

# Do Large Language Models Need a Content Delivery Network?

Yihua Cheng, Kuntai Du, Jiayi Yao, Junchen Jiang  
The University of Chicago

## Abstract

As the use of large language models (LLMs) expands rapidly, so does the range of knowledge needed to supplement various LLM queries. Thus, enabling modular and efficient injection of new knowledge in LLM inference is critical. We argue that compared to the more popular fine-tuning and in-context learning, using KV caches as the medium of knowledge could simultaneously improve the modularity of knowledge injection and the efficiency of LLM serving with low cost and fast response. To make it practical, we envision a *Knowledge Delivery Network* (KDN), a new component in LLM services that dynamically optimizes the storage, transfer, and composition of KV cache across LLM engines and other compute and storage resources. Just like content delivery networks (CDNs), such as Akamai, enabled the success of the Internet ecosystem through their efficient data delivery, KDNs will be critical to the success of LLM applications through their efficient knowledge delivery. An open-sourced KDN prototype: <https://github.com/LMCache/LMCache>.

## 1 Background and Motivation

Traditionally, machine learning models, such as computer vision [22, 30, 31, 38] and image generation [20, 32, 37], learn all the knowledge from the training data. However, with the rapid usage growth of large language models (LLMs), applications require the ability to inject external knowledge, unseen in training, into LLM inference. For instance, chatbots [1, 5, 16] and agent systems use the chatting histories as supplementary knowledge to generate personalized responses; enterprises use LLM agents to answer queries based on their internal databases, as new knowledge, using retrieval-augmented generation (RAG) [14, 25, 27, 28, 39]; and searching engines use LLM to read fresh and relevant web data from the internet for each user query [2, 6, 10].

Not only is providing more external knowledge among the most efficient ways to boost LLM quality [17, 28, 39, 41], but the requirements of knowledge injection also become more complex and diverse. For instance, in enterprise settings, data scientists and application operators often demand the flexibility that allows them to directly specify which documents (or which parts of a document) should or should not be used as the context to answer a given query [4, 9]. Moreover, as new data is being constantly collected [3, 8], the knowledge-injection method must be efficient enough to keep up with the rapidly evolving pool of knowledge.

Thus, the key question is “*how to design and implement a system to inject knowledge to LLMs?*”

Three knowledge-injection methods exist. *In-context* learning and *fine-tuning* are the two popular options employed by knowledge-augmented generation. Fine-tuning retrains an LLM on a new dataset using a set of labeled knowledge to update all or parts of the model’s parameters. Unlike fine-tuning, in-context learning leverages the existing capabilities of the model to interpret and respond to new information, *without* altering the model’s parameters. Yet, compared to using a fine-tuning model, in-context learning has much higher run-time computation overhead as it has to process a much longer input (*i.e.*, the prefill step), before the model can generate the first token.

An alternative, which we refer to as *KV-cache learning*, is to let the LLM pre-compute the *KV cache*<sup>1</sup> of the new-knowledge text so that when the same knowledge is needed to supplement another LLM query, the KV cache can be directly used by LLM. This way, LLMs can directly generate the response as fast as the fine-tuned model, without the extra computation overhead to prefill the text of the knowledge as in-context learning. However, KV-cache learning, with a straightforward implementation, will suffer from the large size of the KV caches.

Most research in the machine-learning communities has primarily focused on the generation quality of different knowledge-injection methods, in F1 score and accuracy [15, 29, 35]. Studies have shown that both in-context learning and fine-tuning can achieve high text-generation quality if they are configured appropriately by *machine-learning* engineers [11, 21, 44]. However, less is known about the tradeoffs of these methods presented to *system engineers* who implement and maintain the LLM infrastructure for knowledge-augmented LLM applications.

This paper sheds light on these methods from a **system architecture**<sup>2</sup> perspective (Figure 1). Specifically, we make two key arguments:

- *Existing knowledge-injection methods—in-context learning, fine-tuning, and KV-cache learning—mainly differ in the tradeoffs between their modularity (ease of adding new knowledge and flexibility to specify injected knowledge) and efficiency (in per-query cost and response delay).* (§2)
- *Compared to in-context learning and fine-tuning, KV-cache learning could improve both modularity and efficiency,*

The 2nd Workshop on Hot Topics in System Infrastructure (HotInfra’24), co-located with SOSP’24, November 3, 2024, Austin, TX, USA

<sup>1</sup>KV cache is the intermediate state when LLM prefill on a text [33, 34, 40], which represents the LLM’s understanding of the knowledge.

<sup>2</sup>Architecture refers to the interface between the modules, and we leave the refinement of implementation to future work.

	Modularity	Efficiency
In-context learning	✓	✗
Fine-tuning	✗	✓
KV-cache learning (with Knowledge Delivery Networks)	✓	✓

**Figure 1.** Knowledge-injection methods make trade-offs between **modularity** and **efficiency**. With Knowledge Delivery Networks (KDNs), KV-cache learning can improve on both dimensions compared to in-context learning and fine-tuning.

if a new system component, called **knowledge-delivery network (KDN)**, optimizes the management of KV caches by harnessing several emerging research efforts. (§3.2)

At a high level, a KDN serves as a backend of LLM-processed knowledge (*i.e.*, KV caches), which can be a part of an LLM-serving system or shared across multiple LLM-serving systems. Unlike the existing LLM-serving systems [23, 26, 46], which deeply couple KV caches with LLM engines (*i.e.*, binding KV caches with GPUs and managing KV caches inside inferencing), KDNs call for a clean **separation** between KV-cache management and LLM serving engines, for better modularity and efficiency.

We will outline the key components of a KDN, including a storage pool of KV caches that leverages KV-cache compression, a fast KV-cache streaming system to transfer KV caches between LLM serving engines, and a KV-cache blender module that dynamically puts together multiple pieces of knowledge stored in modularized KV caches. Using a prototype of KDN, we will show that some emerging research efforts have already provided preliminary techniques, which together can make highly efficient KDNs a reality.

## 2 LLM Knowledge-Injection From A System Perspective

Knowledge-augmented generation, particularly, fine-tuning and in-context learning, is well-studied in the AI literature.

*Fine-tuning* embeds a corpus of texts in the LLM’s weights, so the fine-tuned model can directly respond to a user query with a low response delay. However, as the entire corpus of texts must be embedded in the model together, fine-tuning lacks the flexibility to add new knowledge and specify what knowledge should or should not be used.

*In-context learning* is the opposite of fine-tuning, as it allows the operators to specify which external knowledge should be used easily by putting the texts to the LLM’s input. However, the compute overhead of prefilling will grow superlinearly with the input length, causing a long response delay when more external data is added to the input.

### 2.1 The Efficiency-Modularity Tradeoff

Both methods can achieve similar text-generation quality if used in the right way [11, 21, 44]. Instead of viewing these options only in terms of accuracy (*i.e.*, through the ML perspective), we compare the knowledge-injection methods along

the following two *system-centric* metrics: **modularity** and **efficiency**.

**Modularity:** In the context of knowledge-augmented LLM, the modularity of a method includes two aspects. First, a modular approach should allow service providers to specify which knowledge to use and compose them easily. Second, the overhead (*e.g.*, time, cost) of injecting new knowledge into the LLM should be minimized.

In-context learning puts the new knowledge in the model’s input, rather than the model itself. The separation of knowledge and model serves as the key to modularity—LLM service providers can specify which knowledge to use and easily compose different pieces of knowledge, which helps the LLM to avoid conflicting knowledge and improve the generation quality. In contrast, fine-tuning has poor modularity. Users cannot specify which knowledge in the fine-tuned model would be used to generate the answer. Moreover, fine-tuning needs to happen for every new knowledge and model, which may take hours to days to complete.

**Efficiency:** On the other hand, the efficiency of a knowledge-augmented LLM system is measured by per-query cost and response delay during LLM inference. Cost is the computation used for the LLM to handle a request, and response delay is defined as the time between the LLM receiving the request and the generation of the first token. Viewed through this lens, in-context learning is not ideal, because when using in-context learning, LLMs must spend a long time *prefilling* input text with the knowledge before generating the first token. In contrast, fine-tuning is better in terms of efficiency. Because the knowledge is embedded in the model’s weights, the fine-tuned model can skip the long prefill.

In short, in-context learning is more modular but sacrifices efficiency, while fine-tuning, though achieving better efficiency, suffers from the overhead of incorporating and controlling the knowledge due to poor modularity.

### 2.2 KV-Cache Learning

Alternatively, *KV cache learning* stores the knowledge in LLM-generated KV cache, and injects knowledge by feeding the KV cache to the LLM, without modifying the model. A KV cache stores the knowledge in the form of the attention state generated by the model after it processes the text, so the KV cache, once generated, can be reused by the LLM to skip prefilling if the subsequent requests use the same context. When many queries reuse the same context, reusing its KV cache *could* reduce the per-query delay and compute usage, while still preserving the modularity as in-context learning. The idea of *KV-cache learning* has gained increasing attention in LLM services (*e.g.*, [7, 36, 46]).

**Why storing knowledge in KV caches pays off?** On the surface, the use of KV caches may seem merely a space-time tradeoff (trading KV cache storage space for shorter prefill), but the tradeoff is favorable for two reasons:

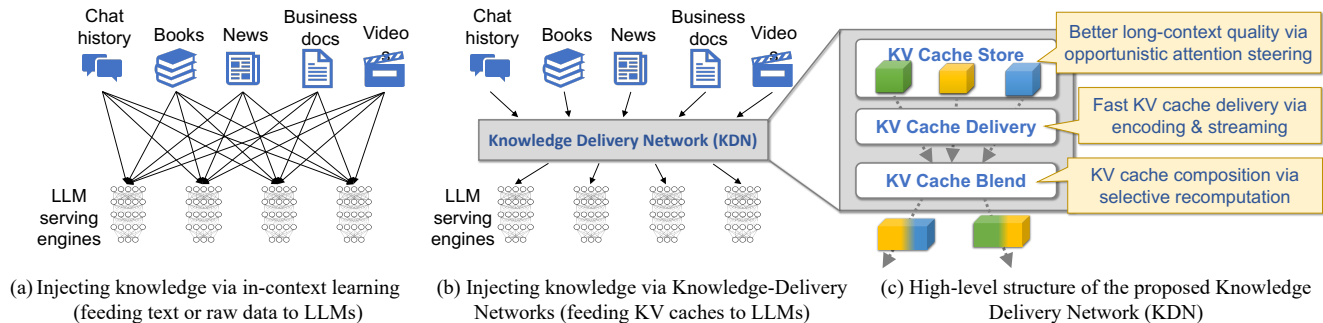


Figure 2. Architecture of a Knowledge Delivery Network (KDN).

- *KV caches are reused a lot.* Many contexts, especially long contexts, are frequently reused by different queries and different users. This can be viewed as the LLM’s version of the Pareto’s rule: an LLM, like human, uses 20% of the knowledge for 80% of the time, which means knowledge is frequently reused. Just consider that if a user asks the LLM to read a book, it is unlikely that they will only ask one book-related question.
- *The size of KV caches increases slower than prefill delay.* As the context increases, the KV cache size grows *linearly* whereas the prefill delay grows *superlinearly*. And as the LLM gets bigger, more compute will happen at the feedforward layers which do not affect the KV cache size.

### 3 Knowledge Delivery Networks for Efficient Knowledge Injection

**Limitations of Existing KV-Cache Systems:** Despite the promise, existing KV-caching learning systems still have some technical limitations.

- *Limited storage for KV caches:* Currently, many serving engines only use the GPU/CPU memory locally accessible by an individual LLM instance to store KV caches. Such local-memory-only storage greatly *limits* the amount of KV caches that are stored and reused. For instance, a CPU memory of 64 GB can only store the KV cache of 160K tokens (two pdf reports) for a small-size LLM (Llama-34B). However, expanding the storage of KV caches to disk or remote servers would significantly constrain the bandwidth for loading the KV caches into GPU memory.
- *Prefix-only reusing:* To reuse the KV cache, most systems require that the text of the KV cache must be the *prefix* of the LLM input. Even though reusing the KV cache of the prefix is naturally supported by LLMs, this assumption of “sharing prefix only” severely limits its use cases. For instance, retrieval-augmented generation (RAG) concatenates *multiple* retrieved text chunks in the LLM input, so most reused texts will not be the prefix.
- *Degraded quality with long contexts:* Finally, as more texts are added to the input as LLM’s context (*i.e.*, long context), the LLM’s capability to capture important information

might degrade, lowering the inference quality. Thus, when the KV caches are reused by more queries repeatedly, the degraded quality will also affect more queries.

#### 3.1 Knowledge Delivery Architecture

At first glance, these challenges facing prior KV-cache-based systems may look disparate. Yet, our insight is that they share a common need—*a separate KV-cache management system, which dynamically compresses, composes, and modifies KV caches to optimize the storage and delivery of KV caches and the LLM inference based on KV caches.* We refer to such a system as a *Knowledge Delivery Network (KDN)*. As depicted in Figure 2, the envisioned architecture of a KDN consists of three main modules:

- The *storage* module stores the KV caches keyed by various texts and offline *modifies* the KV cache content such that the inference quality will be improved when the KV caches are fed to the LLM.
- The *delivery* module transmits the *compressed* KV caches from the storage device to the server running the LLM and decompresses them in GPU to be used by the LLM serving engine.
- The *blending* module dynamically *composes* multiple KV caches corresponding to different texts when these texts are put together in the context of an LLM query.

The existing LLM-serving systems [23, 26, 46] deeply couple KV caches with LLM engines. In contrast, the concept of KDN is to *separate* the management of KV caches and the LLM serving engines. Such separation will enable innovations on the storage, sharing, and use of KV caches, *without* needing to be deeply coupled with the fast-evolving LLM serving engines. By implementing these optimizations separately, a KDN serves as a critical new system module for LLM-serving ecosystems. It enables the decoupling of KV-cache management from LLM serving engines, which allows LLM serving engines to focus on fast query-driven inference, while the KDN can focus on the KV-cache-related optimizations independently.

#### 3.2 Technical Roadmap

The architecture of KDN itself does not directly address the challenges associated with KV caches; it merely breaks a



	Modularity	Efficiency	
	Time to inject new knowledge	Inference cost (\$) per request	Response delay (s) per request
Fine-tuning	10 hours	0.0052	2.63
In-context learning	0	0.0149	10.91
KV-cache learning w/ KDN	0.25 hours	0.0059	2.97

**Table 1.** Comparison between different knowledge-injection methods under a RAG setup. With KDN, KV-cache learning is 40× faster when incorporating new knowledge than fine-tuning, and achieves 3.7× faster and 2.5× cheaper during inference compared to in-context learning.

potential solution into three modules. Fortunately, we observe that emerging research efforts could lead to a sufficient design for each module, thus making KDN practical.

**KV-cache delivery:** The recent KV cache compression techniques make it possible to cheaply store and quickly load KV caches outside GPU and CPU memory. For instance, CacheGen [33] compresses the KV cache by quantizing and then encoding it into binary strings. LLMingua [24] introduces a smaller language model to identify and remove non-essential tokens in the knowledge’s text, thus reducing the size of the corresponding KV cache. H2O [45] directly removes elements in the KV cache based on their importance calculated during the inference. By combining the above techniques, the memory footprint of the KV cache can be reduced by over 10×, drastically improving the loading speed and the storage cost of the KV cache.

**KV-cache blending:** Some recent works also improve the composability of the KV caches. CacheBlend [40], for instance, enables arbitrarily composing different KV caches by recomputing the cross-attention between KV caches, where the recomputation only needs 10% computation of prefilling the full text. PromptCache [19] lets users define a prompt template with different segments, which allows each segment’s KV cache to be reused at different positions rather than prefix-only.

**Offline KV-cache editing:** By separating KV-cache management from the LLM-serving engines, KDNs open up new possibilities to improve inference quality. Recent works have shown that if the attention matrix of an LLM input is appropriately modified, the inference quality can significantly improve [13, 42]. In other words, when a KV cache is “deposited” to KDN, KDN not only stores it but also can actively influence the LLM inference quality by offline editing the KV cache and returning the edited KV cache when it is retrieved next time, all of which is done *without any change to the model itself or the input prompt*.

**Interface with LLM serving engines:** Currently, most LLM serving engines, such as vLLM [26], HuggingFace TGI [23], and SGLang [46], do not readily support the injection of externally provided KV caches as part of the LLM input. Instead, their internal implementation KV-cache management (e.g., paged memory in vLLM) is deeply cobbled with the model inference implementation. One way to deploy KDN is to augment each LLM engine with a KDN as submodule of the

LLM engine. However, as elaborated in §3.2, developing a performant KDN is a substantial undertaking, so it will cause much redundant engineering effort if each LLM engine maintains and evolves its KDNs. To avoid reinventing the wheel, the LLM serving engines could interface with a separate, shared KDN provider, via two APIs: (i) the LLM stores the KV cache with the associated text to KDN, and (ii) the LLM retrieves the KV cache from KDN using some text. Exposing these APIs is feasible, given most popular LLM engines are open-source. Still, several design questions remain. How to leverage heterogeneous links, such as NVLink, RDMA, or PCIe, to transfer KV caches? Can the APIs be shared with other LLM functions, like disaggregated prefilling?

**Early promise:** Based on these techniques, we implemented LMCache (<https://github.com/LMCache/LMCache>), a prototype of KDN. We compare the modularity and efficiency of KV-cache learning with KDN to fine-tuning and in-context learning under a RAG use case with the following setup:

- Total size of knowledge (in text): 2 million tokens.
- Each request: 8K tokens knowledge plus 2K tokens user chatting history<sup>3</sup>.
- Language model: Llama 3.1 70B [18].
- Hardware: 2 Nvidia A40 GPUs.

Modularity is measured by the time<sup>4</sup> of injecting the knowledge base into LLM, and efficiency is measured by inference cost<sup>5</sup> and response delay per request. Table 1 shows that with the KDN, KV-cache learning can be 40× faster than fine-tuning when injecting new knowledge, and it also achieves 3.7× cheaper and 2.5× faster during inference time compared to in-context learning.

## 4 Conclusion

In short, this paper makes a case for (1) the separation between the management of knowledge, in the form of KV caches, and LLM serving engines, and (2) a Knowledge-Delivery Network (KDN) as a new LLM system component that harnesses recent research development to optimize the efficiency (in speed and cost) of KV-cache-based knowledge injection. We hope this paper can inspire more research to tackle the aforementioned problems.

<sup>3</sup>We assume the chatting history cannot be pre-learned by fine-tuning.

<sup>4</sup>The time for fine-tuning the model is estimated from Llama-Adapter [43].

<sup>5</sup>The inference cost is calculated based on the AWS cloud price [12].

## References

- [1] character.ai | personalized ai for every moment of your day. <https://character.ai/>. (Accessed on 09/07/2024).
- [2] Enterprise search: an llm-enabled out-of-the-box search engine. <https://io.google/2023/program/27cce05f-df4c-4ab2-9a59-5b466bdae0f9/>. (Accessed on 09/07/2024).
- [3] How many websites are there in the world? (2024) - siteefy. <https://siteefy.com/how-many-websites-are-there/>. (Accessed on 09/08/2024).
- [4] Instruction to exclude certain information when generating answer - openai developer forum. <https://community.openai.com/t/instruction-to-exclude-certain-information-when-generating-answer/470451>. (Accessed on 09/08/2024).
- [5] Introducing chatgpt | openai. <https://openai.com/index/chatgpt/>. (Accessed on 09/07/2024).
- [6] Perplexity. <https://www.perplexity.ai/>. (Accessed on 09/07/2024).
- [7] Perplexity. <https://lmcache.github.io/2024-09-17-release/>. (Accessed on 10/14/2024).
- [8] Perplexity partners with elevenlabs to launch 'discover daily' podcast. <https://www.perplexity.ai/hub/blog/perplexity-partners-with-elevenlabs-to-launch-discover-daily-podcast>. (Accessed on 09/08/2024).
- [9] Revolutionizing operational efficiency: Unifying analytics and observability for seamless decision-making - conviva improving llm output by combining rag and fine-tuning. <https://www.conviva.com/blog/revolutionizing-operational-efficiency-unifying-analytics-and-observability-for-seamless-decision-making/>. (Accessed on 09/08/2024).
- [10] Search - consensus: Ai search engine for research. <https://consensus.app/search/>. (Accessed on 09/07/2024).
- [11] Simone Alghisi, Massimo Rizzoli, Gabriel Roccabruna, Seyed Mahed Mousavi, and Giuseppe Riccardi. Should We Fine-Tune or RAG? Evaluating Different Techniques to Adapt LLMs for Dialogue. *arXiv preprint arXiv:2406.06399*, 2024.
- [12] Amazon Web Services. Ec2 on-demand instance pricing - amazon web services, 2024. Accessed on September 07, 2024.
- [13] Anonymous. Model tells itself where to attend: Faithfulness meets automatic attention steering. In *Submitted to ACL Rolling Review - June 2024*, 2024. under review.
- [14] Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17754–17762, 2024.
- [15] Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He. Meta-learning via language model in-context tuning. *arXiv preprint arXiv:2110.07814*, 2021.
- [16] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E Gonzalez, et al. Chatbot arena: An open platform for evaluating llms by human preference. *arXiv preprint arXiv:2403.04132*, 2024.
- [17] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.
- [18] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [19] In Gim, Guojun Chen, Seung seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. Prompt cache: Modular attention reuse for low-latency inference, 2023.
- [20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [21] Aman Gupta, Anup Shirgaonkar, Angels de Luis Balaguer, Bruno Silva, Daniel Holstein, Dawei Li, Jennifer Marsman, Leonardo O Nunes, Mahsa Rouzbahman, Morris Sharp, et al. RAG vs Fine-tuning: Pipelines, Tradeoffs, and a Case Study on Agriculture. *arXiv preprint arXiv:2401.08406*, 2024.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [23] HuggingFace. text-generation-inference, 2024.
- [24] Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. LlmLingua: Compressing prompts for accelerated inference of large language models. *arXiv preprint arXiv:2310.05736*, 2023.
- [25] Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. Active retrieval augmented generation. *arXiv preprint arXiv:2305.06983*, 2023.
- [26] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626, 2023.
- [27] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- [28] Jiarui Li, Ye Yuan, and Zehua Zhang. Enhancing llm factual accuracy with rag to counter hallucinations: A case study on domain-specific queries in private knowledge-bases. *arXiv preprint arXiv:2403.10446*, 2024.
- [29] Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.
- [30] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.
- [31] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016.
- [32] Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang, Hanchi Sun, Jianfeng Gao, et al. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*, 2024.
- [33] Yuhan Liu, Hanchen Li, Kuntai Du, Jiayi Yao, Yihua Cheng, Yuyang Huang, Shan Lu, Michael Maire, Henry Hoffmann, Ari Holtzman, et al. CacheGen: Fast context loading for language model applications. *arXiv preprint arXiv:2310.07240*, 2023.
- [34] Zichang Liu, Aditya Desai, Fangshuo Liao, Weitao Wang, Victor Xie, Zhaozhuo Xu, Anastasios Kyrillidis, and Anshumali Shrivastava. Scissorhands: Exploiting the persistence of importance hypothesis for llm kv cache compression at test time. *Advances in Neural Information Processing Systems*, 36, 2024.
- [35] Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation. *arXiv preprint arXiv:2305.16938*, 2023.
- [36] Ruoyu Qin, Zheming Li, Weiran He, Mingxing Zhang, Yongwei Wu, Weimin Zheng, and Xinran Xu. Mooncake: Kimi’s kv-cache-centric

- architecture for llm serving. *arXiv preprint arXiv:2407.00079*, 2024.
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [38] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020.
- [39] Cunxiang Wang, Xiaoze Liu, Yuanhao Yue, Xiangru Tang, Tianhang Zhang, Cheng Jiayang, Yunzhi Yao, Wenyang Gao, Xuming Hu, Zehan Qi, et al. Survey on factuality in large language models: Knowledge, retrieval and domain-specificity. *arXiv preprint arXiv:2310.07521*, 2023.
- [40] Jiayi Yao, Hanchen Li, Yuhao Liu, Siddhant Ray, Yihua Cheng, Qizheng Zhang, Kuntai Du, Shan Lu, and Junchen Jiang. Cacheblend: Fast large language model serving with cached knowledge fusion. *arXiv preprint arXiv:2405.16444*, 2024.
- [41] Zhenrui Yue, Honglei Zhuang, Aijun Bai, Kai Hui, Rolf Jagerman, Hansi Zeng, Zhen Qin, Dong Wang, Xuanhui Wang, and Michael Bendersky. Inference scaling for long-context retrieval augmented generation. *arXiv preprint arXiv:2410.04343*, 2024.
- [42] Qingru Zhang, Chandan Singh, Liyuan Liu, Xiaodong Liu, Bin Yu, Jianfeng Gao, and Tuo Zhao. Tell your model where to attend: Post-hoc attention steering for llms. *arXiv preprint arXiv:2311.02262*, 2023.
- [43] Renrui Zhang, Jiaming Han, Chris Liu, Peng Gao, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, and Yu Qiao. LLaMA-Adapter: Efficient Fine-tuning of Language Models with Zero-init Attention. *arXiv preprint arXiv:2303.16199*, 2023.
- [44] Tianjun Zhang, Shishir G Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E Gonzalez. Raft: Adapting language model to domain specific rag. *arXiv preprint arXiv:2403.10131*, 2024.
- [45] Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Re, Clark Barrett, Zhangyang Wang, and Beidi Chen. H2O: Heavy-Hitter Oracle for Efficient Generative Inference of Large Language Models. In *Workshop on Efficient Systems for Foundation Models @ ICML2023*, 2023.
- [46] Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Jeff Huang, Chuyue Sun, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, et al. Efficiently programming large language models using sglang. *arXiv preprint arXiv:2312.07104*, 2023.