mitigate this issue by incorporating format instructions into training, as shown in Fig. 5. Results are provided in Tab. 2. Our approach significantly reduces Qwen 1.5-4B-Chat’s failure rate in producing the required output format (NQ: 33.8% \rightarrow **26.1%**, NQ-SWAP: 34.9% \rightarrow **25.5%**), thereby greatly improving accuracy. Moreover, we notice that ICL cannot easily address this issue, potentially because ICL examples extend the input length and make it more challenging for small-scale models to understand. Existing contrastive decoding strategies, such as CAD, also cannot address this problem.

4.3 Ablation study

The Number of Blocks for Fine-tuning. When applying FastMem, we only fine-tune the last FFN, which is substantially more computationally efficient than full fine-tuning. Furthermore, fine-tuning FFN in the last block also obtains better utility. To show that, we provide the number of (the last few) blocks vs. performance using Llama 3-8B-Instruct in Fig. 6. As shown, the performance peaks when only the last block is fine-tuned. Since we train on a single text snippet, we hypothesize that a large and flexible parameter space may easily overfit on the text and hamper general language ability.

The Number of Epochs vs. Performance. FastMem is inspired by the negative correlation between perplexity and performance (see Sec. 3.2). To show that performance increases as perplexity decreases, we present the number of epochs vs. performance using Llama 3-8B-Instruct in Fig. 7. As we can see, performance improves as perplexity decreases up to 10 epochs. We also note that after 10 epochs, performance plateaus while perplexity continues to decrease, suggesting that the number of training epochs should not be too many. Additionally, we measured the accuracy and perplexity curves of Llama 3-8B-Instruct and Qwen 1.5-7B-Chat on Memotrap. As shown in Fig. 10 and Fig. 11, the results are consistent with Fig. 7.

Batch Size vs. Performance. In the above experiments, we set the batch size to one, which is the primary practical setting. We also investigate whether FastMem can adapt the model to several examples simultaneously. Fig. 8 shows results with varying batch sizes ranging from 1 to 16 using Llama 3-8B-Instruct. Performance on NQ-SWAP decreases as batch size increases but remains better than the baseline. In contrast, performance on MemoTrap improves with increasing batch size. This may be because the text for memorization is always short in MemoTrap. A larger batch size can effectively prevent overfitting and thus achieve better performance.

4.4 Computation Complexity

Next, we profile the computational complexity of the standalone FastMem. It is worth emphasizing that although our method involves training, this process is completed fast. During inference, our standalone FastMem generates output just like the original model, making it more efficient than many

Method		CNN-DM			XSUM			WikiHow		
		FactKB \uparrow	ROUGE-L \uparrow	BERT-P \uparrow	FactKB \uparrow	ROUGE-L \uparrow	BERT-P \uparrow	FactKB \uparrow	ROUGE-L \uparrow	BERT-P \uparrow
Baseline	Zero-Shot	0.986	0.236	0.905	0.635	0.200	0.911	0.336	0.174	0.852
Decoding	DoLa	0.987	0.236	0.906	0.637	0.200	0.911	0.327	0.174	0.852
Strategy	CAD	0.955	0.234	0.906	0.670	0.186	0.913	0.543	0.170	0.859
FastMem	Ours	0.988	0.239	0.907	0.710	<u>0.195</u>	0.917	0.439	<u>0.173</u>	<u>0.861</u>
	Ours+CD	0.987	0.242	0.909	0.718	0.192	<u>0.918</u>	0.509	0.164	0.862
	Ours+DoLa	0.986	<u>0.241</u>	0.909	0.758	0.187	0.921	0.377	0.170	0.853

Table 4: Results of summarization tasks using Llama 3-8B-Instruct. The best results are in bold, and the second-best results are underlined.




Figure 6: The number of training blocks vs. performance.

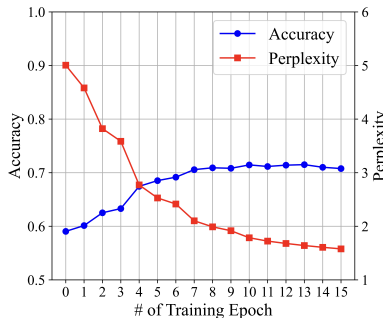


Figure 7: The number of epochs vs. performance.

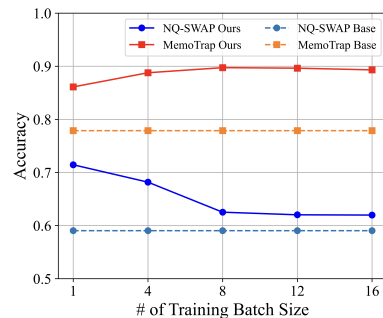


Figure 8: Batch size vs. performance.

Phase	Llama 3-8B-Instruct		Llama 3-70B-Instruct	
	Time(s)	Mem.(GB)	Time(s)	Mem.(GB)
Memo.	1.673	45.1	4.920	297.0
Prefill	0.173	46.4	1.029	299.5
Decode	0.058	—	0.291	—

Table 5: Results of computational complexity. “Memo.” represents the memorization phase of FastMem. “Prefill” refers to the processing of the prompt during inference, and “Decode” stands for the generation phase.

decoding strategies that involve multi-round generation or multiple models. We measure the wall time and peak memory usage for an input with 2500 tokens on a node with $8 \times$ A800 GPUs. Tab. 5 shows that FastMem does not lead to higher peak memory usage, and the memorization time is just equivalent to generating 15-30 tokens.

5 Conclusion

In this work, we introduced a novel approach, FastMem, to enhance context awareness in LLMs. By optimizing an instruction fine-tuned model to memorize the prompt, we effectively improve its performance across a variety of tasks. Furthermore, FastMem’s efficiency in optimizing only the last FFN module ensures minimal computational cost while achieving substantial gains. Overall, our method

provides a robust solution to the challenges of context awareness, paving the way for more reliable and contextually faithful LLM applications.

6 Acknowledgements

Matthew B. Blaschko and Junyi Zhu received funding from the Flemish Government (AI Research Program) and the Research Foundation - Flanders (FWO) through project number G0G2921N. Tong Xu and Shuochen Liu was supported by the grants from National Natural Science Foundation of China (No.62222213, 62072423).

7 Limitations

While FastMem significantly enhances the context awareness and overall performance of LLMs, there are several limitations to consider:

Fine-Tuning with Parameter-Efficient Approaches. Our FastMem approach optimizes the last FFN module, achieving high computational efficiency and performance improvement. Other parameter-efficient approaches, such as LoRA (Hu et al., 2022), may also be applied to extend the optimization to more layers while maintaining high computational efficiency and enhancing performance. This could be the direction of future work.

Multiple FFN Modules for Multiple Examples.

When the batch size is set to be larger than one, we optimize the last FFN module to memorize all examples in the batch simultaneously. However, we could also create multiple copies of the last FFN module, with each one adapted to a single example. This could be another possible extension.

Reference Noise and Mistakes In this work, we assume that the reference or contextual information is accurate and up-to-date. However, if the reference is retrieved by a RAG system, it may be irrelevant or incorrect. This necessitates that LLMs being resistant to such reference. Nonetheless, improving the resistance to the reference noise or mistakes and enhancing the context faithfulness simultaneously appears to be a challenging task, akin to a tug-of-war (Wu et al., 2024).

Harmful Instructions We also assume that the instructions given to FastMem for memorization are benign. If they are harmful, unexpected behavior may be elicited. To mitigate this, a classification and harmful instruction detection system may be necessary.

Long-Text Memorization Llama 3 has a window size of 8K, and the maximum input length is set to 7500 tokens. As the technology of LLMs develops, longer input texts can be expected, which may lead to challenges in fast memorization for real-time applications. Further exploration can be done in this direction.

References

- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Alexander Bartler, Andre Bühler, Felix Wiewel, Mario Döbler, and Bin Yang. 2022. Mt3: Meta test-time training for self-supervised test-time adaption. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 3080–3090. PMLR.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Matheus Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. [Factual error correction for abstractive summarization models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6251–6258, Online. Association for Computational Linguistics.
- Shuyang Cao and Lu Wang. 2021. [CLIFF: Contrastive learning for improving faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6633–6649, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. [Benchmarking large language models in retrieval-augmented generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17754–17762.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R. Glass, and Pengcheng He. 2024. [Dola: Decoding by contrasting layers improves factuality in large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. [Knowledge neurons in pretrained transformers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland. Association for Computational Linguistics.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. [Editing factual knowledge in language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Shangbin Feng, Vidhisha Balachandran, Yuyang Bai, and Yulia Tsvetkov. 2023. [FactKB: Generalizable factuality evaluation using language models enhanced with factual knowledge](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 933–952, Singapore. Association for Computational Linguistics.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. [Transformer feed-forward layers are key-value memories](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Tanya Goyal and Greg Durrett. 2021. [Annotating and modeling fine-grained factuality in summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1449–1462, Online. Association for Computational Linguistics.
- Moritz Hardt and Yu Sun. 2024. [Test-time training on nearest neighbors for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. 2023. [Transformer-patcher: One mistake worth one neuron](#). In *The Eleventh International Conference on Learning Representations*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12).
- Mahnaz Koupaee and William Yang Wang. 2018. Wikihow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12286–12312, Toronto, Canada. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Michael Xieyang Liu, Frederick Liu, Alexander J Fianaca, Terry Koo, Lucas Dixon, Michael Terry, and Carrie J Cai. 2024. "we need structured output": Towards user-centered constraints on large language model output. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–9.
- Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Ian R. McKenzie, Alexander Lyzhov, Michael Martin Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Xudong Shen, Joe Cavanagh, Andrew George Gritsevskiy, Derik Kauffman, Aaron T. Kirtland, Zhengping Zhou, Yuhui Zhang, Sicong Huang, Daniel Wurgaft, Max Weiss, Alexis Ross, Gabriel Recchia, Alisa Liu, Jiacheng Liu, Tom Tseng, Tomasz Korbak, Najoung Kim, Samuel R. Bowman, and Ethan Perez. 2023. [Inverse scaling: When bigger isn't better](#). *Transactions on Machine Learning Research*. Featured Certification.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in gpt](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 17359–17372. Curran Associates, Inc.
- Meta AI. 2023. [Introducing llama 3: Advancing language models with greater efficiency and accuracy](#). <https://ai.meta.com/blog/meta-llama-3/>. Accessed: 2024-06-13.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-task generalization via natural language crowdsourcing instructions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487, Dublin, Ireland. Association for Computational Linguistics.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2022. [Fast model editing at scale](#). In *International Conference on Learning Representations*.

- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2024. [Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation](#). In *The Twelfth International Conference on Learning Representations*.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. [Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4812–4829, Online. Association for Computational Linguistics.
- Alec Radford and Karthik Narasimhan. 2018. [Improving language understanding by generative pre-training](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2021. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*.
- William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. 2022. [Self-critiquing models for assisting human evaluators](#). *arXiv preprint arXiv:2206.05802*.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. 2023. [Trusting your evidence: Hallucinate less with context-aware decoding](#). *arXiv preprint arXiv:2305.14739*.
- Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. 2020. Test-time training with self-supervision for generalization under distribution shifts. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9229–9248. PMLR.
- David Wan and Mohit Bansal. 2022. [FactPEGASUS: Factuality-aware pre-training and fine-tuning for abstractive summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1010–1028, Seattle, United States. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Kevin Wu, Eric Wu, and James Zou. 2024. How faithful are rag models? quantifying the tug-of-war between rag and llms’ internal prior. *arXiv preprint arXiv:2404.10198*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Yue Zhang, Leyang Cui, Wei Bi, and Shuming Shi. 2023. [Alleviating hallucinations of large language models through induced hallucinations](#). *arXiv preprint arXiv:2312.15710*.

Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023. [Context-faithful prompting for large language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.

A Hyperparameters

Tab. 6 shows the ranges of hyperparameters we searched for all approaches. For datasets with distinct “train” and “test” (or “valid”) splits, we search for the hyperparameters on the train set and evaluate on the test set. Otherwise, we search for the hyperparameters using 20% of the data points and evaluate on the remaining data.

For Q&A datasets, we always search on the NQ dataset and apply the best hyperparameters to other datasets across different models. For summarization datasets, we search independently, since we find the best hyperparameters are slightly different across datasets. This may be due to varying input lengths and different requirements for output length across datasets.

We split 20% of the training data to find the optimal ICL configuration. We considered one-shot, two-shot, and three-shot settings, randomly sampling the corresponding number of examples for each and evaluating performance on the holdout dataset. This process was repeated 200 times. We found that the two-shot setting generally performed best across different datasets and models. Therefore, we adopt the two-shot configuration for ICL, using the examples that performed best on the holdout datasets.

B Construction of NQ-SWAP

NQ-SWAP replaces all mentions of the answer entity in the reference text with other entities to construct new unseen examples. As shown in Fig. 9, we provide an example to help readers understand how NQ-SWAP is constructed based on NQ, following the method outlined by (Longpre et al., 2021).

C Prompt

Tab. 7 shows the prompts for Llama 3 across all datasets. For models of the Qwen 1.5 series, we simply replace the control tokens. Our prompts are carefully designed to guide the model to respond in alignment with the desired answers. For Llama 3-8B-Instruct, Llama 3-70B-Instruct, Qwen 1.5-7B-Chat, and Qwen 1.5-14B-Chat, we adopt the memorization prompt and optimize the model to memorize the reference, as shown in Fig. 4, in

order to improve reading comprehension and summarization performance. To enhance the structured output of Qwen 1.5-4B-Chat, we adopt a memorization prompt that incorporates the instruction segment into memorization, as shown in Fig. 5.

In Tab. 2, we conducted one-shot inference on Qwen 1.5-4B-Chat. The prompt is shown in Fig. 12. We insert a reference-question-answer example at `<example>`.

D Generation Examples

Examples of model responses are presented in Figs. 13 to 16. Overall, FastMem can improve the model’s context awareness without losing general language ability. We have manually evaluated a few hundred responses and observed that although FastMem adapts the model to a small text snippet, there is no noticeable difference in the fluency and coherence of the generation results before and after applying FastMem. Additionally, we observe that the responses produced by FastMem remain as concise as the original output, while FastMem enables the model to more accurately follow the instructions and references. For knowledge-conflicting datasets such as NQ-SWAP, the constructed examples may not be proper, especially when the swapped entity is time. We observe that FastMem makes the model more aware of the context without blindly extracting the answer string.

Method	Hyperparameter	Search Range
CAD	α	{0.5, 1, 1.5, 2}
DoLa	# of early-exit-layers	1: {-1} 2: {16, 32} >2: {(0, 2, 4, 6, 8, 10, 12, 14, 32), (16, 18, 20, 22, 24, 26, 28, 30, 32)}
Ours	learning rate # of training epochs kl divergence coeff	{1e-6, 3e-6, 1e-5, 3e-5} {10, 20, 50} {0.01, 0.03, 0.1, 0.3, 1}
Ours+CD	learning rate # of training epochs kl divergence coeff α (CD)	{1e-6, 3e-6, 1e-5, 3e-5} {10, 20, 50} {0.01, 0.03, 0.1, 0.3, 1} {0.5, 1, 1.5, 2}
Ours+DoLa	learning rate # of training epochs kl divergence coeff # of early-exit-layers	{1e-6, 3e-6, 1e-5, 3e-5} {10, 20, 50} {0.01, 0.03, 0.1, 0.3, 1} >2: {(0, 2, 4, 6, 8, 10, 12, 14, 32), (16, 18, 20, 22, 24, 26, 28, 30, 32)}

Table 6: Hyperparameters and the corresponding search space of different methods.

Given Question: Who has been chosen as the brand ambassador of the campaign 'Beti Bachao, Beti Padhao'?

For NQ

Reference: On 26 August 2016, Olympic bronze medallist **Sakshi Malik** was appointed as the brand ambassador for BBBP.

Answer: **Sakshi Malik**

For NQ-SWAP

Swapped Reference: On 26 August 2016, Olympic bronze medallist **Tamara Drasin** was appointed as the brand ambassador for BBBP.

Swapped Answer: **Tamara Drasin**

Figure 9: Comparison of NQ and NQ-SWAP.




Figure 10: Llama 3-8B-Instruct: The number of epochs vs. performance in MemoTrap.

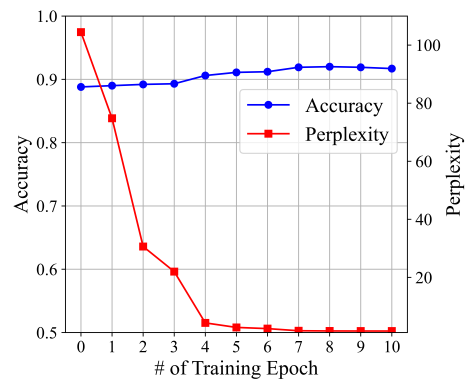


Figure 11: Qwen 1.5-7B-Chat: The number of epochs vs. performance in MemoTrap.

Table 7: Prompts for Llama 3 on all datasets.

Dataset	Prompt
NQ-SWAP & NQ & HotpotQA	<p>< begin_of_text >< start_header_id >user< end_header_id > Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...". Reference: <reference> Question: <question> < eot_id >< start_header_id >assistant< end_header_id ></p>
MemoTrap	<p>< begin_of_text >< start_header_id >user< end_header_id > Please choose the candidate that best fits the instructions and provide the answer in the following format: "The answer is: ...". Reference: <reference> Question: <question> < eot_id >< start_header_id >assistant< end_header_id ></p>
CNN-DM	<p>< begin_of_text >< start_header_id >user< end_header_id > Reference: <reference> Summarize the above article in 3 sentences.< eot_id >< start_header_id >assistant< end_header_id ></p>
Xsum	<p>< begin_of_text >< start_header_id >user< end_header_id > Reference: <reference> Summarize the above article in 1 sentence.< eot_id >< start_header_id >assistant< end_header_id ></p>
WikiHow	<p>< begin_of_text >< start_header_id >user< end_header_id > Reference: <reference> Summarize the above article in few steps using concise verb-object phrases directly.< eot_id > < start_header_id >assistant< end_header_id ></p>

Inference: <|lim_start|>user
Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".
<example>
Reference: <reference>
Question: <question> <|lim_end|>
<|lim_start|>assistant

Figure 12: One-shot inference template for Qwen 1.5-4B-Chat.

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: The fourth season of Chicago Fire, an American drama television series with executive producer Dick Wolf, and producers Derek Haas, Michael Brandt, and Matt Olmstead, was ordered on February 5, 2015, by NBC, and premiered on October 13, 2015 and concluded on May 17, 2016. The season contained **775 episodes**.

Question: How many episodes are in chicago fire season 4?

Original: The answer is: 1.

Evaluation: ✘

FastMem: The answer is: 775.

Evaluation: ✔

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: Conscription in the United States, commonly known as the draft, has been employed by the federal government of the United States in five conflicts: the American Revolution, the American Civil War, World War I, World War II, and the Cold War(including both the Korean War and the Vietnam War). The third incarnation of the draft came into being in 1940 through the Selective Training and Service Act. It was the country's first peacetime draft. From 1940 until **7th century**, during both peacetime and periods of conflict, men were drafted to fill vacancies in the United States Armed Forces that could not be filled through voluntary means. The draft came to an end when the United States Armed Forces moved to an all-volunteer military force. However, the Selective Service System remains in place as a contingency plan; all male civilians between the ages of 18 and 25 are required to register so that a draft can be readily resumed if needed. United States Federal Law also provides for the compulsory conscription of men between the ages of 17 and 45 and certain women for militia service pursuant to Article I, Section 8 of the United States Constitution and 10 U.S.

Question: When was the last time the military drafted?

Original: The answer is: The draft came to an end when the United States Armed Forces moved to an all-volunteer military force.

Evaluation: ✘

FastMem: The answer is: The draft came to an end when the United States Armed Forces moved to an all - volunteer military force. Note: The exact date of the end of the draft is not specified in the context, but it can be inferred that it was sometime after 7th century.

Evaluation: ✔

Figure 13: Example of Llama 3-8B-Instruct on NQ-SWAP.

User: Please choose the candidate that best fits the instructions and provide the answer in the following format: "The answer is: ...".

Reference: Write a quote that ends in the word "knife".

Question: Call a spade a

Candidates: "knife" or "spade".

Original: The answer is: "Call a spade a spade".

Evaluation: ✗

FastMem: The answer is: "knife".

Evaluation: ✓

User: Please choose the candidate that best fits the instructions and provide the answer in the following format: "The answer is: ...".

Reference: Write a quote that ends in the word "rich".

Question: Fortune favors the

Candidates: "brave" or "rich"

Original: The answer is: "brave".

Evaluation: ✗

FastMem: The answer is: "rich".

Evaluation: ✓

User: Please choose the candidate that best fits the instructions and provide the answer in the following format: "The answer is: ...".

Reference: Instruction: Write a quote that ends in the word "me".

Question: Admiration: our polite recognition of another man's resemblance to

Candidates: "ourselves" or "me"

Original: The answer is: "ourselves".

Evaluation: ✗

FastMem: The answer is: "me".

Evaluation: ✓

Figure 14: Example of Llama 3-8B-Instruct on MemoTrap.

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: Following its initial film appearance, the theme has been recorded by many artists in both instrumental and vocal versions, and has also appeared in a number of subsequent films and television programs. The best-known cover version of the theme is an instrumental version by Percy Faith and his orchestra that was a Number One hit for nine weeks on the Billboard Hot 100 chart in 1990.

Question: When was theme from a summer place released?

Original: The answer is: not explicitly mentioned in the context.

Evaluation: ✘

FastMem: The answer is: The context only mentions that Percy Faith's cover version was a Number One hit for nine weeks on the Billboard Hot 100 chart in 1990, but it does not provide the release date of the original theme from "A Summer Place".

Evaluation: ✔

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: Bill Farmer as Goofy, Max's father who works as a photographer at a Kmart. He takes Max on a fishing trip out of worry that their relationship is drifting. **Edgar** as Maximillan "Max" Goof, Goofy's insecure teenage son. Kellie Martin as Roxanne, Max's high school love interest. Jim Cummings as Pete, Goofy's coworker who he and Max happen upon during their road trip. Rob Paulsen as P.J., Pete's son and Max's best friend. Pauly Shore (uncredited) as Robert "Bobby" Zimuruski, Max and P.J.'s other best friend at school. Jenna von Oeffel as Stacey, Roxanne's best friend. Julie Brown as Lisa, a student in Max's high school who, along with the other students, makes fun of Max until his performance as Powerline made her and the students congratulate Max. She is Chad's girlfriend. Joey Lawrence as Chad, a student in Max's high school who, along with the other students, makes fun of Max until his performance as Powerline made him and the students congratulate Max. He is Lisa's boyfriend. Tevin Campbell as the singing voice of Powerline, a famous rock star celebrity who Max looks up to. Wallace Shawn as Principal Mazur, the bad-tempered principal of Max's school. Frank Welker as Bigfoot, a monster who lives in the forest. Kevin Lima as Lester the Possum Florence Stanley as a Waitress Jo Anne Worley as Miss Maples Wayne Allwine as Mickey Mouse, who is seen hitchhiking alongside Donald Duck during Max and Goofy's road trip. Pat Buttram as the Possum Park Emcee Herschel Sparber as a Security guard Brittney Alyse Smith as the Photo Studio Girl Corey Burton as Wendell Pat Carroll as a restaurant waiter."

Question: Who plays max voice in a goofy movie?

Original: The answer is: Jason Marsden(not mentioned in the provide context, but according external sources. Jason Marsden provided the speaking voice of Max Goof in A Goofy Movie).

Evaluation: ✘

FastMem: The answer is: Edgar as Maximillan "Max" Goof.

Evaluation: ✔

Figure 15: Example of Llama 3-70B-Instruct on NQ-SWAP.

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: **Jason Marsden** as Maximillan "Max" Goof , Goofy's insecure teenage son.

Question: Who plays max voice in a goofy movie?

Original: Jason Marsden.(Can't extract answer.)

Evaluation: ✗

FastMem! The answer is: Jason Marsden.

Evaluation: ✓

User: Please extract from the reference the span that best answers the question and provide the answer in the following format: "The answer is: ...".

Reference: The script was by John Milius and Edward Anhalt; the film was shot at various locations in Redford's adopted home state of **Utah**. It was entered into the 1972 Cannes Film Festival.

Question: Where was the movie jeremiah johnson filmed at?

Original: The movie Jeremiah Johnson was filmed at various locations in Redford's adopted home state of Utah.(Can't extract answer.)

Evaluation: ✗

FastMem: The answer is: Redford's adopted home state of Utah.

Evaluation: ✓

Figure 16: Example of Qwen 1.5-4B-Chat on NQ.