Retrieval-Augmented Generation for Natural Language Processing: A Survey

Shangyu Wu City University of Hong Kong, MBZUAI

> Haolun Wu McGill University, Mila

Lianming Huang City University of Hong Kong Ying Xiong* MBZUAI

Can Chen McGill University, Mila

> Xue Liu MBZUAI

Yufei Cui McGill University, Mila

Ye Yuan McGill University, Mila

Tei-Wei Kuo National Taiwan University

Chun Jason Xue MBZUAI

ABSTRACT

Large language models (LLMs) have demonstrated great success in various fields, benefiting from their huge amount of parameters that store knowledge. However, LLMs still suffer from several key issues, such as hallucination problems, knowledge update issues, and lacking domain-specific expertise. The appearance of retrievalaugmented generation (RAG), which leverages an external knowledge database to augment LLMs, makes up those drawbacks of LLMs. This paper reviews all significant techniques of RAG, especially in the retriever and the retrieval fusions. Besides, tutorial codes are provided for implementing the representative techniques in RAG. This paper further discusses the RAG update, including RAG with/without knowledge update. Then, we introduce RAG evaluation and benchmarking, as well as the application of RAG in representative NLP tasks and industrial scenarios. Finally, this paper discusses RAG's future directions and challenges for promoting this field's development.

Nan Guan City University of Hong Kong

Reference Format:

Shangyu Wu, Ying Xiong*, Yufei Cui, Haolun Wu, Can Chen, Ye Yuan, Lianming Huang, Xue Liu, Tei-Wei Kuo, Nan Guan, and Chun Jason Xue. Retrieval-Augmented Generation for Natural Language Processing: A Survey.

1 INTRODUCTION

Large language models (LLMs) [77, 117, 123, 153, 182] have achieved significant advancements in recent years and have become the cornerstone of various applications in the field of natural language processing (NLP). These LLMs are typically pre-trained on a large amount of natural language corpus and then fine-tuned on the specific downstream tasks' datasets. Recent works [4, 57, 115, 127] demonstrate the success of LLMs can be explained by the fact that LLMs act as knowledge bases, which refers to implicitly storing the learned knowledge in the parameters as internal memory and generating responses by retrieving answers from memory. To store more knowledge for better generation performance, existing works

generally enlarge the memory capacity by increasing the volume of parameters [1, 12, 58, 85].

Although existing LLMs have shown great power, several challenges still hinder the development of LLMs. One of the most prominent challenges is the hallucination problem [25, 75, 76], which refers to the tendency of LLMs to generate responses that are coherent and fluent but factually incorrect. Another big challenge is the knowledge update issue. To update the knowledge stored in the LLMs' internal memory [115, 162, 186], it is necessary to retrain/fine-tune LLMs with new data, which is a costly process. Another challenge for general LLMs is lacking domain-specific expertise [22, 148, 149, 184]. Training a domain-specific LLM demands considerable manpower for dataset collection.

To address these challenges, recent works [11, 53, 95] have proposed leveraging an external knowledge database to augment LLMs, known as Retrieval-Augmented Generation (RAG). By supplying LLMs with retrieved relevant factual information, the hallucination problem can be alleviated to some extent. Besides, the knowledge update issue can also be addressed by updating the external knowledge database, which can augment LLMs with up-to-date knowledge. RAG can also convert a general LLM into a domain-specific LLM by constructing and utilizing a domain-specific knowledge database. Therefore, RAG plays an important role in augmenting the functionality of LLMs, making them more accurate, knowledgeable, and reliable in a wide range of applications.

Contributions: This paper reviews all techniques involved in RAG for natural language processing. Although there are several survey papers for RAG [44, 63, 96, 180, 189], our survey still has some key insights,

- This paper systematically introduces each component of RAG, including details about the retriever from building to querying and techniques of the retrieval fusions with tutorial codes.
- (2) This paper exhibits different RAG training strategies, including RAG with/without datastore update.
- (3) This paper further discusses RAG evaluation and benchmarking, as well as the applications of RAG on downstream NLP tasks and practical NLP scenarios.

^{*}Corresponding author.



Figure 1: The overview of retrieval-augmented generation for natural language processing. The inputs as queries are fed into both the retriever for retrieval knowledge and the generator for outputs. There are three kinds of retrieval fusions, including query-based fusion, logits-based fusion, and latent fusion.

(4) This paper finally identifies promising future directions for exploring and the main challenges for addressing.

The remainder of this paper is organized as follows. Section 2

gives an overview of RAG. Section 3 and Section 4 comprehensively

introduce all technical details used in retrievers and retrieval fu-

sions. Section 6 presents how to train the RAG with/without new

knowledge. Section 8 presents the techniques used in representative

NLP tasks. Section 9 shows the applications of RAG in practical

NLP scenarios. Section 10 discusses the future directions of RAG.

4. The overall process is formulated as follows:

$$\mathsf{Retriever}(q, D) \to R_q \tag{1}$$

Generator(Retrieval Fusion (q, R_q)) $\rightarrow answer$ (2)

As shown in Figure 1, RAG typically consists of three modules, the retriever, the generator, and retrieval fusions.

Retriever module usually comprises three components: an encoder for encoding inputs into embeddings, an efficient indexing that supports approximate nearest neighbor search, and a datastore for storing external knowledge in the form of key-value pairs. The main challenge in the retriever module is *finding the optimal trade-off between retrieval efficiency and retrieval quality*. The retrieval efficiency refers to how fast the relevant information can be obtained, which involves accelerating encoding, efficient indexing, batch querying in the datastore, etc. The retrieval quality refers to how relevant the information can be retrieved, which involves chunk representation learning, advanced approximate nearest neighbor search algorithms, etc.

Retrieval Fusions aims to leverage the retrieved information to augment the generation. These fusion techniques can be categorized into three major types: query-based fusion, latent fusion, and logits-based fusion. The query-based fusion augments inputs with retrievals before feeding them into the generators. The logits-based fusion focuses on the output logits of generators and fuses the retrievals logits for more robust logits. The latent fusion refers to

Section 11 makes a final conclusion of this paper.2 OVERVIEW OF RETRIEVAL-AUGMENTED

GENERATION

This section gives an overview of Retrieval-Augmented Generation (RAG) for NLP. An RAG system utilizes the external knowledge base D to enhance the generation system. Taking external documents as an example, D consists of external documents, each of which contains a set of chunks $c_i \in C_i$. These chunks are transformed into vector embedding using an embedding model. When inputting a query q, which is embedded as a vector, then the retriever in the RAG system retrieves top-k chunks $R_q = \{r_1, r_2, \ldots, r_k\}$ most relevant to the query q. The RAG system can use different fusion methods to fuse the retrieved chunks, which is discussed in Section



Figure 2: Two stages of using the retriever.

introducing retrieval representations into the latent representations of generators, thus implicitly improving the models' performance.

Generator modules can be classified into two branches of generators: default generators and retrieval-augmented (RA) generators. The default generators include most pre-trained/fine-tuned large language models, such as GPT-series models [12, 123, 129, 130], Mistral models [77], and Gemini-series models [5, 117, 135]. The RA generators refer to the pre-trained/fine-tuned generators that consist of modules for fusing retrievals, such as RETRO [11, 157] and Enc-Dec [101]. Those generators generate responses or make predictions.

The workflow of RAG involves three steps: 1) retrieving the relevant information from external databases based on given inputs; 2) fusing the retrieved information with inputs or intermediate states based on the fusion techniques; 3) making predictions by generators based on the input and corresponding retrievals.

3 RETRIEVER

Figure 2 shows the two stages for using the retriever, which involves first building the retriever and then querying the retriever. The following sections will introduce details about each stage.

3.1 Building the Retriever

This section will explain how to build a retriever using a large natural language corpus. As shown in Figure 2 (a), the process involves three steps: chunking corpus, encoding chunks, and building the vector database. Specifically, building the vector database includes building the ANN index and storing the data with key-value pairs.

3.1.1 Chunking Corpus. Chunking techniques generally refer to dividing large documents into small text chunks [11, 14, 46, 70, 120], which is an indispensable key step in the process of building the retriever. The intuitions behind chunking techniques are, (1) The texts or embeddings used for the indexing should be semantically independent, containing one core idea for models to encode. Short

texts are more likely to be ambiguous, for example, the word "apple" can refer to a fruit or a company. (2) Encoding a long sequence document would result in considerable resource overheads when using existing transformer-based models, while processing shorter text chunks can significantly accelerate the encoding process and save memory costs. Therefore, the main challenge of the chunking techniques is to find the best chunking size to make a better tradeoff between text semantics and encoding efficiency.

To solve the above challenge, three key points need to be considered when determining the chunking size:

- Task's preference. Different tasks may benefit from different kinds of retrieval chunks. For example, question-answer tasks may prefer short phrases, while summarization tasks may prefer long documents.
- (2) Encoder's preference. Different encoder models have varying encoding capabilities on texts with different lengths. For example, models in the sentence-transformer [136] behave better on a single sentence, while the text-embeddingada-002 [122] is good at longer texts.
- (3) Query's preference. The length of the user's queries should be aligned with the chunking size, which implicitly aligns the amount of contextual information in chunks with that in queries, thus improving the relevance between queries and retrievals. For example, a retrieval database built on short phrases may be useless for queries with long documents.

Overall, there is no golden rule for determining the chunking size, and it depends on the specific RAG scenarios.

There are basically three types of chunking techniques, including **the chunking with fixed length**, **the semantic chunking**, and **the content-based chunking**. Chunking with fixed length is the simplest way to split documents sequentially using a length hyperparameter. The semantic chunking cuts documents based on semantics, such as the period character or the newline character that represents the end of the sentence. Existing state-of-the-art natural language processing toolkits, such as NLTK [121] and spaCy [35], have provided convenient sentence-cutting methods. The contentbased chunking segments documents according to the unique structural characteristics. For example, electronic medical records can be easily segmented based on the sections, or programming codes can be segmented based on function blocks.

3.1.2 Encoding Chunks. Encoding refers to numericalizing textual chunks as vector representations (embeddings). These embeddings generally capture the semantics of the chunks, enabling the retriever to perform similarity searches based on content relevance rather than just keyword matching.

According to the sparsity of the embeddings, there are two kinds of encoding methods, i.e., **sparse encoding** and **dense encoding**. The sparse encoding represents text by creating high-dimensional vectors where most elements are zero. The basic sparse encoding is one-hot encoding [54], which represents a word with a highdimensional vector as large as the vocabulary table size but only marks the value corresponding to the presence of the word as one. The embeddings produced by such encodings are called the one-hot vector. Other common sparse encodings include:

- (1) Bag of Words (BoW) [55]. This encoding improves onehot encoding by replacing the zero-one counting with the frequency counting. However, BoW ignores the syntax and word order in the documents and focuses on statistical information, thus only expressing limited semantics.
- (2) **Term Frequency-Inverse Document Frequency (TF-IDF)** [131]. This encoding not only counts the occurrence (frequency) of each word but also adjusts these counts based on how common the word is across all documents (inverse document frequency). TF-IDF helps emphasize words that are more descriptive of the document's content.
- (3) BM25 [139] is a probabilistic ranking algorithm used in information retrieval to estimate the relevance of documents to a search query by balancing term frequency, inverse document frequency, and document length normalization, ensuring robust scoring even for long or short documents. BM25 focuses on lexical matches and is computationally efficient, making it a cornerstone of traditional search engines.

Sparse encoding is an efficient way to encode textual chunks. However, such encoding methods may not capture deeper semantic meanings well.

The dense encoding generates vectors where each dimension can capture a range of semantic features, and most elements are nonzero floating points. The dense embeddings are generally produced by (deep) neural network models,

- (1) BERT [31] and Variants. Bidirectional Encoder Representation from Transformers (BERT) is a typical pre-trained transformer model, generating dense semantic embeddings that capture the contextual information. Other BERT variants, such as RoBERTa [107], DistilBERT [142], and ELEC-TRA [21], further improve the semantic representations with advanced learning techniques.
- (2) Siamese Encoders. This is a type of neural network designed to learn the similarity between inputs, which is

usually trained with contrastive learning. Existing stateof-the-art siamese encoders are DPR [86], SimCSE [43], Contriever [71].

(3) LLM-based Encoders. This type of encoder benefits from the powerful representation capability of LLMs. LLMs, which contain billions of parameters and are pre-trained on vast amounts of data covering a wide range of topics, have advanced semantic language understanding capabilities. Typical LLM-based encoders are text-embedding-ada-002 [122], bge-embedding [172], mxbai-embedding [144], MedCPT [83].

Compared to sparse encoding, dense encoding leverages deep neural networks, especially transformers [155], to capture broader linguistic and semantic information. Dense encodings are widely used in most representation scenarios. Moreover, some works also utilized hybrid methods to encode text for leveraging both lexical and semantic information [94, 109].

3.1.3 Building the Index. Indexing in the vector database aims to accelerate the search process for data similar to high-dimensional query embedding. Unlike common indexing in databases, indexing in the vector database mainly focuses on supporting efficient approximate nearest neighbor (ANN) search [34, 51, 84] rather than transaction operations like insertion, deletion, and update. The key challenge of indexing is making a good trade-off between search quality and search efficiency. To solve the challenge, there are various specific optimizations in both algorithmic aspects and systematic aspects to be explored, including choices of similarity metrics, dimension reduction (DR) on embeddings, advanced ANN indexing, system-level optimizations, hardware-aware optimization, and so on. Due to the page limits, this section discusses the optimizations that significantly affect the search quality and efficiency.

Choice of Similarity Metrics. The similarity metrics are the basic components in the retriever, which measures the degree of relevance between query embeddings and chunk embeddings. The similarity metrics would affect the search quality. Typical similarity metrics include cosine similarity, Euclidean similarity, and Manhattan distance.

Dimension Reduction on Embeddings. Reducing the dimensionality of embeddings can improve search efficiency but at the risk of harming the semantic representations. The basic but effective dimension reduction (DR) is the principal component analysis (PCA). The PCA is a simple statistical technique that transforms the original data into a new coordinate system while retaining the most important features. Another popular and advanced dimension reduction is locality-sensitive hashing (LSH). LSH significantly reduces the dimensionality by mapping the data into buckets but preserves the similarity of the original input data. The intuition behind LSH is that the nearest neighbors will be mapped into the same buckets. Unlike LSH, product quantization (PQ) [74] is another popular and effective DR technique for ANN search. The core idea of the PQ is to divide the high-dimensional space into smaller, independently quantized subspaces. Each subspace creates a codebook of different quantized integers to form the representative and compact vectors. The above techniques enable efficient storage and fast approximate search but may lose semantic information. Recent work [19] proposed a new technique named AutoCompressor that

reduces the dimension of embeddings by compressing the original context into semantically shorter embeddings.

Advanced ANN Indexing. ANN Indexing generally refers to the methods or structures used to organize and manage data so that the approximate-nearest-neighbor search process is optimized for retrieval quality and retrieval efficiency. This paper will introduce several advanced ANN indexing techniques.

- (1) The InVerted File system with Product Quantization (IVFPQ) [34] is a simple but effective indexing framework that combines two powerful techniques to enable an efficient and scalable ANN search process. The main idea of IVFPQ is first to cluster the data for coarse-grained partition and then to compress the data within each cluster into subvectors for fine-grained quantization. The coarse-grained clustering (the IVF component) significantly reduces the search space, while the fine-grained quantization (the PQ component) ensures a high retrieval performance.
- (2) The Hierarchical Navigable Small World (HNSW) [111] uses a hierarchical graph structure to perform ANN search in high-dimensional spaces efficiently. Specifically, HNSW treats high-dimensional vectors as nodes and connects them with their nearest neighbors. The multi-layer graph structure is determined probabilistically to ensure fewer nodes at higher layers for efficient search.
- (3) Tree-based Indexing aims to organize high-dimensional vectors in tree-liked structures, such as KD-Trees [134], Ball Trees [66] and VP-Trees [106]. Typical tree-based indexing is Approximate Nearest Neighbors Oh Yeah (Annoy) [150], which uses a forest of trees built based on random projections to separate the vector space into multiple hyperplanes for efficient ANN search.

3.1.4 Building the Datastore with Key-Value Pairs. The datastore used in the vector database is a specialized database that stores and manages data as a collection of key-value pairs, where keys are the unique identifiers of high-dimensional embeddings and values are the domain-specific knowledge. Since the amount of the data stored in the datastore may be quite large, the storage engine, such as LMDB [108] or RocksDB [37], should be capable of efficient retrieval and data persistence. The key point in the datastore for ANN search is what should be used to be stored as values. For example, for question-answer tasks, when adding retrievals to prompts, the naive but effective way is to store the question embedding as the key and question-answer pairs as the value. This can help the generation process as retrievals are used as demonstrations for models. Recent works have proposed various state-of-the-art vector databases including the indexing and the datastore, such as Milvus [50, 159], FAISS [34, 84], LlamaIndex [103], etc.

3.1.5 Code Demonstrations. Algorithm 3.1 shows detailed steps to build the retriever. Lines 2-8 present the chunking and the encoding process for a natural language corpus containing multiple documents. In line 6, algorithm 3.1 takes the concatenation of the current chunk and the next chunk as the value. Notably, the choice of value can vary for different tasks. Another practical issue is that the memory cost of all keys and values may exceed the memory capacity of

Algorithm 3.1 Building the retriever.

Input: A natural language corpus $D = \{d_1, \ldots, d_n\}$ for building the knowledge database, an encoder \mathcal{E} for encoding chunks.

Output: The index I and the key-value store S.

1: $\mathcal{K} = \{\}, \mathcal{V} = \{\};$

2: for $d_i \in D$ do

3:

4:

5:

- $c_i^1,\ldots,c_i^m = Chunk(d_i);$ /* Split each data d_i */
- for j from 1 to m do
- $e_i^j = \mathcal{E}(c_i^j);$ /* Encode each chunk $c_i^j *$ / Add e_i^j into \mathcal{K} and $c_i^j + c_i^{j+1}$ into $\mathcal{V};$ /* Take next chunk as an 6 exampless */
- The \mathcal{K} and \mathcal{V} persist in the storage (e.g., SSD) if necessary; 7:
- 8: end for
- 9: end for
- 10: Build the index I with \mathcal{K} :
- 11: Store \mathcal{K} and \mathcal{V} into the key-value store \mathcal{S} ;

12: return I and S;

Algorithm 3.2 Query the retriever.

Input: A query input q, an encoder \mathcal{E} for encoding chunks, the index I, the key-value store S, the parameter k.

Output: Top-k nearest neighbor knowledge.

- 1: $e = \mathcal{E}(q);$
- 2: $\{idx_1, \ldots, idx_k\} = I.Search(e, k);$ /* Search the top-k nearest neighbors */
- 3: $\{v_1, \ldots, v_k\} = S.Fetch(\{idx_1, \ldots, idx_k\}); /*$ Fetch the values of the neighbors */
- 4: $\{v_{j_1}, ..., v_{j_k}\} = PostProcess(\{v_1, ..., v_k\})$
- 5: **return** $\{v_{j_1}, \ldots, v_{j_k}\};$

the server in the practical scenario. Thus, it is recommended that the keys and values persist in the storage if necessary.

3.2 **Querying the Retriever**

This section will explain how to query the pre-built retriever, which basically includes three steps as shown in Figure 2(b): encoding queries, ANN search, and post-processing.

3.2.1 Encoding Queries and ANN Search. To align with the prebuilt embedding space, the retriever uses the same encoder to encode queries during the querying stage. The ANN search leverages the pre-built indexing and datastore to find similar data via ANN searching algorithms and then retrieves the corresponding values.

Searching the index refers to searching the pre-built index, finding the top-k nearest neighbors, and returning the unique identifiers of k nearest neighbors. The nearest neighbor search process depends on indexing algorithms or structures. Taking IVFPQ as an example, the search process first compares the query embedding with cluster embeddings and selects several candidate clusters for further search. Then, within each cluster, the search process performs the same product quantization on the query embedding and finds the top-k nearest neighbors based on the distance. Finally, the search process merges all nearest neighbor candidates and re-orders all candidates for the final top-k nearest neighbors.

Retrieving values from datastore fetches the corresponding values based on the nearest key identifiers.



Figure 3: The categories of fusion methods in RAG.

Algorithm 4.1 Query-based Fusions.

Input: A query input *q*, top-*k* nearest neighbor knowledge $\{v_1, \ldots, v_k\}$, an encoder \mathcal{E}_f and a decoder \mathcal{D}_f for feature concatenation, the generator \mathcal{G} for text concatenation.

Output: Generated response *y*.

1: if Use the text concatenation then

2:
$$x = v_1 \oplus \ldots \oplus v_k \oplus q$$
; /* Concatenate neighbor texts and query */

3: $y = \mathcal{G}(x);$

- 4: else
- 5: $e_q = \mathcal{E}_f(q), e_{v_j} = \mathcal{E}_f(v_j), j \in \{1, \dots, k\};$
- 6: $e_x = e_q \oplus e_{v_1} \oplus \ldots \oplus e_{v_k}; /^*$ Concatenate embeddings of neighbors and query */
- 7: $y = \mathcal{D}_f(e_x)$
- 8: **end if**
- 9: **return** *y*;

3.2.2 *Post-Processing*. The post-processing involves a set of techniques after the initial retrieval step. These techniques aim to refine, enhance, or adapt the retrievals based on the specific task objectives. This section will list some typical post-processing techniques.

Reranking aims to reorder the retrieved knowledge based on task-specific objectives. The intuition is that the knowledge is retrieved based on task-agnostic metrics, such as Euclidean distance. Existing reranking methods [20, 60, 92, 156] mostly design different architectures or strategies to reorder the retrieved knowledge.

3.2.3 *Code Demonstrations.* After building the retriever, this section demonstrates the detailed steps of querying the retriever to obtain the top-*k* nearest neighbor knowledge in algorithm 3.2, including encoding the query (line 1), performing the approximate nearest neighbor search (line 2), and fetching the knowledge for fusion (line 3). These three steps depend on the specific APIs of encoders, indexing, and datastore. After obtaining the top-*k* retrievals, optimizations for post-processing are applied (line 4).

4 RETRIEVAL FUSIONS

Retrieval fusions refer to how to leverage the retrieved knowledge to improve generators' performance. Basically, there are three types of retrieval fusions: query-based fusions, logits-based fusions, and latent fusions. Figure 3 shows the detailed categorization of fusions and representative works of each retrieval fusion in RAG.

4.1 Query-based Fusion

The simplest and most direct fusion technique is query-based fusion, which integrates the retrieved information with input queries to generate responses. The query-based fusion can be further categorized into two sub-classes according to the type of concatenated information, i.e., text concatenation and feature concatenation.

Text concatenation involves performing query-based fusion with raw texts, making it particularly suitable for contemporary LLMs like GPT-4. These models function as black-box systems with limited interaction capabilities, typically offering only API access to users. Existing works [53, 95, 133] directly concatenate the input with the top-*k* retrieved sentences/documents to form the query for generators. To better use the in-context learning capability of LLMs, some works [36, 98, 156, 161] design effective prompt templates to integrate retrieved information and inputs. To address the issue of lengthy inputs after concatenating retrievals, recent studies [6, 104, 110, 166, 176] have introduced methods for assigning importance weights to elements within the retrieved knowledge base and filtering out less relevant contexts based on these weights.

The feature concatenation involves merging the encoded retrievals with the input features. A simple yet effective approach is FID [72], which first encodes the retrieved passages into sparse or dense representations and then takes the concatenated features as the input for a generator. The state-of-the-art performance of the FID demonstrates the efficacy of feature concatenation. The follow-up works [27, 52, 73, 105, 141] further improve the FID by jointly tuning the retriever and the encoder, which can enhance the retrieved knowledge's representations. Besides, Chen et al. [16] concatenate the representations of related knowledge as demonstrations for prompt learning, yielding better generalization.

Algorithm 4.1 presents how to leverage query-based fusions to fuse retrieved knowledge. For those using text concatenation [53, 133], algorithm 4.1 first concatenates the retrieved texts and inputs (line 2), then feeds the concatenated input into the generator. Notably, since there is a limit to the maximum input length of existing language models, concatenating too many retrievals would result in a truncation of the concatenated input, which may cut the given input. Therefore, designing the prompt template is the key step for this branch of work. For those using feature concatenation [52, 72], algorithm 4.1 first leverages an encoder to obtain the feature (line 5), then concatenates the feature of input and retrievals (line 6), finally passes the concatenated feature into a decoder model (line

Algorithm 4.2 Logits-based Fusions.

Input: A query input q, top-k nearest neighbor knowledge $\{v_1, \ldots, v_k\}$, the generator \mathcal{G} .

Output: Generated response y. 1: $y_q = \mathcal{G}(q);$ 2: **for** *j* from 1 to *k* **do** 3: $y_{v_i} = \mathcal{G}(v_j)$ 4: end for 5: if Use ensemble then $y = \lambda \sum_{i} y_{v_i} + (1 - \lambda) y_q;$ 6: 7: else 8: $\lambda_t = Calibrate(y_q, y_{v_1}, \dots, y_{v_k})$ $y = \lambda_t \sum_j y_{v_j} + (1 - \lambda_t) y_q;$ 9: 10: end if 11: return y;

7). This branch of work generally incurs high memory costs due to the long sequence length.

4.2 Logits-based Fusion

The logits-based fusion refers to incorporating the retrieved knowledge into the output layers. Basically, retrieved knowledge would be fed into the same model to obtain the logits for enhancing or calibrating the predictions. Therefore, logits-based fusion can be categorized into two branches, i.e., ensemble-based fusion and calibration-based fusion.

Ensemble-based fusion treats the logits from the retrieved knowledge as part of an ensemble of predictions. Such ensemble-based fusion can significantly improve the generalization and robustness of the model [87, 88, 175]. One notable work of ensemble-based fusion is kNN-LM [88], which aggregates the logits of the top-*k* nearest neighbors' targets and then interpolates the final predictions. Similar to kNN-LM, Khandelwal et al. [87] propose kNN-MT to enhance the machine translation using retrievals' logits, which is also followed by a branch of works [68, 190].

Different from ensemble-based fusion, calibration-based fusion uses the logits from the retrieved knowledge as a form of calibration for the model's predictions. Specifically, Jiang et al. [78] propose a confidence-enhanced kNN-MT that refines the kNN distribution and interpolation weights with the neural machine translation confidence. Li et al. [97] propose to leverage the source context to calibrate the retrieval-augmented neural machine translation.

Algorithm 4.2 demonstrates the detailed steps of using the logitsbased fusion to integrate the retrieved knowledge. This branch of work first treats retrievals as similar data to augment the model (lines 2-4). For ensemble, algorithm 4.2 leverages a hyperparameter to fuse the retrieval logits and the output logits (line 6). For calibration, algorithm 4.2 dynamically determines the parameter based on the retrieval logits and the output logits (line 8). Then, algorithm 4.2 performs the same fusion with the computed parameter (line 9).

4.3 Latent Fusion

The latent fusion investigates merging the retrieved knowledge into the hidden states of generators for a better generation. Based on the introduction method, latent fusion can be further classified into two categories: attention-based and weighted-addition.

Algorithm 4.3 Latent Fusions.

Input: A query input <i>q</i> , top- <i>k</i> nearest neighbors {	v_1, \ldots, v_k , the encoder
\mathcal{E} , the generator \mathcal{G} containing l pairs of mode	ules $\{(\mathcal{M}_1^A, \mathcal{M}_1^F), \ldots\}$
where \mathcal{M}_{i}^{A} and \mathcal{M}_{i}^{F} are the attention modul	le and the FFN module
at layer <i>i</i> , \mathcal{M}^{C} is the cross-attention module i	used in attention-based
latent fusions.	
Output: Generated response <i>u</i>	
1: if Use the attention then	
2: $h_{r}^{F} = a$	
3: for <i>i</i> from 1 to <i>l</i> do	
4: $h_i^A = \mathcal{M}_i^A(h_{i-1}^F);$	
5: $e_{v_1}, \dots, e_{v_k} = \mathcal{E}(v_1, \dots, v_k, h_i^A)$	
6: $h_i^R = \mathcal{M}_i^C(h_i^A, e_{v_1}, \dots, e_{v_k});$	/* Cross-attention */
7: $h_i^F = \mathcal{M}_i^F(h_i^R)$	
8: end for	
9: $y = LM_HEAD(h_l^F)$	
10: else	
11: $e_{v_1}, \ldots, e_{v_k} = \mathcal{E}(v_1, \ldots, v_k)$	
12: $h_0^F = q;$	
13: for <i>i</i> from 1 to <i>l</i> do	
14: $h_i^A = \mathcal{M}_i^A(h_{i-1}^F);$	
15: $h_i^R = h_i^A + \frac{1}{k} \sum_j w_j e_{v_j}$	/* Weighted sum */
16: $h_i^F = \mathcal{M}_i^F(h_i^R)$	
17: end for	
18: $y = LM_HEAD(h_1^F)$	
19: end if	
20. return 11:	

One notable contribution of attention-based fusion is the Retrieval-Enhanced Transformer (RETRO) [11]. RETRO represents a pioneering effort in pre-training retrieval-based LLMs, introducing a new cross-attention module to integrate retrieved knowledge directly into the model's hidden states. A significant finding from this work is demonstrating a scaling law for the retrieval database, where RETRO, with a 2 trillion token database, attains performance comparable to that of major models like GPT-3 and Jurassic-1, albeit with 25 times fewer parameters. Customizing the transformer model in RETRO highlights the potential of pre-trained, retrieval-enhanced architectures in improving the efficiency and scalability of LLMs.

In addition to RETRO, other studies [13, 28, 101, 167, 170] have contributed to the field by leveraging new attention modules to introduce external knowledge. Typically, Li et al. [101] have extended the RETRO model by decoupling the context encoding from the model inference. Wu et al. [170], Wang et al. [163] store the hidden attention keys and values into external memory and retrieve the knowledge from the memory using an attention mechanism.

Due to the high complexity of the attention mechanism, another branch of work adopts lightweight (weighted) additions to introduce retrieved knowledge. Fevry et al. [40] propose the EAE model that retrieves top-k related entities' embeddings from a learnable external memory and adds entities' embeddings to the hidden states of the model. Wu et al. [169] propose ReFusion, which explores various learnable reranking schemes to first re-weight the retrieved knowledge's embeddings, then use weighted addition to incorporate them into the hidden states of the model. Those approaches signify a growing trend towards models that dynamically select and integrate relevant knowledge, paving the way for more sophisticated and nuanced language generation and understanding.

Algorithm 4.3 shows the steps of using latent fusion to introduce the retrieved knowledge into the hidden states of the generator. For attention-based latent fusion, algorithm 4.3 first encodes the retrievals with the output states of the attention module (line 5), then uses a cross-attention module to fuse the retrieval features into the hidden state (line 6). Different from attention-based latent fusion, weighted-addition-based latent fusion adopts a more lightweight way to incorporate retrieved knowledge (lines 10-19). Algorithm 4.3 first encodes the retrievals before feeding them into the generator (line 11), which can be done offline and directly stored as values in the datastore. Then, algorithm 4.3 learns a set of weights to add the retrieval features on the hidden states of generators (line 15).

4.4 Comparison of Different Fusion Methods

Each fusion method exhibits distinct advantages and limitations. The query-based fusion method is a simple and straightforward approach that preserves the raw information from retrieved information. However, it significantly increases the input sequence length, resulting in higher computational costs. Additionally, this kind of method suffers from limited contextual integration capabilities and scalability, as it struggles to effectively combine retrieved information with the input context.

The logits-based fusion method combines the output logits (or probabilities) of the base LLM with those from a separate model (or shared LLM) that processes the retrieved information. This approach is computationally efficient, as it can process retrievals and inputs in batches, avoiding increasing the input sequence length. It decouples the retrieval and generation processes, enabling independent improvements to each component. Nevertheless, this kind of method is limited by its shallow contextual integration, as fusion occurs only at the output level. Consequently, it may underperform in complex reasoning tasks that require fine-grained interaction between the input and retrieved information.

In contrast, the latent fusion method integrates retrieved information directly into the hidden states (or intermediate representations) of the LLM. This approach is both scalable and efficient, as it avoids expanding the input sequence length. Moreover, it facilitates deep, context-aware integration of retrieved information and supports fine-grained control over how the information is fused. However, the latent fusion method necessitates significant architectural modifications and careful training to achieve optimal performance. In summary, the choice of fusion method depends on the specific application, available resources, and the desired trade-off between simplicity and performance. Each method offers unique strengths and weaknesses, making them suitable for different use cases and operational constraints.

5 GENERATORS

This section introduces representative generators, which are generally pre-trained on large datasets. Existing generators are mostly large language models that adopt or modify the transformer-based architecture [155]. For example, DeepSeek [49], Llama-series models [153, 154], GPT-series models [12, 123, 129, 130], and Geminiseries models [5, 117, 135] remove all encoder modules, retaining only the decoder module, which includes an attention module and a feed-forward network module. Other advanced techniques, such as root mean square layer normalization [183], rotary position embedding [151], and group query attention mechanisms [3], have been incorporated into the design of existing large language models to enhance their performance.

Existing generators can be generally categorized into two groups, close-sourced LLMs [12, 123] and open-sourced LLMs [49, 153, 154]. The formers can only adopt query-based fusions to introduce external knowledge. While the latter can adapt to all kinds of fusions. Especially generators using the latent fusion typically would train the generators from scratch. Those generators integrate novel modules designed for retrieved information fusion, such as cross-attention-based module [11, 101, 170]. Their training paradigm combines pre-training on massive textual datasets with integrating external retrieved knowledge to establish robust information grounding. As detailed in Section 4.3, those generators enable dynamic, context-aware retrieved information incorporation while further improving the LLM's generative capabilities.

6 RAG TRAINING AND DATASTORE UPDATE

This section introduces RAG training, which can be categorized into two main classes: **RAG without datastore update** and **RAG with datastore update**. The former refers to the case where only trainable parameters in each module of RAG would be updated, and the knowledge in the datastore would remain the same during the training stage. The latter refers to the case where the knowledge in the datastore would be updated, then each module's parameters in RAG would be updated in a similar way as the former case.

6.1 RAG without Datastore Update

The goal of training RAG without datastore update is to update the knowledge stored in the short-term memory of generators based on the existing knowledge datastore. As shown in Figure 4 (a)-(c), there are three training cases, i.e., training the retriever, training the generator, and jointly training the retriever and generator.

6.1.1 Training retriever. Considering the case of no datastore update, training the retriever generally refers to training the retriever encoder and rebuilding the indexing. Since sparse encodings rely on statistical methods without parameters, training the encoder pertains only to dense encoding methods. Different training methods may have different goals, such as improving the semantic representations, accelerating the encoding process, or learning the domainspecific representations. The first two goals are often achieved by replacing the original encoder with a more powerful or tiny encoder, such as DistilBERT [142], or TinyBERT [81]. The last requires training the original encoder on the domain-specific corpus. RE-PLUG [146] updated the retriever by minimizing KL divergence between retrieval and LM scores, and asynchronously refreshing the datastore index. After training the retriever encoder, the vector database's embeddings that serve as keys will also change. Thus, all indexes should be rebuilt with new embeddings. Besides, if the encoder remains unchanged, the indexing can be updated using new ANN searching algorithms or re-tuning the hyperparameters. After the retriever is trained, it can be directly incorporated into the RAG without updating the generator.



Figure 4: Different RAG training strategies with/without datastore update.

6.1.2 Training generator. Training the generator involves updating its parameters or those in the retrieval fusion modules. Since the generator is generally an LLM, training the LLM is a resource- and time-consuming process. Fortunately, several parameter-efficient fine-tuning techniques, such as LoRA [61], are proposed to address the fine-tuning problem of LLMs. Although the parameters in the retrieval fusion modules are less than those in the generator, only fine-tuning those parameters may encounter some training problems, such as low convergence and overfitting. Jointly tuning the parameters in the generator and the retrieval fusion modules is a better way to train the generator and the retrieval fusion modules if there are sufficient and powerful resources.

6.1.3 Jointly training the retriever and generator. Apart from independently training the retriever and the generator, jointly training the retriever and the generator can be another good choice for better performance on downstream tasks. The primary challenges involve computational complexity from large-scale retrieval and marginalization over latent documents, alongside training instability due to potential retrieval collapse and interdependent retriever-generator feedback loops. Additionally, sparse gradients from discrete retrieval steps and approximation errors from fixed top-*k* document selection limit effective end-to-end optimization. Typically, complex indexes, such as FAISS [34], are not a suitable choice during the fine-tuning stage.

Existing works generally leverage the complex indexes to preselect a small subset of nearest neighbors as candidates, then choose the final top-*k* nearest neighbors by performing the matrix-multiplication operations. REALM [53] and RAG [95] both unify retriever and language modeling by treating retrieved documents as latent variables and optimizing end-to-end via gradient descent. They both used Maximum Inner Product Search (MIPS). RAG [95] only updated

the encoder in the retriever, but REALM [53] also asynchronous refreshed the indexing. Atlas [73] jointly pre-trained the retriever and generator, using different loss functions (such as attention distillation and perplexity distillation) to optimize the retriever. Besides, it also asynchronous refreshed the indexing. Experimental results in these works demonstrate that joint training is an end-to-end optimization that can lead to better coordination between the retriever and the generator and improve the contextual understanding of the generator.

6.2 RAG with Datastore Update

As shown in Figure 4 (d), the scenario involves two stages: updating the knowledge database, then training the retriever and the generator. There are three cases for updating the knowledge database, i.e., updating with trainable embeddings, updating with new values, and updating with new corpus. In the first case, values generally are trainable embeddings and are simultaneously/asynchronously updated with parameters in the RAG [16]. The last two cases usually refer to updating the knowledge database with up-to-date information. Taking question-answer corpus as an example, updating with new values refers to updating the answer to existing questions, while updating with new corpus refers to adding new questionanswer pairs. To update the value of existing keys requires first querying the existing key-value pairs and then performing in-place updates. For a new corpus, the datastore first needs to perform insertion operations, then rebuilds or updates the indexes as new keys are added. After updating the datastore, training the retriever and the generator is similar to RAG without datastore updates. However, this training step is not always necessary, thanks to the in-context learning capability of LLMs.

7 RAG EVALUATION AND BENCHMARK

Retrieval-Augmented Generation (RAG) systems integrate the capabilities of large language models (LLMs) with external knowledge retrieval to advance text generation tasks in natural language processing (NLP). Evaluating RAG systems requires a dual focus on retrieval quality and generation performance to ensure the system effectively leverages external knowledge and produces accurate, contextually appropriate outputs. Recent research has introduced domain-specific benchmarks and comprehensive evaluation frameworks to address these challenges.

For retrieval quality, Pipitone et al. [128] proposed a benchmark tailored for RAG systems in the legal domain, employing key retrieval metrics such as Precision@K and Recall@K to assess the relevance and ranking of retrieved documents. Adlakha et al. [2] expanded the evaluation scope by incorporating correctness, faithfulness, and human evaluation to assess retrieved information, alongside novel token-overlap metrics for finer-grained analysis.

On the generation side, Hui et al. [69] introduced a benchmark suite for RAG systems in real-world document analysis, utilizing metrics such as accuracy, F1 score, and exact match score to evaluate answer quality. Similarly, Xiong et al. [174] developed a medical RAG benchmark to systematically compare the performance of medical RAG systems, emphasizing domain-specific evaluation criteria.

Recent advancements have further enriched the evaluation landscape by proposing more systematic and comprehensive methodologies. Chen et al. [15] identified four fundamental RAG abilities—noise robustness, negative rejection, information integration, and counterfactual robustness—as key evaluation dimensions. Saad-Falcon et al. [140] introduced ARES, an automated RAG evaluation system, which assesses RAG systems using context relevance, answer faithfulness, and answer relevance. Building on this, Friel et al. [41] extended the evaluation framework by incorporating additional metrics such as context utilization and answer completeness, providing a more holistic assessment of RAG system performance. These efforts collectively highlight the growing emphasis on rigorous, domain-specific, and multi-dimensional evaluation methods to ensure RAG systems meet the demands of diverse applications while maintaining high retrieval and generation quality standards.

8 TASKS

This section lists several classical tasks in the NLP domain and introduces advanced RAG techniques used to solve these tasks.

8.1 Language Modeling

Language modeling is the task that requires the prediction of the probability distribution of the next word or character given a sequence of words or characters, which is also named the next-token prediction task. Language modeling has become the fundamental task for pre-training large language models, which can measure the models' generation capability using the perplexity metric. The formal definition is as follows: given such a sequence of tokens x_1, \ldots, x_n called *Prefix*, the language modeling task aims to model

its probability via next-token prediction,

$$p(x_1, \dots, x_n) = p(x_1) \cdot \prod_{i=2}^n p(x_i | x_1, \dots, x_{i-1}),$$
 (3)

where the conditional probabilities $p(x_i|x_1, ..., x_{i-1})$ are modeled by a parameterized language model.

Recent works mainly leverage RAG further to improve language modeling capability in the pre-training stage. A branch of works [11, 101, 163, 170] modifies the architecture of generators by adding a new cross-attention module in each transformer block for introducing retrieval knowledge. The intuition of those works is that given the similar Prefixes and their next tokens (retrieving stage), the pre-trained model can calibrate the model's prediction using the cross-attention module to capture the pattern between the next token and prefix (model forwarding stage). Zhong et al. [191] propose to augment the language model with three types of retrieval memories/databases (local memory, long-term memory, and external memory) and optimize the next-token probability distribution with nearest neighbors retrieved from the memories/databases. Another branch of works [53, 68, 88, 133, 177] focuses on augmenting the inputs or outputs of generators with retrievals. Guu et al. [53] and Ram et al. [133] concatenate the retrieved knowledge with inputs and feed the retrieval-augmented inputs into the generators. Other works [68, 88, 177] fuse the logits of inputs as well as retrievals at the final output layer and generate the final probability distribution based on the interpolated results. Those works believe that the concatenated/fused retrievals can provide useful context information on inputs/outputs to improve models' robustness during the pre-training stage. Besides, Doostmohammadi et al. [33] focus on pre-training models with a semantic retriever (BM25) and achieve a better language modeling performance.

8.2 Machine Translation

Machine translation (MT) leverages computational linguistics algorithms to translate text or speech from one language to another automatically. The goal of MT is to produce an accurate and fluent translation, preserving the meaning of the original text while adhering to the grammatical and stylistic norms of the target language. MT systems have evolved from rule-based machine translation (RBMT) to statistical machine translation (SMT) and, more recently, to neural machine translation (NMT). In particular, NMT methods have significantly improved translation quality by leveraging deep learning techniques, which thus will be the focus of this section.

RAG techniques can further enhance MT by incorporating external knowledge into the translation process. The simplest way is to concatenate the similar translation examples into the inputs or fuse the logits of similar translation examples at the output layer. For example, some works [18, 161] retrieve similar translations according to the source text and concatenate corresponding target texts or pairs of source and target texts as examples into inputs. Other works [60, 87, 190] feed the retrieved source text into the models and obtain the logits of the next target tokens, then aggregate all logits to generate the final predictions. Moreover, Jiang et al. [78] and Li et al. [97] use the logits of retrieved examples to calibrate the aggregated logits, improving the robustness of the generation. Another branch of works [191, 192] injects external knowledge into the objective function during the training stage, refining the representation space with similar translations. Besides, Cai et al. [13] encode similar translations and store them as the translation memory, then introduce the knowledge from memory with a cross-attention module. Instead of improving the performance, a branch of work focuses on accelerating the generation efficiency on MT tasks, such as searching from a pre-built subset [29, 116] or a dynamic datastore [24], searching by chunks [113].

8.3 Text Summarization

Text summarization is the process of condensing a larger text document into a shorter version, preserving key information and the overall message. This task can be broadly categorized into two types: extractive summarization, which involves selecting and compiling parts of the original text, and abstractive summarization, which entails rewriting the essence of the text in a new, concise form. The goal is to produce a coherent and fluent summary that encapsulates the most critical information from the source material.

RAG techniques can significantly enhance text summarization tasks by leveraging external knowledge and similar documents to inform the summarization process. [18, 38, 98, 161] simply concatenates the retrieved similar summaries into inputs to generate summarizations. Instead of concatenating texts, other works fuse features at the intermediate layers by cross-attention [10], or at the output layers by logits ensemble [60]. Besides, Jiang et al. [80] argue that retrieving for every generation may not always be the best choice and propose to retrieve external knowledge during the generation process adaptively.

8.4 Question Answering

Question Answering (QA) is a fundamental task in NLP that involves building systems capable of automatically answering human questions in natural language. QA systems can be broadly classified into two categories: open-domain, where the system answers questions about virtually anything, and closed-domain, focusing on a specific area of knowledge. The primary challenge in QA is understanding the question's intent and retrieving accurate, relevant information from a vast collection of data to provide a concise answer. Due to the page limits, this paper only discusses the works of open-domain QA systems.

RAG techniques combine information retrieval with model-based generation, which is highly suitable for QA systems. In particular, open-domain QA systems usually first require searching for knowledge from the Internet or large-scale databases, then generate the corresponding answers according to the retrieved knowledge. Naturally, given similar questions and corresponding answers as demonstrations which are concatenated into inputs [64, 98, 161], generators in RAG can learn the pattern between questions and answers and infer what answers should be. For some specific QA tasks where a set of reference documents is given, retrievers in RAG would retrieve the relevant documents for concatenation, and then generators in RAG would read the context then generate the final answers via the self-attention mechanism [7, 53, 93, 133], which is similar to solving a reading comprehension problem. Besides, Fabbri et al. [36] focus on designing effective templates for re-organizing the concatenated contexts. Baek et al. [8] leverage the knowledge

graph to retrieve the related facts for the input questions, then feed their concatenation and inputs into the generators. Instead of directly concatenating texts, another branch of works focuses on joining the retrieval embeddings with input embeddings for the encoder-decoder models [27, 72, 73, 141].

Some works incorporate the external knowledge in the hidden states or the final logits of generators. For the fusion in the hidden states, the key is what kind of knowledge representation would be injected, such as entities [28, 40], chunks [11, 158], documents [17]. For the fusion in the logits, most works combine the logits of retrievals and inputs by ensemble techniques [53, 95, 119, 146].

Instead of designing different knowledge fusions for QA systems, existing works also improve QA systems with RAG from other aspects. Some works [48, 102, 125] use retrieved questionanswering pairs as extra training data. Some works optimize the retriever module, e.g., improving the keys' representation when building the retriever database [132], replacing the indexing with a pre-trained ranking model [181], or enabling retrieving phrases with two queries [118]. Other works focus on accelerating the generation efficiency of RAG. Jong et al. [26] propose the layer-sparse cross-attention to speed up the decoding. Some works [7, 80, 165] observe that the retrievals may not always provide useful information during the generation process and learn to determine when to retrieve. Moreover, Sun et al. [152] combine the RAG with agents to iteratively reason the final results.

8.5 Information Extraction

Information Extraction (IE) is a critical task in NLP to automatically extract structured information from unstructured and semistructured text sources. This task encompasses several sub-tasks, including Named Entity Recognition (NER), Entity Linking (EL), Coreference Resolution (CR), Relation Extraction (RE), etc. The goal is to identify and classify key elements from text and understand the relationships between them, thereby converting textual data into a structured format amenable to analysis and interpretation.

With RAG techniques, addressing IE tasks can be significantly improved in terms of not only performance but also interpretability. In NER tasks, Wang et al. [164] first retrieve similar sentences and then concatenate the ranked retrievals for better semantic representations. Ren et al. [138] show that naive RAG may not address Event Argument Extraction (EAE) tasks. Thus, they adopt a sampling-based method to guarantee the same distribution of event labels between retrievals and inputs then concatenate retrieval texts into inputs for better performance in EAE tasks. Table augmentation is also a challenging task, which requires extracting information from tables. Glass et al. [45] propose to extract information in a retrieval-augmented manner.

8.6 Text Classification

Text classification tasks are common in NLP applications. Sentiment analysis, a prominent text classification task in NLP, entails identifying and categorizing the emotional tone conveyed in a text. For example, given a sentence of "I love to watch movies", the analysis models should determine whether it has a positive attitude or a negative attitude. The attitude in sentiment analysis can range from positive to negative or can be neutral, nuanced, and even mixed. The sentiment analysis task is crucial for understanding consumer feedback, monitoring brand reputation, and gaining insights into public opinion on various issues.

RAG techniques can significantly enhance sentiment analysis with different external knowledge fusion strategies. Li et al. [98] concatenate the retrieved options and corresponding prompt-based labels with input options. Other works [16, 52] concatenate the retrieval embeddings with input embeddings before feeding them into the decoder. Some works fuse the retrieval features into the hidden states of generators via cross-attention [17, 163] or rankingbased addition [169]. Besides, other works focus on fusing the logits of retrievals with the output logit using ensemble techniques [179, 185]. Except for knowledge fusions, Min et al. [118] enable locating knowledge in phrases more accurately via two queries.

8.7 Dialogue Systems

Dialogue systems, also known as conversational agents or chatbots, are designed to simulate conversation with human users, either in text or speech form. These systems can be categorized into two main types: task-oriented systems [65], which assist users in completing specific tasks such as booking tickets or ordering food, and opendomain systems, which aim to carry on a general conversation on a wide range of topics [147]. The core challenge in developing effective dialogue systems lies in understanding user intent, maintaining context, and generating coherent, relevant responses.

Existing works improve the dialogue system with RAG mostly via the query-based fusions. Some works [18, 90, 100] concatenate the retrieved history conversations with current inputs. Other works [18, 39, 105] first leverage an encoder to encode the history responses, then feed the concatenated embeddings into a decoder to generate new responses.

9 APPLICATIONS

9.1 LLM-based Autonomous Agents

LLM-based autonomous agents are intelligent software systems that leverage the power of LLMs to perform tasks without the need for continuous human intervention [99, 160, 171]. These agents use LLMs as a brain or controller [67], and extend their abilities through multimodal perception [173], tool utilization [143] and external memory [124]. Especially, external long-term memory for agents functions as the knowledge datastore in RAG, which provides agents with the capability to incorporate external knowledge over extended periods. Therefore, applying RAG would be beneficial to access a broader range of information, improving agents' decision-making and problem-solving abilities [187]. This section explores how LLM-based agents can leverage RAG from two perspectives.

Using RAG to Retrieve from External Memory. LLM-based agents can utilize RAG to access and retrieve information from their own external memory [56, 114, 188]. This external memory serves as a knowledge base that the agent can draw upon to enhance its understanding and decision-making. When faced with a query or a task, the agent can use RAG to retrieve relevant information from this memory, which is then integrated into the generation process of the LLM. This allows the agent to produce responses or solutions that are informed by a wider range of knowledge, leading to more accurate and contextually relevant outcomes.

The ability to tap into a vast external memory enables the agent to continuously learn and adapt based on new information, making it more effective over time. Using Tools to Search the Web and RAG for Up-to-Date Information. In addition to retrieving information from its own memory, an LLM-based agent can use tools to search the web for the most current information [143]. This capability is particularly useful for tasks that require up-to-date knowledge, such as news summarization, market analysis, or responding to rapidly evolving situations. Once the agent retrieves the latest information from the web, it can use RAG to integrate this data into its generation process. By combining the LLM's natural language understanding with real-time data from the web, the agent can generate responses that are not only contextually relevant but also reflect the latest developments. This approach enhances the agent's ability to provide accurate and timely information, improving its effectiveness in dynamic environments.

In both cases, RAG plays a crucial role in augmenting the capabilities of LLM-based agents by enabling them to access and leverage a wider range of information, whether it's from their own external memory or from real-time sources on the web. This leads to more informed decision-making and enhances the overall performance of the agents.

9.2 Frameworks

Frameworks like Langchain [91] and LLaMAindex [103] pose a significant impact on enhancing the practical implementation of RAG. Langchain and LLaMAindex exemplify the integration of sophisticated retrieval mechanisms with generative models, facilitating the seamless incorporation of external data into the language generation process. This section will introduce these two representative RAG frameworks in details.

Langchain is a framework designed to augment the capabilities of language models by integrating them with external knowledge sources and databases. It acts as a middleware that facilitates the interaction between language models and various data retrieval systems, enabling a more informed and accurate generation of responses. The core functionality of Langchain involves orchestrating the flow of information from external databases into the generative process of language models, enhancing their ability to leverage context and specific knowledge in their responses. This integration plays a crucial role in enabling language models to perform tasks that require access to up-to-date or detailed information that is not contained within the model's initial training data.

LLaMAindex is a specialized data framework that focuses on organizing and indexing vast amounts of data to improve the retrieval capabilities of language models. This framework supports efficient querying mechanisms, allowing language models to quickly access relevant information from a structured repository. LLaMAindex is designed to be highly scalable and can handle diverse data types, from text documents to structured databases. The indexed data supports a wide range of applications, from simple fact retrieval to complex analytical tasks, making it an indispensable tool for enhancing the information retrieval phase in language models. Both Langchain and LLaMAindex are deeply connected to the concept of RAG. Langchain enhances RAG by providing a structured way for language models to interact with external databases and knowledge sources during the generation process. On the other hand, LLaMAindex serves as a powerful backend for RAG systems by ensuring that the retrieval process is both fast and relevant. Together, Langchain and LLaMAindex enhance the capabilities of RAG by ensuring that the language models are not only generating text based on their internal knowledge but are also capable of pulling in external data to provide responses that are contextually enriched and informationally robust.

10 DISCUSSION AND FUTURE DIRECTION

Despite the success of the RAG for natural language processing, some challenges should be considered. This paper highlights these challenges to inspire future research and provides possible future research directions in RAG for NLP.

10.1 Retrieval Quality

The retrieval quality largely determines the performance of the generator. Park et al. [126] investigated imperfect retrieval's effect and demonstrated that unanswerability or contradiction of a document set, which frequently leads to hallucinations. However, Ren et al. [137] demonstrated that the improvement brought about by the increase in the number of supporting documents is not entirely due to the increase in recall. To this end, using various evaluation methods referred to in Section 7 for retrieval information will be more beneficial to model generation. The retrieval quality involves the following four key factors that must be designed. The first consideration is **determining the optimal key** to use in the vector database. This process typically involves subjective decisionmaking and requires human effort to design effectively. The naive idea is to choose inputs for the given tasks, treating each task as a QA problem.

The second is **the choice of embedding model**. After determining the key, the next step is leveraging embedding models to convert text into vector representations. Models such as BERT [31], RoBERTa [107], or domain-specific embeddings can be crucial to determine how well nuances and contextual meanings are captured. Adapting the embedding model to suit specific data types or queries better can significantly enhance retrieval quality. This requires training the model on domain-specific corpora, including the types of queries and documents the system will encounter.

Thirdly, **designing effective similarity metrics** is also crucial to improve retrieval quality. The goal of similarity metrics is to measure the relevance between the query and the retrieved information. Some classical similarity metrics, such as cosine similarity or Euclidean distance, used for ranking in the recommender system can also be used in RAG [47]. Apart from these metrics, some works explored more complex similarity metrics, such as optimal transport distance [23], to obtain a task-specific similarity.

Finally, **approximate nearest neighbor (ANN) searching** is also a key step in determining what knowledge should be returned as nearest neighbors. Advanced ANN searching aims to accelerate the retrieval efficiency at the cost of sacrificing the retrieval quality. Choosing a suitable ANN algorithm, such as product quantization [74] or HNSW [111], requires a good trade-off between retrieval efficiency and retrieval quality. All of these factors collectively contribute to the retrieval quality of the retriever.

10.2 RAG Efficiency

RAG efficiency is crucial for downstream NLP applications, which limits the volume of data that can be retrieved. There are two simple ways to guarantee RAG efficiency without new algorithms, i.e., reducing the volume of data or adding more powerful computing and memory resources. However, the former may impact the retrieval quality, while the latter requires more resource cost.

RAG efficiency encompasses the efficiency of the retriever and the efficiency of retrieval fusions. Retriever efficiency refers to the time cost of retrieving relevant information, which can be divided into three parts, i.e., encoding time, ANN searching time, and data fetching time of the datastore. It is unnecessary to jointly optimize all three components as the bottleneck would vary from different database sizes. For smaller retrieval databases, such as those with fewer than 1 million entries, the encoding phase is often the primary bottleneck, as the vector database can be all stored in the memory. Several topics, such as model quantization [9, 89], distillation [32, 81], or model pruning [42], are used to accelerate the encoding.

In contrast, for larger databases, the time cost of searching in the index and fetching data from the datastore becomes the major bottleneck, as the searching is over a considerable amount of data, and the fetching involves I/O overheads. In this case, efficient ANN searching algorithms [34, 51, 84] and system-level optimizations [79, 82] are the main focus.

Retrieval fusion efficiency, which aims to enhance the inference efficiency when integrating retrievals, is worth to be optimized for improving the RAG efficiency. For example, the computational overhead of query-based fusion is often non-negligible due to the long sequence length. Some works, such as Fid-light [59] and Re-Fusion [169], mainly target reducing the computations while integrating the retrieved information.

10.3 Discussion for Query-based Fusion in RAG

Two main query-based fusion RAG concerns must be discussed. The first is the performance comparison towards long-context LLMs. Hui et al. [69] demonstrated that RAG system and long-context LLMs show similar performance for free-form or knowledge-based tasks (e.g., paper-based and wiki-based Q&A). However, in tasks requiring numerical reasoning (e.g., financial Q&A), long-context LLMs underperform compared to RAG, as the verbose content in long contexts can obscure key facts and hinder precise reasoning. Additionally, RAG tends to outperform long-context mechanisms in smaller LLMs, which struggle to process large volumes of data effectively. The second concern is improving the query-based fusion method, which is prevalent in modern RAG systems. Recent works used prompt engineering strategies [52, 95], query refinement techniques [178], calibration-based strategies [7, 80], and iterative RAG [145] to improve the query-based fusion RAG. Among these strategies, prompt engineering and calibration offer improved results with minimal overhead, while iterative retrieval-augmented

generation (RAG) and query refinement prioritize higher computational costs in exchange for potential accuracy gains. Prompt engineering and query refinement depend on manual tuning, whereas calibration and iterative RAG are more automation-friendly but require high-quality, robust data. Simpler strategies, such as prompting, are easier to deploy but lack flexibility, while more complex methods like iterative RAG provide greater adaptability at the expense of stability. These approaches can be combined (e.g., refined queries with iterative loops and calibrated prompts), though such synergies must carefully balance the compounded drawbacks, such as increased latency and complexity.

10.4 Choices of Fusions

This paper introduces three kinds of retrieval fusions, where each fusion is worth further exploring. Query-based fusions concatenate the texts or embeddings of retrieved knowledge with inputs. These methods have better interpretability and are easy to apply even only when the API of LLMs is provided. However, concatenation leads to a long sequence of inputs, thus resulting in a large computational overhead in the attention and truncation of inputs. Some works [6, 169] aim to improve efficiency when integrating retrievals, while others [10, 163] focus on improving the efficiency when increasing the model input length.

Conversely, latent-based fusions amalgamate information at a deeper, more abstract level, which may capture more nuanced relationships between the retrieved information and the query. However, these fusions significantly lack interpretability and often require pre-training or fine-tuning to adjust the retrieval embeddings or reweight the retrievals. Therefore, enhancing the interpretability of such latent-based fusions is also worth exploring in the future.

Logits-based fusions incorporate information at the decision level, thereby offering a potentially more flexible and robust integration of data from various sources. Nonetheless, these fusions may oversimplify the fusion process, diminishing the richness of the retrieved information by reducing them to logit values. Meanwhile, such fusions require performing all inference of retrievals, which is also a time-consuming process.

Apart from applying one kind of fusion in practical applications, combining different fusions is also worth exploring for better performance. These fusion methods are not mutually exclusive, as they focus on augmenting the different stages of generators, i.e., inputs, hidden states, and outputs. Besides, during the generation, when to fuse retrieved knowledge is also a significant problem worthy of further exploration [112].

10.5 RAG Training

As introduced in Section 6, RAG training includes two branches of works, RAG with/without datastore update. For RAG without a datastore update, the main challenge is how to jointly optimize all parameters in RAG. This may involve new loss functions with multiple objectives, new optimizations for efficient tuning parameters in the retriever and generator, or other training strategies.

For RAG with datastore update, one challenge is how to align the retrieval representations with the generator's representations. Although the time cost of the update operation in datastore cannot be ignored, some works [16] reduce the update frequency by asynchronously updating, thus achieving the alignment of knowledge representation and model's representation. Another challenge is when to retrain/fine-tune the generator in RAG when a new corpus is added. Due to the in-context learning capability of existing LLM-based generators and high training overhead, retraining/finetuning the generator or directly inferring the generator becomes a challenging choice for different scenarios. Recently, some efficient training strategies [30, 61] have been proposed to accelerate the fine-tuning process, which can be taken into consideration.

10.6 Cross-Modality Retrieval

Retrieving cross-modality information in NLP tasks can greatly enhance the quality and richness of the representations, leading to improved performance. First, cross-modality information, such as combining text with images, videos, or audio, provides a richer context to the content [62]. For instance, when language is ambiguous, accompanying images can clarify meanings difficult to convey through text alone. Second, different modalities can contribute various types of information that are not accessible from a single source. For example, visual data can provide spatial, color, and action cues, while textual data can offer detailed descriptions, emotions, or abstract concepts. Combining these can lead to a more comprehensive understanding of the data. Moreover, Models trained on multi-modal data typically exhibit increased robustness and generalizability [168]. These models are adept at associating information across diverse inputs, mitigating overfitting to the peculiarities of a single modality. This attribute is particularly valuable in real-world applications of NLP, such as in autonomous vehicles, where systems must interpret textual information from signs or dialogues and sensory data from the surrounding environment to make informed decisions. Furthermore, multi-modal data can resolve ambiguities that cannot be resolved within a single modality. For example, the phrase "bank" can refer to either a financial institution or the side of a river, and visual context can help disambiguate this. Last, human communication is inherently multi-modal, incorporating elements such as gestures, facial expressions, and tone of voice. Systems capable of processing multiple modes of communication can interact with humans in a manner that is both more natural and intuitive. In conclusion, integrating cross-modality information in RAG for NLP tasks not only enhances the richness and quality of data representations but also significantly improves the systems' comprehension, interaction capabilities, and adaptability to diverse applications.

11 CONCLUSION

In this survey, we delve into the development of RAG within the field of natural language processing. First, this paper introduces the components of RAG and their functionalities. Subsequently, this paper elaborates on each step involved in retriever, discussing the diverse techniques. Furthermore, this paper categorizes the retrieval fusions, evaluating the strengths and weaknesses inherent in each retrieval fusion technique. Besides, this paper discusses the RAG training, including RAG with/without datastore update. Then, this paper presents RAG evaluation and benchmarking, and explores how RAG can be adapted for various NLP tasks and provides practical applications of RAG in real-world scenarios. Conclusively, this paper identifies ongoing challenges and suggests directions for future research to foster advancements in this evolving area.

REFERENCES

- Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. 2022. Exploring the Limits of Large Scale Pre-training. In The Tenth International Conference on Learning Representations (ICLR).
- [2] Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. 2024. Evaluating Correctness and Faithfulness of Instruction-Following Models for Question Answering. *Trans. Assoc. Comput. Linguistics* 12 (2024), 681–699.
- [3] Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrón, and Sumit Sanghai. 2023. GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP), Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 4895–4901.
- [4] Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona T. Diab, and Marjan Ghazvininejad. 2022. A Review on Language Models as Knowledge Bases. *CoRR* abs/2204.06031 (2022).
- [5] Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy P. Lillicrap, An-geliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul Ronald Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, and et al. 2023. Gemini: A Family of Highly Capable Multimodal Models. *CoRR* abs/2312.11805 (2023).
- [6] Md. Adnan Arefeen, Biplob Debnath, and Srimat Chakradhar. 2023. LeanContext: Cost-Efficient Domain-Specific Question Answering Using LLMs. CoRR abs/2309.00841 (2023).
- [7] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.
- [8] Jinheon Baek, Alham Fikri Aji, and Amir Saffari. 2023. Knowledge-Augmented Language Model Prompting for Zero-Shot Knowledge Graph Question Answering. *CoRR* abs/2306.04136 (2023).
- [9] Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jin Jin, Xin Jiang, Qun Liu, Michael R. Lyu, and Irwin King. 2021. BinaryBERT: Pushing the Limit of BERT Quantization. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/IJCNLP). Association for Computational Linguistics, 4334–4348.
- [10] Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew R. Gormley. 2023. Unlimiformer: Long-Range Transformers with Unlimited Length Input. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [11] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre. 2022. Improving Language Models by Retrieving from Trillions of Tokens. In Proceedings of the 39th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research), Vol. 162. 2206–2240.
- [12] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In Advances in Neural Information Processing Systems 33 (NeurIPS).
- [13] Deng Cai, Yan Wang, Huayang Li, Wai Lam, and Lemao Liu. 2021. Neural Machine Translation with Monolingual Translation Memory. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/IJCNLP). 7307-7318.
- [14] Junying Chen, Qingcai Chen, Dongfang Li, and Yutao Huang. 2022. SeDR: Segment Representation Learning for Long Documents Dense Retrieval. CoRR abs/2211.10841 (2022).

- [15] Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. Benchmarking Large Language Models in Retrieval-Augmented Generation. In Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada, Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (Eds.). AAAI Press, 17754–17762.
- [16] Xiang Chen, Lei Li, Ningyu Zhang, Xiaozhuan Liang, Shumin Deng, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2022. Decoupling Knowledge from Memorization: Retrieval-augmented Prompt Learning. In Advances in Neural Information Processing Systems 35 (NeurIPS).
- [17] Xin Cheng, Yankai Lin, Xiuying Chen, Dongyan Zhao, and Rui Yan. 2023. Decouple knowledge from paramters for plug-and-play language modeling. In Findings of the Association for Computational Linguistics (ACL). 14288–14308.
- [18] Xin Cheng, Di Luo, Xiuying Chen, Lemao Liu, Dongyan Zhao, and Rui Yan. 2023. Lift Yourself Up: Retrieval-augmented Text Generation with Self-Memory. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [19] Alexis Chevalier, Alexander Wettig, Anirudh Ajith, and Danqi Chen. 2023. Adapting Language Models to Compress Contexts. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP). 3829– 3846.
- [20] Yung-Sung Chuang, Wei Fang, Shang-Wen Li, Wen-tau Yih, and James R. Glass. 2023. Expand, Rerank, and Retrieve: Query Reranking for Open-Domain Question Answering. In *Findings of the Association for Computational Linguistics* (ACL). Association for Computational Linguistics, 12131–12147.
- [21] Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In 8th International Conference on Learning Representations (ICLR). OpenReview.net.
- [22] Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, André F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and Michael Desa. 2024. SaulLM-7B: A pioneering Large Language Model for Law. CoRR abs/2403.03883 (2024).
- [23] Yufei Cui, Ziquan Liu, Yixin Chen, Yuchen Lu, Xinyue Yu, Xue (Steve) Liu, Tei-Wei Kuo, Miguel Rodrigues, Chun Jason Xue, and Antoni B. Chan. 2023. Retrieval-Augmented Multiple Instance Learning. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [24] Yuhan Dai, Zhirui Zhang, Qiuzhi Liu, Qu Cui, Weihua Li, Yichao Du, and Tong Xu. 2023. Simple and Scalable Nearest Neighbor Machine Translation. In The Eleventh International Conference on Learning Representations (ICLR).
- [25] David Dale, Elena Voita, Loïc Barrault, and Marta R. Costa-jussà. 2023. Detecting and Mitigating Hallucinations in Machine Translation: Model Internal Workings Alone Do Well, Sentence Similarity Even Better. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). Association for Computational Linguistics, 36–50.
- [26] Michiel de Jong, Yury Zemlyanskiy, Joshua Ainslie, Nicholas FitzGerald, Sumit Sanghai, Fei Sha, and William W. Cohen. 2023. FiDO: Fusion-in-Decoder optimized for stronger performance and faster inference. In *Findings of the Association for Computational Linguistics (ACL)*. 11534–11547.
- [27] Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Joshua Ainslie, Sumit Sanghai, Fei Sha, and William W. Cohen. 2023. Pre-computed memory or onthe-fly encoding? A hybrid approach to retrieval augmentation makes the most of your compute. In Proceedings of the 40th International Conference on Machine Learning (ICML). 7329–7342.
- [28] Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, and William W. Cohen. 2022. Mention Memory: incorporating textual knowledge into Transformers through entity mention attention. In *The Tenth International Conference* on Learning Representations (ICLR).
- [29] Hiroyuki Deguchi, Taro Watanabe, Yusuke Matsui, Masao Utiyama, Hideki Tanaka, and Eiichiro Sumita. 2023. Subset Retrieval Nearest Neighbor Machine Translation. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 174–189.
- [30] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [31] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT). 4171-4186.
- [32] Zixiang Ding, Guoqing Jiang, Shuai Zhang, Lin Guo, and Wei Lin. 2023. SKD-BERT: Compressing BERT via Stochastic Knowledge Distillation. In *Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 7414– 7422.
- [33] Ehsan Doostmohammadi, Tobias Norlund, Marco Kuhlmann, and Richard Johansson. 2023. Surface-Based Retrieval Reduces Perplexity of Retrieval-Augmented Language Models. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 521–529.

- [34] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The Faiss library. CoRR abs/2401.08281 (2024).
- [35] explosion. 2016. Spacy. https://spacy.io/
- [36] Alexander R. Fabbri, Patrick Ng, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. Template-Based Question Generation from Retrieved Sentences for Improved Unsupervised Question Answering. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL). 4508–4513.
- [37] Facebook. 2013. RocksDB. https://github.com/facebook/rocksdb
- [38] Angela Fan and Claire Gardent. 2022. Generating Full Length Wikipedia Biographies: The Impact of Gender Bias on the Retrieval-Based Generation of Women Biographies. CoRR abs/2204.05879 (2022).
- [39] Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2021. Augmenting Transformers with KNN-Based Composite Memory for Dialog. Trans. Assoc. Comput. Linguistics 9 (2021), 82–99.
- [40] Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. Entities as Experts: Sparse Memory Access with Entity Supervision. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). 4937–4951.
- [41] Robert Friel, Masha Belyi, and Atindriyo Sanyal. 2024. Ragbench: Explainable benchmark for retrieval-augmented generation systems. arXiv preprint arXiv:2407.11005 (2024).
- [42] Prakhar Ganesh, Yao Chen, Xin Lou, Mohammad Ali Khan, Yin Yang, Hassan Sajjad, Preslav Nakov, Deming Chen, and Marianne Winslett. 2021. Compressing Large-Scale Transformer-Based Models: A Case Study on BERT. Trans. Assoc. Comput. Linguistics 9 (2021), 1061–1080.
- [43] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, 6894–6910.
- [44] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. CoRR abs/2312.10997 (2023).
- [45] Michael R. Glass, Xueqing Wu, Ankita Rajaram Naik, Gaetano Rossiello, and Alfio Gliozzo. 2023. Retrieval-Based Transformer for Table Augmentation. In Findings of the Association for Computational Linguistics (ACL). 5635–5648.
- [46] Hongyu Gong, Yelong Shen, Dian Yu, Jianshu Chen, and Dong Yu. 2020. Recurrent Chunking Mechanisms for Long-Text Machine Reading Comprehension. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL). Association for Computational Linguistics, 6751–6761.
- [47] Asela Gunawardana and Guy Shani. 2009. A Survey of Accuracy Evaluation Metrics of Recommendation Tasks. J. Mach. Learn. Res. 10 (2009), 2935-2962.
- [48] Daya Guo, Duyu Tang, Nan Duan, Ming Zhou, and Jian Yin. 2019. Coupling Retrieval and Meta-Learning for Context-Dependent Semantic Parsing. In Proceedings of the 57th Conference of the Association for Computational Linguistics (ACL). 855–866.
- [49] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025).
- [50] Rentong Guo, Xiaofan Luan, Long Xiang, Xiao Yan, Xiaomeng Yi, Jigao Luo, Qianya Cheng, Weizhi Xu, Jiarui Luo, Frank Liu, Zhenshan Cao, Yanliang Qiao, Ting Wang, Bo Tang, and Charles Xie. 2022. Manu: A Cloud Native Vector Database Management System. Proc. VLDB Endow. 15, 12 (2022), 3548–3561.
- [51] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. In Proceedings of the 37th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research), Vol. 119. PMLR, 3887–3896.
- [52] Zhicheng Guo, Sijie Cheng, Yile Wang, Peng Li, and Yang Liu. 2023. Prompt-Guided Retrieval Augmentation for Non-Knowledge-Intensive Tasks. In Findings of the Association for Computational Linguistics (ACL). 10896–10912.
- [53] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Retrieval Augmented Language Model Pre-Training. In Proceedings of the 37th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research), Vol. 119. 3929–3938.
- [54] David Harris and Sarah Harris. 2010. Digital design and computer architecture. Morgan Kaufmann.
- Zellig S Harris. 1954. Distributional structure. Word 10, 2-3 (1954), 146–162.
 Kostas Hatalis, Despina Christou, Joshua Myers, Steven Jones, Keith Lambert, Adam Amos-Binks, Zohreh Dannenhauer, and Dustin Dannenhauer. 2023. Memory Matters: The Need to Improve Long-Term Memory in LLM-Agents. In Proceedings of the AAAI Symposium Series, Vol. 2. 277–280.
- [57] Qiyuan He, Yizhong Wang, and Wenya Wang. 2024. Can Language Models Act as Knowledge Bases at Scale? CoRR abs/2402.14273 (2024).
- [58] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes

Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. Training Compute-Optimal Large Language Models. *CoRR* abs/2203.15556 (2022).

- [59] Sebastian Hofstätter, Jiecao Chen, Karthik Raman, and Hamed Zamani. 2023. FiD-Light: Efficient and Effective Retrieval-Augmented Text Generation. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR). ACM, 1437–1447.
- [60] Nabil Hossain, Marjan Ghazvininejad, and Luke Zettlemoyer. 2020. Simple and Effective Retrieve-Edit-Rerank Text Generation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL). 2532– 2538.
- [61] 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 *The Tenth International Conference on Learning Representations (ICLR).*
- [62] Xuming Hu. 2023. Multimodal Named Entity Recognition and Relation Extraction with Retrieval-Augmented Strategy. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR). ACM, 3488.
- [63] Yucheng Hu and Yuxing Lu. 2024. RAG and RAU: A Survey on Retrieval-Augmented Language Model in Natural Language Processing. CoRR abs/2404.19543 (2024).
- [64] Jie Huang, Wei Ping, Peng Xu, Mohammad Shoeybi, Kevin Chen-Chuan Chang, and Bryan Catanzaro. 2023. RAVEN: In-Context Learning with Retrieval Augmented Encoder-Decoder Language Models. *CoRR* abs/2308.07922 (2023).
- [65] Qiushi Huang, Shuai Fu, Xubo Liu, Wenwu Wang, Tom Ko, Yu Zhang, and Lilian H. Y. Tang. 2023. Learning Retrieval Augmentation for Personalized Dialogue Generation. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, 2523–2540.
- [66] Qiang Huang and Anthony K. H. Tung. 2023. Lightweight-Yet-Efficient: Revitalizing Ball-Tree for Point-to-Hyperplane Nearest Neighbor Search. In 39th IEEE International Conference on Data Engineering (ICDE). IEEE, 436–449.
- [67] Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. 2024. Understanding the planning of LLM agents: A survey. *CoRR* abs/2402.02716 (2024).
- [68] Yangsibo Huang, Daogao Liu, Zexuan Zhong, Weijia Shi, and Yin Tat Lee. 2023. kNN-Adapter: Efficient Domain Adaptation for Black-Box Language Models. *CoRR* abs/2302.10879 (2023).
- [69] Yulong Hui, Yao Lu, and Huanchen Zhang. 2024. UDA: A Benchmark Suite for Retrieval Augmented Generation in Real-World Document Analysis. In Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024, Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (Eds.).
- [70] Shonosuke Ishiwatari, Jingtao Yao, Shujie Liu, Mu Li, Ming Zhou, Naoki Yoshinaga, Masaru Kitsuregawa, and Weijia Jia. 2017. Chunk-based Decoder for Neural Machine Translation. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL). Association for Computational Linguistics, 1901–1912.
- [71] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. *Trans. Mach. Learn. Res.* 2022 (2022).
- [72] Gautier Izacard and Edouard Grave. 2021. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL). 874–880.
- [73] Gautier Izacard, Patrick S. H. Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot Learning with Retrieval Augmented Language Models. J. Mach. Learn. Res. 24 (2023), 251:1–251:43.
- [74] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. 2011. Product Quantization for Nearest Neighbor Search. *IEEE Trans. Pattern Anal. Mach. Intell.* 33, 1 (2011), 117–128.
- [75] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. ACM Comput. Surv. 55, 12 (2023), 248:1–248:38.
- [76] Ziwei Ji, Zihan Liu, Nayeon Lee, Tiezheng Yu, Bryan Wilie, Min Zeng, and Pascale Fung. 2023. RHO: Reducing Hallucination in Open-domain Dialogues with Knowledge Grounding. In *Findings of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, 4504–4522.
- [77] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. CoRR abs/2310.06825 (2023).

- [78] Hui Jiang, Ziyao Lu, Fandong Meng, Chulun Zhou, Jie Zhou, Degen Huang, and Jinsong Su. 2022. Towards Robust k-Nearest-Neighbor Machine Translation. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP). 5468–5477.
- [79] Wenqi Jiang, Shuai Zhang, Boran Han, Jie Wang, Bernie Wang, and Tim Kraska. 2024. PipeRAG: Fast Retrieval-Augmented Generation via Algorithm-System Co-design. CoRR abs/2403.05676 (2024).
- [80] Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active Retrieval Augmented Generation. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP). 7969–7992.
- [81] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. TinyBERT: Distilling BERT for Natural Language Understanding. In Findings of the Association for Computational Linguistics (EMNLP) (Findings of ACL), Vol. EMNLP 2020. Association for Computational Linguistics, 4163–4174.
- [82] Chao Jin, Zili Zhang, Xuanlin Jiang, Fangyue Liu, Xin Liu, Xuanzhe Liu, and Xin Jin. 2024. RAGCache: Efficient Knowledge Caching for Retrieval-Augmented Generation. *CoRR* abs/2404.12457 (2024).
- [83] Qiao Jin, Won Kim, Qingyu Chen, Donald C. Comeau, Lana Yeganova, W. John Wilbur, and Zhiyong Lu. 2023. MedCPT: Contrastive Pre-trained Transformers with large-scale PubMed search logs for zero-shot biomedical information retrieval. *Bioinform*. 39, 10 (2023).
- [84] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-Scale Similarity Search with GPUs. IEEE Trans. Big Data 7, 3 (2021), 535–547.
- [85] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling Laws for Neural Language Models. *CoRR* abs/2001.08361 (2020).
- [86] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). 6769–6781.
- [87] Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Nearest Neighbor Machine Translation. CoRR abs/2010.00710 (2020).
- [88] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Generalization through Memorization: Nearest Neighbor Language Models. In The 8th International Conference on Learning Representations (ICLR).
- [89] Sehoon Kim, Amir Gholami, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. 2021. I-BERT: Integer-only BERT Quantization. In Proceedings of the 38th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research), Vol. 139. PMLR, 5506–5518.
- [90] Brendan King and Jeffrey Flanigan. 2023. Diverse Retrieval-Augmented In-Context Learning for Dialogue State Tracking. In Findings of the Association for Computational Linguistics (ACL). 5570–5585.
- [91] LangChain. 2023. LangChain. https://www.langchain.com/
- [92] Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. 2022. Internet-augmented language models through few-shot prompting for open-domain question answering. *CoRR* abs/2203.05115 (2022).
- [93] Kyungjae Lee, Sang-eun Han, Seung-won Hwang, and Moontae Lee. 2023. When to Read Documents or QA History: On Unified and Selective Open-domain QA. In Findings of the Association for Computational Linguistics (ACL). 6420–6432.
- [94] Patrick S. H. Lewis, Barlas Oguz, Wenhan Xiong, Fabio Petroni, Scott Yih, and Sebastian Riedel. 2022. Boosted Dense Retriever. In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022, Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz (Eds.). Association for Computational Linguistics, 3102-3117.
- [95] Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In Advances in Neural Information Processing Systems 33 (NeurIPS).
- [96] Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemao Liu. 2022. A Survey on Retrieval-Augmented Text Generation. CoRR abs/2202.01110 (2022).
- [97] Xuanhong Li, Peng Li, and Po Hu. 2023. Revisiting Source Context in Nearest Neighbor Machine Translation. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP).
- [98] Xiaonan Li, Kai Lv, Hang Yan, Tianyang Lin, Wei Zhu, Yuan Ni, Guotong Xie, Xiaoling Wang, and Xipeng Qiu. 2023. Unified Demonstration Retriever for In-Context Learning. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 4644–4668.
- [99] Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, Rui Kong, Yile Wang, Hanfei Geng, Jian Luan, Xuefeng Jin, Zilong Ye, Guanjing Xiong, Fan Zhang, Xiang Li, Mengwei Xu, Zhijun Li, Peng Li, Yang Liu, Ya-Qin Zhang, and Yunxin Liu. 2024. Personal LLM Agents: Insights and Survey about the Capability, Efficiency and Partice Proceedings.

Security. CoRR abs/2401.05459 (2024).

- [100] Yunhao Li, Yunyi Yang, Xiaojun Quan, and Jianxing Yu. 2021. Retrieve & Memorize: Dialog Policy Learning with Multi-Action Memory. In Findings of the Association for Computational Linguistics (ACL/IJCNLP). 447–459.
- [101] Zonglin Li, Ruiqi Guo, and Sanjiv Kumar. 2022. Decoupled Context Processing for Context Augmented Language Modeling. In Advances in Neural Information Processing Systems 35 (NeurIPS).
- [102] Bill Yuchen Lin, Kangmin Tan, Chris Miller, Beiwen Tian, and Xiang Ren. 2022. Unsupervised Cross-Task Generalization via Retrieval Augmentation. In Advances in Neural Information Processing Systems 35 (NeurIPS).
- [103] Jerry Liu. 2022. LlamaIndex. https://doi.org/10.5281/zenodo.1234
- [104] Junyi Liu, Liangzhi Li, Tong Xiang, Bowen Wang, and Yiming Qian. 2023. TCRA-LLM: Token Compression Retrieval Augmented Large Language Model for Inference Cost Reduction. In Findings of the Association for Computational Linguistics (EMNLP). 9796–9810.
- [105] Shuai Liu, Hyundong Cho, Marjorie Freedman, Xuezhe Ma, and Jonathan May. 2023. RECAP: Retrieval-Enhanced Context-Aware Prefix Encoder for Personalized Dialogue Response Generation. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 8404–8419.
- [106] Shi-guang Liu and Yin-wei Wei. 2015. Fast nearest neighbor searching based on improved VP-tree. Pattern Recognit. Lett. 60-61 (2015), 8-15.
- [107] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. CoRR abs/1907.11692 (2019).
- [108] LMDB. 2014. LMDB. https://github.com/LMDB/lmdb
- [109] Man Luo, Shashank Jain, Anchit Gupta, Arash Einolghozati, Barlas Oguz, Debojeet Chatterjee, Xilun Chen, Chitta Baral, and Peyman Heidari. 2023. A Study on the Efficiency and Generalization of Light Hybrid Retrievers. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2023, Toronto, Canada, July 9-14, 2023, Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics. 1617–1626.
- [110] Xiaozhong Lyu, Stefan Grafberger, Samantha Biegel, Shaopeng Wei, Meng Cao, Sebastian Schelter, and Ce Zhang. 2023. Improving Retrieval-Augmented Large Language Models via Data Importance Learning. CoRR abs/2307 03027 (2023)
- Language Models via Data Importance Learning. CoRR abs/2307.03027 (2023).
 [111] Yury A. Malkov and Dmitry A. Yashunin. 2020. Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs. IEEE Trans. Pattern Anal. Mach. Intell. 42, 4 (2020), 824–836.
- [112] Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 9802–9822.
- [113] Pedro Henrique Martins, Zita Marinho, and André F. T. Martins. 2022. Chunkbased Nearest Neighbor Machine Translation. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP). 4228–4245.
- [114] Kai Mei, Zelong Li, Shuyuan Xu, Ruosong Ye, Yingqiang Ge, and Yongfeng Zhang. 2024. AIOS: LLM Agent Operating System. CoRR abs/2403.16971 (2024).
- [115] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and Editing Factual Associations in GPT. In Advances in Neural Information Processing Systems 35 (NeurIPS).
- [116] Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang, and Jiwei Li. 2022. Fast Nearest Neighbor Machine Translation. In Findings of the Association for Computational Linguistics (ACL). 555–565.
- [117] Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Cristian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, and et al. 2024. Gemma: Open Models Based on Gemini Research and Technology. CoRR abs/2403.08295 (2024).
- [118] Sewon Min, Weijia Shi, Mike Lewis, Xilun Chen, Wen-tau Yih, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2023. Nonparametric Masked Language Modeling. In Findings of the Association for Computational Linguistics (ACL). 2097–2118.
- [119] Aaron Mueller, Kanika Narang, Lambert Mathias, Qifan Wang, and Hamed Firooz. 2023. Meta-training with Demonstration Retrieval for Efficient Fewshot Learning. In *Findings of the Association for Computational Linguistics (ACL)*. 6049–6064.
- [120] Ewa Muszynska. 2016. Graph- and surface-level sentence chunking. In Proceedings of the ACL 2016 Student Research Workshop. Association for Computational Linguistics, 93–99.

- [121] NLTK. 2001. NLTK. https://www.nltk.org/
- [122] OpenAI. 2022. Text-Emb-Ada. https://platform.openai.com/docs/guides/ embeddings
- [123] OpenAI. 2023. GPT-4 Technical Report. CoRR abs/2303.08774 (2023).
- [124] Charles Packer, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph E. Gonzalez. 2023. MemGPT: Towards LLMs as Operating Systems. *CoRR* abs/2310.08560 (2023).
- [125] Bhargavi Paranjape, Matthew Lamm, and Ian Tenney. 2022. Retrieval-guided Counterfactual Generation for QA. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL). 1670–1686.
- [126] Seong-Il Park and Jay-Yoon Lee. 2024. Toward Robust RALMs: Revealing the Impact of Imperfect Retrieval on Retrieval-Augmented Language Models. Trans. Assoc. Comput. Linguistics 12 (2024), 1686–1702.
- [127] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language Models as Knowledge Bases?. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IfCNLP), Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, 2463– 2473.
- [128] Nicholas Pipitone and Ghita Houir Alami. 2024. LegalBench-RAG: A Benchmark for Retrieval-Augmented Generation in the Legal Domain. *CoRR* abs/2408.10343 (2024).
- [129] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *OpenAI blog* (2018).
- [130] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI* blog 1, 8 (2019), 9.
- [131] Anand Rajaraman and Jeffrey David Ullman. 2011. Data Mining. Cambridge University Press, 1–17.
- [132] Ori Ram, Liat Bezalel, Adi Zicher, Yonatan Belinkov, Jonathan Berant, and Amir Globerson. 2023. What Are You Token About? Dense Retrieval as Distributions Over the Vocabulary. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 2481–2498.
- [133] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-Context Retrieval-Augmented Language Models. Trans. Assoc. Comput. Linguistics 11 (2023), 1316–1331.
- [134] Parikshit Ram and Kaushik Sinha. 2019. Revisiting kd-tree for Nearest Neighbor Search. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD), Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 1378-1388.
- [135] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, and et al. 2024. Gemin 1.5: Unlocking multimodal understanding across millions of tokens of context. *CoRR* abs/2403.05530 (2024).
- [136] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). 3980–3990.
- [137] Ruiyang Ren, Yuhao Wang, Yingqi Qu, Wayne Xin Zhao, Jing Liu, Hua Wu, Ji-Rong Wen, and Haifeng Wang. 2025. Investigating the Factual Knowledge Boundary of Large Language Models with Retrieval Augmentation. In Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025, Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (Eds.). Association for Computational Linguistics, 3697–3715.
- [138] Yubing Ren, Yanan Cao, Ping Guo, Fang Fang, Wei Ma, and Zheng Lin. 2023. Retrieve-and-Sample: Document-level Event Argument Extraction via Hybrid Retrieval Augmentation. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 293–306.
- [139] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. Found. Trends Inf. Retr. 3, 4 (2009), 333–389.
- [140] Jon Saad-Falcon, Omar Khattab, Christopher Potts, and Matei Zaharia. 2024. ARES: An Automated Evaluation Framework for Retrieval-Augmented Generation Systems. In Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21,

2024, Kevin Duh, Helena Gómez-Adorno, and Steven Bethard (Eds.). Association for Computational Linguistics, 338–354.

- [141] Devendra Singh Sachan, Siva Reddy, William L. Hamilton, Chris Dyer, and Dani Yogatama. 2021. End-to-End Training of Multi-Document Reader and Retriever for Open-Domain Question Answering. In Advances in Neural Information Processing Systems 34 (NeurIPS). 25968–25981.
- [142] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distil-BERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR* abs/1910.01108 (2019).
- [143] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [144] Darius Koenig Julius Lipp Sean Lee, Aamir Shakir. 2024. Open Source Strikes Bread - New Fluffy Embeddings Model. https://www.mixedbread.ai/blog/mxbaiembed-large-v1
- [145] Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Enhancing Retrieval-Augmented Large Language Models with Iterative Retrieval-Generation Synergy. In Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 9248–9274.
- [146] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. REPLUG: Retrieval-Augmented Black-Box Language Models. CoRR abs/2301.12652 (2023).
- [147] Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval Augmentation Reduces Hallucination in Conversation. In *Findings of the Association for Computational Linguistics (EMNLP)*. Association for Computational Linguistics, 3784–3803.
- [148] Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Kumar Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Nathaneal Schärli, Aakanksha Chowdhery, Philip Andrew Mansfield, Blaise Agüera y Arcas, Dale R. Webster, Gregory S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle K. Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023. Large Language Models Encode Clinical Knowledge. Nature 620, 7972 (2023), 172–180.
- [149] Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaekermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Andrew Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle K. Barral, Dale R. Webster, Gregory S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2023. Towards Expert-Level Medical Question Answering with Large Language Models. CoRR abs/2305.09617 (2023).
- [150] Spotify. 2017. Annoy. https://github.com/spotify/annoy
- [151] Jianlin Su, Murtadha H. M. Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. RoFormer: Enhanced transformer with Rotary Position Embedding. *Neurocomputing* 568 (2024), 127063.
- [152] Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model with Knowledge Graph. CoRR abs/2307.07697 (2023).
- [153] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. CoRR abs/2302.13971 (2023).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, [154] Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Mova Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenvin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goval, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. CoRR abs/2307.09288 (2023).
- [155] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In Advances in Neural Information Processing Systems 30 (NeurIPS). 5998–6008.

- [156] Tu Vu, Mohit Iyyer, Xuezhi Wang, Noah Constant, Jerry W. Wei, Jason Wei, Chris Tar, Yun-Hsuan Sung, Denny Zhou, Quoc V. Le, and Thang Luong. 2023. FreshLLMs: Refreshing Large Language Models with Search Engine Augmentation. *CoRR* abs/2310.03214 (2023).
- [157] Boxin Wang, Wei Ping, Lawrence McAfee, Peng Xu, Bo Li, Mohammad Shoeybi, and Bryan Catanzaro. 2023. InstructRetro: Instruction Tuning post Retrieval-Augmented Pretraining. *CoRR* abs/2310.07713 (2023).
- [158] Boxin Wang, Wei Ping, Peng Xu, Lawrence McAfee, Zihan Liu, Mohammad Shoeybi, Yi Dong, Oleksii Kuchaiev, Bo Li, Chaowei Xiao, Anima Anandkumar, and Bryan Catanzaro. 2023. Shall We Pretrain Autoregressive Language Models with Retrieval? A Comprehensive Study. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP). 7763–7786.
- [159] Jianguo Wang, Xiaomeng Yi, Rentong Guo, Hai Jin, Peng Xu, Shengjun Li, Xiangyu Wang, Xiangzhou Guo, Chengming Li, Xiaohai Xu, Kun Yu, Yuxing Yuan, Yinghao Zou, Jiquan Long, Yudong Cai, Zhenxiang Li, Zhifeng Zhang, Yihua Mo, Jun Gu, Ruiyi Jiang, Yi Wei, and Charles Xie. 2021. Milvus: A Purpose-Built Vector Data Management System. In SIGMOD '21: International Conference on Management of Data. ACM, 2614–2627.
- [160] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024. A survey on large language model based autonomous agents. *Frontiers Comput. Sci.* 18, 6 (2024), 186345.
- [161] Shuohang Wang, Yichong Xu, Yuwei Fang, Yang Liu, Siqi Sun, Ruochen Xu, Chenguang Zhu, and Michael Zeng. 2022. Training Data is More Valuable than You Think: A Simple and Effective Method by Retrieving from Training Data. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL). 3170–3179.
- [162] Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. 2023. Knowledge Editing for Large Language Models: A Survey. CoRR abs/2310.16218 (2023).
- [163] Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. 2023. Augmenting Language Models with Long-Term Memory. In Advances in Neural Information Processing Systems 36 (NeurIPS).
- [164] Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. Improving Named Entity Recognition by External Context Retrieving and Cooperative Learning. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/IJCNLP). 1800–1812.
- [165] Yile Wang, Peng Li, Maosong Sun, and Yang Liu. 2023. Self-Knowledge Guided Retrieval Augmentation for Large Language Models. In Findings of the Association for Computational Linguistics (EMNLP). 10303–10315.
- [166] Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md. Rizwan Parvez, and Graham Neubig. 2023. Learning to Filter Context for Retrieval-Augmented Generation. *CoRR* abs/2311.08377 (2023).
- [167] Zichao Wang, Weili Nie, Zhuoran Qiao, Chaowei Xiao, Richard G. Baraniuk, and Anima Anandkumar. 2023. Retrieval-based Controllable Molecule Generation. In The Eleventh International Conference on Learning Representations, (ICLR).
- [168] Mike Wu and Noah D. Goodman. 2018. Multimodal Generative Models for Scalable Weakly-Supervised Learning. In Advances in Neural Information Processing Systems 31 (NeurIPS). 5580–5590.
- [169] Shangyu Wu, Ying Xiong, Yufei Cui, Xue Liu, Buzhou Tang, Tei-Wei Kuo, and Chun Jason Xue. 2024. ReFusion: Improving Natural Language Understanding with Computation-Efficient Retrieval Representation Fusion. CoRR abs/2401.02993 (2024).
- [170] Yuhuai Wu, Markus Norman Rabe, DeLesley Hutchins, and Christian Szegedy. 2022. Memorizing Transformers. In *The Tenth International Conference on Learning Representations (ICLR).*
- [171] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huang, and Tao Gui. 2023. The Rise and Potential of Large Language Model Based Agents: A Survey. CoRR abs/2309.07864 (2023).
- [172] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighof. 2023. C-Pack: Packaged Resources To Advance General Chinese Embedding. CoRR abs/2309.07597 (2023).
- [173] Junlin Xie, Zhihong Chen, Ruifei Zhang, Xiang Wan, and Guanbin Li. 2024. Large Multimodal Agents: A Survey. CoRR abs/2402.15116 (2024).
- [174] Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. Benchmarking Retrieval-Augmented Generation for Medicine. In *Findings of the Association*

for Computational Linguistics: ACL 2024, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.).

- [175] Ying Xiong, Xin Yang, Linjing Liu, Ka-Chun Wong, Qingcai Chen, Yang Xiang, and Buzhou Tang. 2023. EARA: Improving Biomedical Semantic Textual Similarity with Entity-Aligned Attention and Retrieval Augmentation. In Findings of the Association for Computational Linguistics (EMNLP). Association for Computational Linguistics, 8760–8771.
- [176] Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. RECOMP: Improving Retrieval-Augmented LMs with Compression and Selective Augmentation. *CoRR* abs/2310.04408 (2023).
- [177] Frank F. Xu, Uri Alon, and Graham Neubig. 2023. Why do Nearest Neighbor Language Models Work?. In Proceedings of the 40th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research), Vol. 202. 38325–38341.
- [178] Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective Retrieval Augmented Generation. CoRR abs/2401.15884 (2024).
- [179] Guoxin Yu, Lemao Liu, Haiyun Jiang, Shuming Shi, and Xiang Ao. 2023. Retrieval-Augmented Few-shot Text Classification. In Findings of the Association for Computational Linguistics (EMNLP). 6721–6735.
- [180] Hao Yu, Aoran Gan, Kai Zhang, Shiwei Tong, Qi Liu, and Zhaofeng Liu. 2024. Evaluation of Retrieval-Augmented Generation: A Survey. *CoRR* abs/2405.07437 (2024).
- [181] Zichun Yu, Chenyan Xiong, Shi Yu, and Zhiyuan Liu. 2023. Augmentation-Adapted Retriever Improves Generalization of Language Models as Generic Plug-In. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 2421–2436.
- [182] Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, Weng Lam Tam, Zixuan Ma, Yufei Xue, Jidong Zhai, Wenguang Chen, Zhiyuan Liu, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023. GLM-130B: An Open Bilingual Pre-trained Model. In *The Eleventh International Conference on Learning Representations (ICLR)*. OpenReview.net.
- [183] Biao Zhang and Rico Sennrich. 2019. Root Mean Square Layer Normalization. In Advances in Neural Information Processing Systems 32 (NeurIPS). 12360–12371.
- [184] Hongbo Zhang, Junying Chen, Feng Jiang, Fei Yu, Zhihong Chen, Guiming Chen, Jianquan Li, Xiangbo Wu, Zhiyi Zhang, Qingying Xiao, Xiang Wan, Benyou Wang, and Haizhou Li. 2023. HuatuoGPT, Towards Taming Language Model to Be a Doctor. In Findings of the Association for Computational Linguistics (EMNLP). Association for Computational Linguistics, 10859–10885.
- [185] Jianyi Zhang, Aashiq Muhamed, Aditya Anantharaman, Guoyin Wang, Changyou Chen, Kai Zhong, Qingjun Cui, Yi Xu, Belinda Zeng, Trishul Chilimbi, and Yiran Chen. 2023. ReAugKD: Retrieval-Augmented Knowledge Distillation For Pre-trained Language Models. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 1128–1136.
- [186] Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi, Shengyu Mao, Jintian Zhang, Yuansheng Ni, Siyuan Cheng, Ziwen Xu, Xin Xu, Jia-Chen Gu, Yong Jiang, Pengjun Xie, Fei Huang, Lei Liang, Zhiqiang Zhang, Xiaowei Zhu, Jun Zhou, and Huajun Chen. 2024. A Comprehensive Study of Knowledge Editing for Large Language Models. CoRR abs/2401.01286 (2024).
- [187] Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Adrian de Wynter, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. 2024. LLM as a Mastermind: A Survey of Strategic Reasoning with Large Language Models. *CoRR* abs/2404.01230 (2024).
- [188] Zeyu Zhang, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Quanyu Dai, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2024. A Survey on the Memory Mechanism of Large Language Model based Agents. *CoRR* abs/2404.13501 (2024).
- [189] Penghao Zhao, Hailin Zhang, Qinhan Yu, Zhengren Wang, Yunteng Geng, Fangcheng Fu, Ling Yang, Wentao Zhang, and Bin Cui. 2024. Retrieval-Augmented Generation for AI-Generated Content: A Survey. CoRR abs/2402.19473 (2024).
- [190] Xin Zheng, Zhirui Zhang, Junliang Guo, Shujian Huang, Boxing Chen, Weihua Luo, and Jiajun Chen. 2021. Adaptive Nearest Neighbor Machine Translation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/I/CNLP). 368–374.
- [191] Zexuan Zhong, Tao Lei, and Danqi Chen. 2022. Training Language Models with Memory Augmentation. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP). 5657–5673.
- [192] Wenhao Zhu, Jingjing Xu, Shujian Huang, Lingpeng Kong, and Jiajun Chen. 2023. INK: Injecting kNN Knowledge in Nearest Neighbor Machine Translation. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL). 15948–15959.