

Retrieving Multimodal Information for Augmented Generation: A Survey

Ruochen Zhao¹ Hailin Chen¹ Weishi Wang¹ Fangkai Jiao¹
Xuan Long Do^{1*} Chengwei Qin¹ Bosheng Ding¹ Xiaobao Guo¹
Minzhi Li² Xingxuan Li¹ Shafiq Joty^{1,3†}

¹Nanyang Technological University, Singapore

²National University of Singapore, Singapore

³Salesforce Research

{ruochen002, hailin001, xuanlong001, weishi001, fangkai002, chengwei003, bosheng001}@e.ntu.edu.sg

{xiaobao001, xingxuan001}@e.ntu.edu.sg, li.minzhi@u.nus.edu, srjoty@ntu.edu.sg

Abstract

As Large Language Models (LLMs) become popular, there emerged an important trend of using multimodality to augment the LLMs’ generation ability, which enables LLMs to better interact with the world. However, there lacks a unified perception of at which stage and how to incorporate different modalities. In this survey, we review methods that assist and augment generative models by retrieving multimodal knowledge, whose formats range from images, codes, tables, graphs, to audio. Such methods offer a promising solution to important concerns such as factuality, reasoning, interpretability, and robustness. By providing an in-depth review, this survey is expected to provide scholars with a deeper understanding of the methods’ applications and encourage them to adapt existing techniques to the fast-growing field of LLMs.

1 Introduction

Generative Artificial Intelligence (GAI) has demonstrated impressive performances in tasks such as text generation (Ouyang et al., 2022; Chowdhery et al., 2022; Brown et al., 2020) and text-to-image generation (Ramesh et al., 2021a; Poole et al., 2022). The recent advancements in Multimodal Large Language Models (MLLMs) (Driess et al., 2023; OpenAI, 2023; Huang et al., 2023b) have further improved the models’ capabilities to handle multi-format information, opening up possibilities for developing general-purpose learners.

Nevertheless, generative models are not exempt from inherent limitations, including the tendency for generating hallucinations (Ye and Durrett, 2022), struggling with arithmetic tasks (Patel et al., 2021), and lacking interpretability. Consequently, a promising solution for enhancing their capabilities lies in enabling them to interact with the external world and acquire knowledge in diverse formats and modalities, thereby improving

the factuality and rationality of the generated content. Recently, there have been emerging studies focusing on retrieval-augmented approaches (Mialon et al., 2023), which aim to provide information that is more grounded and factually dependent. Among them, most (Nakano et al., 2021; Guu et al., 2020b) retrieves textual information, which matches the data format used during pre-training and offers a natural medium for interaction. However, there is more world knowledge stored in different structures and modalities, such as images and videos, which is often inaccessible, unavailable, or not describable in traditional textual corpora.

Therefore, there arises an important research intersection that retrieves multimodal knowledge to augment generative models. It offers a promising solution to current challenges such as factuality, reasoning, interpretability, and robustness. As this field is very recent, there lacks a unified understanding of recognizing these methods as a specific group, visualizing their innate connections, connecting their methodologies, and outlining their applications.

Therefore, we survey recent advancements in multimodal retrieval-augmented generation (RAG). Specifically, we discuss current research by grouping them into different modalities, including image, code, structured knowledge, audio, and video. For each modality, we systematically search the ACL Anthology and Google Scholar with relevant keywords and perform manual filtering to determine their relevance to the survey. As a result, we collect 146 papers for detailed analysis. In Appendix A.1, we include search details, statistics, and a trend analysis figure, which shows that multimodal RAG papers have indeed developed very fastly since the emergence of large-scale general-purpose models. Within each modality, we discuss relevant papers by grouping them under different applications. By providing an in-depth survey, we hope to help researchers recognize the importance

*Now affiliated with NUS

† Work done while the author is on leave from NTU

of incorporating knowledge in different formats and encourage adaptation and advancements on existing techniques to the fast-growing field of LLMs.

In summary, our contributions are as follows:

- We establish retrieval augmented generation with multi-modality as an important group of methods that emerges with the recent advances in LLMs.
- For common modalities, we provide an in-depth review of research papers by contextualizing their innate connections and shared challenges.
- We provide an informative analysis of future directions, which could contain promising solutions to many current challenges.

2 Definitions and Background

To better understand the state and advancements that inspired multimodal retrieval augmentation, we first define and discuss the background of two key concepts: multimodal learning and retrieval-augmented generation (RAG).

2.1 Multimodal Learning

Multimodal learning refers to learning a unified representation of data from different modalities. It aims at extracting complementary information to facilitate compositional tasks (Baltrušaitis et al., 2018; Gao et al., 2020). In this survey, we include all modalities whose formats are different from natural language, including image, code, structured knowledge (*e.g.* tables, knowledge graphs), audio, and video.

Multimodal generative models have a wide range of applications, such as text-image generation, creative writing generation, and multilingual translation. For instance, the image recognition task can benefit from analyzing images and videos in conjunction with textual descriptions (Ju et al., 2022; Alayrac et al., 2022a; Jia et al., 2021; Radford et al., 2021b). Conversely, incorporating visual information also aids language understanding and generation (Zhou et al., 2020; Lei et al., 2021; ?). Moreover, they have the potential to significantly improve machine learning systems across various domains by enabling models to learn from and integrate multiple sources of information (Tsai et al., 2019; Acosta et al., 2022; Nagrani et al., 2021). Additionally, there has been growing interest in developing generative models that can output multiple modalities of data (Ramesh et al., 2021b; Crowson et al., 2022; Lin and Byrne, 2022a; Chen et al., 2022a). However, there remain challenges for multimodal generative models, such as gaining access

to a large amount of multimodal data and designing a network that produces semantically meaningful outputs.

2.2 Retrieval-Augmented Generation (RAG)

RAG typically consists of two phases: retrieving contextually relevant information, and guiding the generation process using the retrieved knowledge.

Recently, RAG has gained popularity in Natural Language Processing (NLP) due to the rise of general-purpose Large Language Models (LLMs) (Chowdhery et al., 2022; OpenAI, 2023), which have boosted performances in a wide range of NLP tasks. However, there are two primary challenges: Firstly, because generative models rely on the inner knowledge (weights), they result in a high amount of hallucinations (Zhao et al., 2023b). Secondly, due to their large parameter sizes and the high costs of updating, the traditional pretraining and finetuning approaches have become infeasible. As a solution, RAG methods (Gu et al., 2018; Weston et al., 2018; Cai et al., 2019b; Lewis et al., 2020) offer a promising solution for LLMs to effectively interact with the external world.

RAG is applied to a wide range of downstream NLP tasks, including machine translation (Gu et al., 2018; Zhang et al., 2018; Xu et al., 2020; He et al., 2021), dialogue generation (Weston et al., 2018; Wu et al., 2019; Cai et al., 2019a), abstractive summarization (Peng et al., 2019), and knowledge-intensive generation (Lewis et al., 2020; Izacard and Grave, 2021). Among them, most methods focus on retrieving textual information. For example, Guu et al. (2020b); Lewis et al. (2020); Borgeaud et al. (2022); Izacard et al. (2022) jointly train a retrieval system with an encoder or sequence-to-sequence LM, achieving comparable performance to larger LMs that use significantly more parameters. Recent research also proposes combining a retriever with chain-of-thought (CoT) prompting for reasoning to augment language models (He et al., 2022a; Trivedi et al., 2022; Zhao et al., 2023c).

3 Multimodal Retrieval-Augmented Generation

For each modality, there are different retrieval and synthesis procedures, targeted tasks, and challenges. Therefore, we discuss relevant methods by grouping them in terms of modality, including image, code, structured knowledge, audio, and video.

3.1 Image

Recent advances on pretrained models shed light on general image-text multi-modal models (Ramesh et al., 2021a; Alayrac et al., 2022b; Aghajanyan et al., 2022; Yu et al., 2022; Dou et al., 2022; Li et al., 2023a). However, these models require huge computational resources for pretraining and large amounts of model parameters — as they need to memorize vast world knowledge. More critically, they cannot efficiently deal with new or out-of-domain knowledge. To this end, multiple retrieval-augmented methods have been proposed to better incorporate external knowledge from images and text documents. In general text generation tasks, image retrieval can also improve generation quality by expanding text generation contexts with more “imagination”.

Visual question answering (VQA) To tackle open-domain VQA, RA-VQA (Lin and Byrne, 2022b) jointly trains the document retriever and answer generation module by approximately marginalizing predictions over retrieved documents. It first uses existing tools of object detection, image captioning, and optical character recognition (OCR) to convert target images to textual data. Then, it performs dense passage retrieval (DPR) (Karpukhin et al., 2020a) to fetch text documents relevant to the target image in the database. Finally, each retrieved document is concatenated with the initial question to generate the final prediction, similar to RAG (Lewis et al., 2020). Besides external documents, PICa (Yang et al., 2022b) and KAT (Gui et al., 2022) also consider LLMs as implicit knowledge bases and extract relevant implicit information from GPT-3. Plug-and-Play (Tiong et al., 2022) retrieves relevant image patches by using GradCAM (Selvaraju et al., 2017) to localize relevant parts based on the initial question. It then performs image captioning on retrieved patches to acquire augmented context. Beyond text-only augmented context, MuRAG (Chen et al., 2022b) retrieves both text and image data and incorporates images as visual tokens. RAMM (Yuan et al., 2023) retrieves similar biomedical images and captions and encodes them through different networks.

Image captioning To generate multi-style captions, Zhou and Long (2023) uses a style-aware visual encoder to retrieve image contents before generating captions. Beyond simply encoding visual information, Cho et al. (2022) further uses the multimodal similarity between image-text pairs as

a reward function to train a more fine-grained captioning model. Beyond retrieving image elements, Sarto et al. (2022); Shi et al. (2021); Ramos et al. (2023); Yang et al. (2023b) retrieves relevant captions to the inputs. Zhou et al. (2022a) tackles news image captioning by retrieving visually grounded entities in news articles.

Visually grounded dialogue This task (Lee et al., 2021b) requires retrieving visual information to produce relevant dialog responses. Fan et al. (2021) augments generative models with KNN-based Information Fetching (KIF) modules that retrieve images and wiki knowledge. Liang et al. (2021) retrieves a correlated image to the dialog from an image index to ground the response generator. Shen et al. (2021) trains a word-image mapping model to retrieve response visual impressions and then uses both textual and visual information for response generation.

Text generation For general text generation tasks, image retrieval can also help expand contexts. Yang et al. (2022a) augments a text model’s “imagination” by retrieving existing images and synthesizing newly generated images. As a result, fueling language models with imagination can improve performances in many downstream natural language tasks. Similarly, Zhu et al. (2023) compares “imagination” augmentation with synthetic and retrieved images and argues that machine-generated images could provide better guidance due to better consideration of the contexts. Moreover, Fang and Feng (2022) shows that machine translation can be significantly improved by retrieving visual information at the phrase level, especially when the textual context is limited. Image RAG can also help low-resource tasks such as medical report generation (Wu et al., 2022a) and architectural description generation (Mille et al., 2020).

Beyond retrieving images before generating text, Re-Imagen (Chen et al., 2022c) leverages a multi-modal knowledge base to retrieve image-text pairs to facilitate image generation. RA-CM3 (Yasunaga et al., 2022) can generate mixtures of images and text. It shows that retrieval-augmented image generation performs much better on knowledge-intensive generation tasks and opens up new capabilities such as multimodal in-context learning.

3.2 Code

Software developers attempt to search for relevant information to improve their productivity from large amounts of available resources such as expla-

nations for unknown terminologies, reusable code patches, and solutions to common programming bugs (Xia et al., 2017). Inspired by the progress of deep learning in NLP, a general retrieval-augmented generation paradigm has benefited a wide range of code intelligent tasks, including code completion (Lu et al., 2022b), code generation (Zhou et al., 2022b), and automatic program repair (APR) (Nashid et al., 2023). However, these approaches often treat programming languages and natural languages as equivalent sequences of tokens and ignore the rich semantics inherent to source code. To address these limitations, recent research work has focused on improving code generalization performance via multimodal learning, which incorporates additional modalities such as code comments, identifier tags, and abstract syntax trees (AST) into code pretrained models (Wang et al., 2021b; Guo et al., 2022; Li et al., 2022d). To this end, multimodal retrieval-augmented generation approach has demonstrated its feasibility in a variety of code-specific tasks.

Text-to-Code Generation Numerous research studies have investigated the utilization of relevant codes and associated documents to benefit code generation models. A prominent example is REDCODER (Parvez et al., 2021), which retrieves the top-ranked code snippets or summaries from an existing codebase, and aggregates them with source code sequences to enhance the generation or summarization capabilities. As another such approach, DocPrompting (Zhou et al., 2022b) uses a set of relevant documentation as in-context prompts to generate corresponding code via retrieval. In addition to these lexical modalities, Hayati et al. (2018) proposes a syntax-based code generation approach to reference existing subtree from the AST as templates to direct code generation explicitly.

Code-to-Text Generation Retrieval-based code summarization methods are studied extensively. For example, RACE (Shi et al., 2022) leverages relevant code differences and their associated commit messages to enhance commit message generation. Besides, RACE calculates the semantic similarity between source code differences and retrieved ones to weigh the importance of different input modalities. Rencos (Zhang et al., 2020) retrieves two similar code snippets based on the aspects of syntactic-level similarity and semantic-level similarity of a given query code. These similar contexts are then incorporated into the summarization model during

the decoding phase. This idea is further explored by Liu et al. (2021), where retrieved code-summary pairs are used to augment the original code property graph (Yamaguchi et al., 2014) of source code via local attention mechanisms. To capture the global semantics for better code structural learning, a global structure-aware self-attention mechanism (Zhu et al., 2019) is further employed.

Code Completion Recent advances in retrieval-based code completion tasks (McConnell, 2004) have gained increasing attention. Notably, Hashimoto et al. (2018) adapts the retrieve-and-edit framework to improve the model’s performance in code auto-completion tasks. To address practical code completion scenarios, ReACC (Lu et al., 2022b) takes both lexical and semantic information of the unfinished code snippet into account, utilizing a hybrid technique to combine a lexical-based sparse retriever and a semantic-based dense retriever. First, the hybrid retriever searches for a relevant code from the codebase based on the given incomplete code. Then, the unfinished code is concatenated with the retrieval, and an auto-regressive code completion generator will generate the completed code based on them. In order to address project relations, CoCoMIC (Ding et al., 2022) decomposes a code file into four components: *files*, *global variables*, *classes*, and *functions*. It constructs an in-file context graph based on the hierarchical relations among all associated code components, forming a project-level context graph by considering both in-file and cross-file dependencies. Given an incomplete program, CoCoMIC retrieves the most relevant cross-file entities from its project-level context graph and jointly learns the incomplete program and the retrieved cross-file context for code completion.

Automatic Program Repair (APR) Inspired by the nature that a remarkable portion of commits is comprised of existing code commits (Martinez et al., 2014), APR is typically treated as a search problem by traversing the search space of repair ingredients to identify a correct fix (Qi et al., 2014), based on a redundancy assumption (White et al., 2019) that the target fix can often be reconstructed in the search space. Recent studies have shown that mining relevant bug-fix patterns from existing search space (Jiang et al., 2018) and external repair templates from StackOverflow (Liu and Zhong, 2018) can significantly benefit APR

models. Joshi et al. (2022) intuitively ranks a collection of bug-fix pairs based on the similarity of error messages to develop few-shot prompts. They incorporate compiler error messages into a large programming language model Codex (Chen et al., 2021) for multilingual APR. CEDAR (Nashid et al., 2023) further extends this idea to retrieval-based prompts design using relevant code demonstrations, comprising more modalities such as unit test, error type, and error information. Additionally, Jin et al. (2023) leverage a static analyzer Infer to extract error type, error location, and syntax hierarchies (Clement et al., 2021) to prioritize the focal context. Then, they retrieve semantically-similar fixes from an existing bug-fix codebase and concatenate the retrieved fixes and focal context to form the instruction prompts for program repair.

Reasoning over Codes as Intermediate Steps

While large language models (LLMs) have recently demonstrated their impressive capability to perform reasoning tasks, they are still prone to logical and arithmetic errors (Gao et al., 2022a; Chen et al., 2022d; Madaan et al., 2022). To mitigate this issue, emerging research papers have focused on using LLMs of code (e.g., Codex (Chen et al., 2021)) to generate the code commands for solving logical and arithmetic tasks and calling external interpreters to execute the commands to obtain the results. Notably, Gao et al. (2022a) proposes to generate Python programs as intermediate reasoning steps and offload the solution step to a Python interpreter. Additionally, Chen et al. (2022d) explore generating chain-of-thought (CoT) (Wei et al., 2022) for not only text but also programming language statements as reasoning steps to solve the problem. During the inference phase, answers are obtained via an external interpreter. Similarly, Lyu et al. (2023) propose Faithful CoT that first translates the natural language query to a symbolic reasoning chain and then solves the reasoning chain by calling external executors to derive the answer. Another example is Ye et al. (2023), which utilizes LLMs to decompose table-based reasoning tasks into subtasks, decouples logic and numerical computations in each step through SQL queries by Codex, and calls SQL interpreters to solve them (a process called "parsing-execution-filling").

LLMs of code are also known as good-structured commonsense reasoners, and even better-structured reasoners than LLMs (Madaan et al., 2022). As a result, prior studies have also investigated the

idea of transforming structured commonsense generation tasks into code generation problems and employing LLMs of code as the solvers. One such work is CoCoGen (Madaan et al., 2022) which converts each training sample (consisting of textual input and the output structure) into a Tree class in Python. The LLMs of code then perform few-shot reasoning over the textual input to generate Python code, and the Python code is then converted back to the original structure for evaluation. Besides, the success of LLMs of code such as Codex in synthesizing computer code also makes them suitable for generating formal codes. Motivated by this, Wu et al. (2022b) propose to prove mathematical theorems by adopting Codex to generate formalized theorems from natural language mathematics for the interactive theorem prover Isabelle (Wenzel et al., 2008).

3.3 Structured Knowledge

An open challenge in generative models is hallucination, where the model is likely to output false information. Thus, A potential solution is to ground generation with retrieved structured knowledge, such as knowledge graphs, tables, and databases.

Question Answering (QA) A natural setting to use knowledge is QA. To augment *Knowledge Base (KB) QA* by extracting the most relevant knowledge, Hu et al. (2022b) uses dense retrieval and Liu et al. (2022b) uses a cross-encoder ranker. Shu et al. (2022) employs multi-grained retrieval to extract relevant KB context and uses constrained decoding to control the output space. In *table QA*, Nan et al. (2022) proposes a dataset that requires retrieving relevant tables for answer generation. Pan et al. (2021) then proposes a method that uses a transformer-based system to retrieve the most relevant tables and locate the correct cells. Moreover, to improve *Video QA*, Hu et al. (2023) retrieves from Knowledge Graph (KG) encodings stored in the memory. The most prominent RAG usage remains in *open-domain QA*, where multiple datasets are proposed (Lin et al., 2023; Ramnath et al., 2021). For retrieval, Ma et al. (2022) verbalizes the KG and then uses dense passage retrieval. Fan et al. (2019); Gupta et al. (2018) encodes KG information into dense representations. Pramanik et al. (2021); Jin et al. (2022) builds graph embeddings to retrieve question-relevant evidence. Xu et al. (2021); Baek et al. (2023) use semantic similarity and text-matching methods. Synthesis can occur at different stages. At the input stage, Xu

et al. (2021); Baek et al. (2023) feed in the retrieved contexts as additional inputs or prompts to the PLM. (Ma et al., 2022; Fan et al., 2019) adapt the generator to accept the context representations as inputs. At model inference stage, an interesting work is Hu et al. (2022c), which inserts an interaction layer into PLMs to guide an external KG reasoning module.

General text generation External knowledge retrieval can improve general text generation to be more factually grounded. Liu et al. (2022a) presents a memory-augmented approach to condition an autoregressive language model on a knowledge graph (KG). During inference, Tan et al. (2022) selects knowledge entries through dense retrieval and then injects them into the input encoding and output decoding stages in pretrained language models (PLMs). For *domain-specific text generation*, Frisoni et al. (2022); Yang et al. (2021); Li et al. (2019) retrieve medical report chunks or report templates to augment input prompts. Then, they use self-devised decoders or graph transformers to generate grounded reports. To improve interpretability, RAG could be used to select facts as interpretable reasoning paths (Aggarwal et al., 2021; Jansen and Ustalov, 2019). Moreover, RAG is especially useful for low-resource generation tasks, such as question generation (Yu and Jiang, 2021; Xin et al., 2021; Gu et al., 2019), document-to-slide (Sun et al., 2021), table-to-text (Su et al., 2021), counterargument generation (Jo et al., 2021), entity description generation (Cheng et al., 2020) and text-based games (Murugesan et al., 2021).

Recent research has attempted to reduce hallucinations in LLMs by leveraging external structured knowledge. For example, during fine-tuning, LaMDA (Thoppilan et al., 2022) learns to consult external knowledge sources before responding to the user, including an information retrieval system that can retrieve knowledge triplets and web URLs. Some papers treat the generative model (often large language models) as black-box and retrieve structured information without fine-tuning. For example, BINDER (Cheng et al., 2023) uses in-context learning to output designed API calls that retrieve question-relevant columns from tables.

Reasoning with knowledge By selecting knowledge, reasoning tasks can be solved in a more grounded and interpretable way. To generate an entailment tree explanation for a given hypothesis, Neves Ribeiro et al. (2022) retrieves from tex-

tual premises iteratively and combines them with generation. Yang et al. (2022c) proposes a math reasoner that first retrieves highly-correlated algebraic knowledge and then passes them as prompts to improve the semantic representations for the generation task. With the recent advances in LLMs, He et al. (2022a); Li et al. (2023b) retrieve from KG and KB, such as Wikidata, based on reasoning steps obtained from the chain-of-thought (CoT) prompting (Wei et al., 2022).

Knowledge-grounded dialogue Dialogue generation based on relevant tables and knowledge bases has been a practical research application (Wu et al., 2020b; Li et al., 2022b; Nakamura et al., 2022; Gao et al., 2022b; Lu et al., 2023). To tackle the challenge, Li et al. (2022c) and Galetzka et al. (2021) retrieve relevant knowledge, process it into a dense representation and incorporate it into dialogue generation. On top of dense representations, Gu et al. (2020) and Jung et al. (2020) leverage attention mechanisms to flexibly adjust which knowledge to depend on during generation. Some methods (Zhang et al., 2021; Dziri et al., 2021; Chen et al., 2020) first generate subgoals or responses and then use them to retrieve relevant knowledge. The retrieved knowledge then helps amend previous responses. Besides knowledge, Cai et al. (2019b) and Wu et al. (2020a) improve dialogue response generation by retrieving templates or prototype dialogues to augment inputs. Recently, Kang et al. (2023) retrieves relevant subgraphs from KGs, and then utilizes contrastive learning to ensure that the generated texts have high similarity to the subgraphs.

By retrieving from relevant sources, RAG not only improves factuality but also provides the grounding contexts while generating, thus addressing interpretability and robustness concerns. With the potential to handle more information types with recent advances in LLMs (OpenAI, 2023), RAG with structured knowledge could be further enhanced. There are still challenges to be addressed. For example, there could be new designs for better retrieval systems that could promote efficient interactions suitable for diverse knowledge bases. Synthesizing this information correctly is also an open challenge, where it is hard to decide which parts need augmenting in the textual outputs.

3.4 Audio

Audio RAG is useful in incorporating audio information in specific audio-language tasks, such as

music captioning, music and text generation, and speech recognition. Moreover, using audio RAG for audio data augmentation has also been proven useful in mitigating the lack of audio-text training data. It could be a promising future direction (Li et al., 2022a).

Text-audio data augmentation For text-audio tasks, one of the most important challenges is the lack of training data on audio-text pairs. Therefore, retrieving audio and textual cues can alleviate the data scarcity problem and improve performance. In audio captioning, which aims at translating the input audio into its description, Koizumi et al. (2020) retrieves guidance captions similar to the input audio from the training set. Then, the retrieved guidance captions are fed into a PLM to help generate new captions, which improves generation performance. To augment scarce speech translation (ST) data, Zhao et al. (2023a) proposes Spoken-Vocab, a technique to convert machine translation (MT) data to synthetic ST data. To form synthetic speech, SpokenVocab retrieves and stitches audio snippets, corresponding to words in an MT sentence. Experiments show that stitched audio snippets can improve translation quality. Kim et al. (2023) leverages a PLM to tackle the data scarcity issue. It retrieves features from the input audio, maps them to continuous vectors using mapping networks, and uses vectors as prefixes for prefix tuning the PLM. With the additional information from retrieved audio, it outperforms previous methods. In text-to-audio generation, Huang et al. (2023a) applies audio-text retrieval to get pseudo text prompts, which enhance audio generation in data-scarce scenarios. To augment the argumentation mining (AM) task in political debates, Mestre et al. (2023) integrates audio features into PLMs, which improves performance when data is scarce.

Music captioning Music captioning is the task of generating a text description or lyrics given the music audio. And RAG is explored to learn better audio-lyric alignment. Manco et al. (2021) proposes the first music audio captioning model, MusCaps. Firstly, a pretrained multimodal encoder obtains audio representations that retrieve musical features in the input. As the pretraining bridges the gap between the audio modality and textual understanding, the method improves task performance. He et al. (2022b) learns an audio-lyric alignment through contrastive learning, which results in a higher-quality generation of captions for music.

Music generation Royal et al. (2020) uses deep neural hashing to retrieve music building blocks and then performs generation by using the current music segment to retrieve the next. In automatic speech recognition (ASR), Chan et al. (2023) uses a k-nearest neighbor (KNN) approach to retrieve external knowledge related to the audio and text embeddings. The retrieved knowledge significantly reduces domain adaptation time for ASR.

The audio modality is closely intertwined with other modalities, such as video. Therefore, recent advancements in using audio features for text-video retrieval (Falcon et al., 2022; Mithun et al., 2018) can benefit RAG tasks involving other modalities. Moreover, although audio-text retrieval has been a long-standing task (Liu et al., 2015; Milde et al., 2016a,b), exploring recently discovered techniques (Hu et al., 2022a; Lou et al., 2022; Koepke et al., 2022) could lead to further improvements.

3.5 Video

Retrieving video snippets for generation is used primarily in two tasks: video-grounded dialogue and video captioning. Recently, augmenting LLMs with video retrieval also demonstrates good performances, especially in few-shot settings.

Video-grounded dialogue Given video contexts, the model learns to engage in a relevant dialogue. Pasunuru and Bansal (2018) introduces a video-context, many-speaker dialogue dataset, which challenges researchers to develop visually-grounded dialogue models that generate relevant responses from live videos. Similarly, Lei et al. (2020) proposes TVQA+, a dataset that requires retrieving relevant video moments to answer textual questions about videos. Then, it proposes a unified framework that encodes video segments into representations, uses an attention mechanism to locate relevant information, and produces textual answers. To better perform visually-grounded dialogue tasks, Le et al. (2020) retrieves visual cues from prior user queries. The cues are then used as contextual information to construct relevant responses. On video QA, it substantially outperforms prior approaches. Recently, Le et al. (2022) extracts visual cues from the video to augment video-grounded dialogues. The video retrieval is performed with neural module networks, which are instantiated with entities and actions in previous dialogues.

Video captioning Sharing a similar motivation to RAG, Long et al. (2018) first proposes to use attention layers to automatically select the most

salient visual or semantic features and use them to augment caption generation. As a result, it stably outperforms previous methods. (Whitehead et al., 2018) then develops a retrieval-based approach for video description generation. For news videos, it retrieves topically related news documents and then generates a description using a knowledge-aware video description network.

LLM augmentation Wang et al. (2022) attempts to augment an LLM to generalize to various video-to-text tasks from a few examples. As the LLMs cannot accept video inputs, it first translates video contents into attributes using image-language models and then prompts the retrieved content to instruct the LLM. It demonstrates good few-shot performances on a wide range of video-language tasks.

Currently, the video-text research bottleneck mainly lies in the representation gap between different modalities. Research has been attempting to learn a better mapping between video-text via joint learning (Xu et al., 2015; Sun et al., 2019). Recent studies on dense video representation learning can also be useful for future video RAG. Besides, some papers (Yang et al., 2023a; Wang et al., 2021a) try to introduce fine-grained interaction between different modalities to learn better aligned representations. Zeng et al. (2022) encourages multiple pretrained models in different modalities to exchange information with each other in a zero-shot manner. Most recently, Zhang et al. (2023a) trains Video-Llama to better align pretrained video and audio encoders with LLM’s embedding space.

4 Future Directions

With the development of multi-modal LLMs, retrieving multimodal information to augment text generation will be a promising direction to better ground textual generation in real-world contexts, contributing towards building a model that is fully aware and can better interact with the world. Specifically, we describe some potential directions that can be of benefit to the community.

4.1 Retrieval Augmented Multimodal Reasoning

One potential application of multimodal RAG is multimodal reasoning. Lu et al. (2022a) first introduces ScienceQA, a large-scale multimodal science question dataset annotated with lectures and explanations. Then, Zhang et al. (2023b) proposes Multimodal Chain-of-Thought (Multimodal-CoT) which incorporates language and vision modalities

into a two-stage (rationale generation and answer inference) framework, surpassing GPT-3.5 by a large margin with a much smaller fine-tuned model. Similar to Zhang et al. (2023b), kosmos-1 (Huang et al., 2023b) breaks down multimodal reasoning into two steps. It first generates intermediate content as the rationale based on visual information and then uses the generated rationale to induce the result. However, both methods may have difficulties understanding certain types of images (e.g., maps), which could be mitigated by retrieving informative image-text pairs.

4.2 Building a Multimodal Knowledge Index

In order to facilitate multimodal RAG, one of the most fundamental aspects is building a multimodal knowledge index. The goal is twofold: Firstly, dense representations should support low storage, dynamic updating of the knowledge base, and accurate search. Secondly, it could enable faster search speed with the help of local sensitive hashing (Leskovec et al., 2014), which combats scaling and robustness concerns when the knowledge base is scaled up extremely.

Currently, the dense representations for text snippets are widely studied for documents (Karpukhin et al., 2020b; Gao and Callan, 2021; Gao et al., 2021), entities (Sciavolino et al., 2021; Lee et al., 2021a), and images (Radford et al., 2021a). Besides, there are studies optimizing dense representations in an end-to-end manner (Lewis et al., 2020). Nevertheless, few papers (Chen et al., 2022a) have explored building a multimodal index at the same time for downstream generation tasks. How to map a multimodal knowledge index into a unified space remains a long-term challenge.

4.3 Pretraining with Multimodal Retrieval

To better align the abilities to handle different modalities in a pre-trained model, future work could be built on employing retrieval-based approaches during pre-training. Currently, some methods fine-tune the pre-trained generative model to learn to retrieve from different modalities. For example, LaMDA (Thoppilan et al., 2022) calls an external toolset for fine-tuning, including an information retrieval system. Similarly, during fine-tuning, Toolformer (Schick et al., 2023) augments models with API calls to tools including a QA system and a Wikipedia search engine.

When similar retrieval abilities are leveraged during pretraining, the generative models can interact

with retrieval tools much better. Then, instead of relying solely on internal weights, they could effectively use an external base to output more grounded information, provide relevant contexts to users, and update their information accordingly. Such pre-training techniques would also greatly improve robustness for out-of-domain tasks. As an example, Guu et al. (2020a) augments pretraining with an external knowledge retriever, which outperforms previous methods.

To incorporate retrieval with pretraining, there remains the challenge of developing appropriate datasets labeled with retrieval API calls. To tackle this challenge, LaMDA (Thoppilan et al., 2022) uses labels developed by human annotators, which could be expensive to collect. Toolformer (Schick et al., 2023) uses a sampling and filtering approach for automatic labeling, which is inexpensive but could induce bias. A potential solution is to use a neuro-symbolic approach (Davoudi and Komeili, 2021), which uses prototype learning and deep-KNN to find nearest neighbors during training.

5 Conclusions

This survey reviews research that augments generative models by retrieving multi-modal information. Specifically, we categorize the current domain into enhancing with different modalities, including image, code, structured knowledge, speech, and video. With the emergence of large multi-modal models, we believe that this survey could serve as a comprehensive overview of an emerging and promising field. Moreover, we hope it could encourage future research in the domain, including retrieval-augmented multimodal reasoning, building a multi-modal knowledge index, and combining retrieval with pretraining.

6 Limitations

RAG also has some limitations. For example, there exists an attribution-fluency tradeoff (Aksitov et al., 2023) where the output quality is affected due to the added constraints of the retrieved knowledge.

References

Julián N Acosta, Guido J Falcone, Pranav Rajpurkar, and Eric J Topol. 2022. Multimodal biomedical ai. *Nature Medicine*, 28(9):1773–1784.

Shourya Aggarwal, Divyanshu Mandowara, Vishwa-jeet Agrawal, Dinesh Khandelwal, Parag Singla, and

Dinesh Garg. 2021. *Explanations for CommonsenseQA: New Dataset and Models*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3050–3065, Online. Association for Computational Linguistics.

Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, and Luke Zettlemoyer. 2022. CM3: A causal masked multimodal model of the internet. *CoRR*, abs/2201.07520.

Renat Aksitov, Chung-Ching Chang, David Reitter, Siamak Shakeri, and Yunhsuan Sung. 2023. *Characterizing attribution and fluency tradeoffs for retrieval-augmented large language models*.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022a. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022b. Flamingo: a visual language model for few-shot learning. *CoRR*, abs/2204.14198.

Jinheon Baek, Alham Fikri Aji, and Amir Saffari. 2023. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. *arXiv preprint arXiv:2306.04136*.

Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pages 2206–2240. PMLR.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

- Deng Cai, Yan Wang, Wei Bi, Zhaopeng Tu, Xiaojiang Liu, Wai Lam, and Shuming Shi. 2019a. [Skeleton-to-response: Dialogue generation guided by retrieval memory](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1219–1228, Minneapolis, Minnesota. Association for Computational Linguistics.
- Deng Cai, Yan Wang, Wei Bi, Zhaopeng Tu, Xiaojiang Liu, and Shuming Shi. 2019b. [Retrieval-guided dialogue response generation via a matching-to-generation framework](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1866–1875, Hong Kong, China. Association for Computational Linguistics.
- David M Chan, Shalini Ghosh, Ariya Rastrow, and Björn Hoffmeister. 2023. Using external off-policy speech-to-text mappings in contextual end-to-end automated speech recognition. *arXiv preprint arXiv:2301.02736*.
- Chieh-Yang Chen, Pei-Hsin Wang, Shih-Chieh Chang, Da-Cheng Juan, Wei Wei, and Jia-Yu Pan. 2020. [Air-Concierge: Generating task-oriented dialogue via efficient large-scale knowledge retrieval](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 884–897, Online. Association for Computational Linguistics.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William Cohen. 2022a. [MuRAG: Multimodal retrieval-augmented generator for open question answering over images and text](#). In *EMNLP*, pages 5558–5570. ACL.
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William Cohen. 2022b. [MuRAG: Multimodal retrieval-augmented generator for open question answering over images and text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5558–5570, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Wenhu Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. 2022c. [Re-Imagen: Retrieval-augmented text-to-image generator](#). *arXiv preprint arXiv:2209.14491*.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022d. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.
- Liying Cheng, Dekun Wu, Lidong Bing, Yan Zhang, Zhanming Jie, Wei Lu, and Luo Si. 2020. [ENT-DESC: Entity description generation by exploring knowledge graph](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1187–1197, Online. Association for Computational Linguistics.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. 2023. Binding language models in symbolic languages. *ICLR*.
- Jaemin Cho, Seunghyun Yoon, Ajinkya Kale, Franck Dernoncourt, Trung Bui, and Mohit Bansal. 2022. [Fine-grained image captioning with CLIP reward](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 517–527, Seattle, United States. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Colin B Clement, Shuai Lu, Xiaoyu Liu, Michele Tufano, Dawn Drain, Nan Duan, Neel Sundaresan, and Alexey Svyatkovskiy. 2021. Long-range modeling of source code files with ewash: Extended window access by syntax hierarchy. *arXiv preprint arXiv:2109.08780*.
- Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. 2022. Vqgan-clip: Open domain image generation and editing with natural language guidance. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII*, pages 88–105. Springer.
- Seyed Omid Davoudi and Majid Komeili. 2021. Toward faithful case-based reasoning through learning prototypes in a nearest neighbor-friendly space. In *International Conference on Learning Representations*.
- Yangruibo Ding, Zijian Wang, Wasi Uddin Ahmad, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, Dan Roth, and Bing Xiang. 2022. Cocomic: Code completion by jointly modeling in-file and cross-file context. *arXiv preprint arXiv:2212.10007*.
- Zi-Yi Dou, Aishwarya Kamath, Zhe Gan, Pengchuan Zhang, Jianfeng Wang, Linjie Li, Zicheng Liu, Ce Liu, Yann LeCun, Nanyun Peng, et al. 2022. Coarse-to-fine vision-language pre-training with fusion in the backbone. *arXiv preprint arXiv:2206.07643*.

- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. 2023. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*.
- Nouha Dziri, Andrea Madotto, Osmar Zaiane, and Avishek Joey Bose. 2021. **Neural path hunter: Reducing hallucination in dialogue systems via path grounding**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2197–2214, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alex Falcon, Giuseppe Serra, and Oswald Lanz. 2022. A feature-space multimodal data augmentation technique for text-video retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4385–4394.
- Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2019. Using local knowledge graph construction to scale seq2seq models to multi-document inputs. *arXiv preprint arXiv:1910.08435*.
- Angela Fan, Claire Gardent, Chloé Braud, and Antoine Bordes. 2021. **Augmenting transformers with KNN-based composite memory for dialog**. *Transactions of the Association for Computational Linguistics*, 9:82–99.
- Qingkai Fang and Yang Feng. 2022. **Neural machine translation with phrase-level universal visual representations**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5687–5698, Dublin, Ireland. Association for Computational Linguistics.
- Giacomo Frisoni, Miki Mizutani, Gianluca Moro, and Lorenzo Valgimigli. 2022. **BioReader: a retrieval-enhanced text-to-text transformer for biomedical literature**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5770–5793, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Fabian Galetzka, Jewgeni Rose, David Schlangen, and Jens Lehmann. 2021. **Space efficient context encoding for non-task-oriented dialogue generation with graph attention transformer**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7028–7041, Online. Association for Computational Linguistics.
- Jing Gao, Peng Li, Zhikui Chen, and Jianing Zhang. 2020. A survey on deep learning for multimodal data fusion. *Neural Computation*, 32(5):829–864.
- Luyu Gao and Jamie Callan. 2021. **Condenser: a pre-training architecture for dense retrieval**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 981–993, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2022a. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*.
- Silin Gao, Jena D. Hwang, Saya Kanno, Hiromi Wakaki, Yuki Mitsufuji, and Antoine Bosselut. 2022b. **ComFact: A benchmark for linking contextual common-sense knowledge**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1656–1675, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jia-Chen Gu, Zhenhua Ling, Quan Liu, Zhigang Chen, and Xiaodan Zhu. 2020. **Filtering before iteratively referring for knowledge-grounded response selection in retrieval-based chatbots**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1412–1422, Online. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. 2018. Search engine guided neural machine translation. In *AAAI*, volume 32.
- Yunfan Gu, Yang Yuqiao, and Zhongyu Wei. 2019. **Extract, transform and filling: A pipeline model for question paraphrasing based on template**. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 109–114, Hong Kong, China. Association for Computational Linguistics.
- Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander Hauptmann, Yonatan Bisk, and Jianfeng Gao. 2022. **KAT: A knowledge augmented transformer for vision-and-language**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 956–968, Seattle, United States. Association for Computational Linguistics.
- Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. **Unixcoder: Unified cross-modal pre-training for code representation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, *ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 7212–7225. Association for Computational Linguistics.
- Vishal Gupta, Manoj Chinnakotla, and Manish Shrivastava. 2018. **Retrieve and re-rank: A simple and effective IR approach to simple question answering over knowledge graphs**. In *Proceedings of the*

- First Workshop on Fact Extraction and VERification (FEVER)*, pages 22–27, Brussels, Belgium. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020a. [Realm: Retrieval-augmented language model pre-training](#).
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020b. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Tatsunori B Hashimoto, Kelvin Guu, Yonatan Oren, and Percy S Liang. 2018. A retrieve-and-edit framework for predicting structured outputs. *NeurIPS*, 31.
- Shirley Anugrah Hayati, Raphael Olivier, Pravalika Avvaru, Pengcheng Yin, Anthony Tomasic, and Graham Neubig. 2018. [Retrieval-based neural code generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 925–930, Brussels, Belgium. Association for Computational Linguistics.
- Hangfeng He, Hongming Zhang, and Dan Roth. 2022a. Rethinking with retrieval: Faithful large language model inference. *arXiv preprint arXiv:2301.00303*.
- Qiuxiang He, Guoping Huang, Qu Cui, Li Li, and Lemao Liu. 2021. Fast and accurate neural machine translation with 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 (Volume 1: Long Papers)*, pages 3170–3180.
- Zihao He, Weituo Hao, and Xuchen Song. 2022b. Re-cap: Retrieval augmented music captioner. *arXiv preprint arXiv:2212.10901*.
- Tao Hu, Xuyu Xiang, Jiaohua Qin, and Yun Tan. 2022a. Audio-text retrieval based on contrastive learning and collaborative attention mechanism.
- Xixin Hu, Xuan Wu, Yiheng Shu, and Yuzhong Qu. 2022b. [Logical form generation via multi-task learning for complex question answering over knowledge bases](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1687–1696, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, Kai-Wei Chang, Yizhou Sun, Cordelia Schmid, David A Ross, and Alireza Fathi. 2023. Reveal: Retrieval-augmented visual-language pre-training with multi-source multimodal knowledge memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23369–23379.
- Ziniu Hu, Yichong Xu, Wenhao Yu, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Kai-Wei Chang, and Yizhou Sun. 2022c. [Empowering language models with knowledge graph reasoning for open-domain question answering](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9562–9581, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. 2023a. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. *arXiv preprint arXiv:2301.12661*.
- Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. 2023b. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*.
- 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: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. Atlas: Few-shot learning with retrieval augmented language models. *arXiv preprint arXiv*, 2208.
- Peter Jansen and Dmitry Ustalov. 2019. [TextGraphs 2019 shared task on multi-hop inference for explanation regeneration](#). In *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*, pages 63–77, Hong Kong. Association for Computational Linguistics.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR.
- Jiajun Jiang, Yingfei Xiong, Hongyu Zhang, Qing Gao, and Xiangqun Chen. 2018. Shaping program repair space with existing patches and similar code. In *ISSTA*, pages 298–309. ACM.
- Bowen Jin, Yu Zhang, Qi Zhu, and Jiawei Han. 2022. Heterformer: A transformer architecture for node representation learning on heterogeneous text-rich networks. *arXiv preprint arXiv:2205.10282*.
- Matthew Jin, Syed Shahriar, Michele Tufano, Xin Shi, Shuai Lu, Neel Sundaresan, and Alexey Svyatkovskiy. 2023. Inferfix: End-to-end program repair with llms. *arXiv preprint arXiv:2303.07263*.

- Yohan Jo, Haneul Yoo, JinYeong Bak, Alice Oh, Chris Reed, and Eduard Hovy. 2021. [Knowledge-enhanced evidence retrieval for counterargument generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3074–3094, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Harshit Joshi, José Cambronero, Sumit Gulwani, Vu Le, Ivan Radicek, and Gust Verbruggen. 2022. Repair is nearly generation: Multilingual program repair with llms. *arXiv preprint arXiv:2208.11640*.
- Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. 2022. Prompting visual-language models for efficient video understanding. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 105–124. Springer.
- Jaehun Jung, Bokyung Son, and Sungwon Lyu. 2020. [AttnIO: Knowledge Graph Exploration with In-and-Out Attention Flow for Knowledge-Grounded Dialogue](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3484–3497, Online. Association for Computational Linguistics.
- Minki Kang, Jin Myung Kwak, Jinheon Baek, and Sung Ju Hwang. 2023. [Knowledge graph-augmented language models for knowledge-grounded dialogue generation](#).
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020a. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020b. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Minkyu Kim, Kim Sung-Bin, and Tae-Hyun Oh. 2023. Prefix tuning for automated audio captioning. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- A Sophia Koepke, Andreea-Maria Oncescu, Joao Henriques, Zeynep Akata, and Samuel Albanie. 2022. Audio retrieval with natural language queries: A benchmark study. *IEEE Transactions on Multimedia*.
- Yuma Koizumi, Yasunori Ohishi, Daisuke Niizumi, Daiki Takeuchi, and Masahiro Yasuda. 2020. Audio captioning using pre-trained large-scale language model guided by audio-based similar caption retrieval. *arXiv preprint arXiv:2012.07331*.
- Hung Le, Nancy Chen, and Steven Hoi. 2022. [Vgmn: Video-grounded neural module networks for video-grounded dialogue systems](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3377–3393.
- Hung Le, Doyen Sahoo, Nancy Chen, and Steven C.H. Hoi. 2020. [BiST: Bi-directional spatio-temporal reasoning for video-grounded dialogues](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1846–1859, Online. Association for Computational Linguistics.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021a. [Learning dense representations of phrases at scale](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6634–6647, Online. Association for Computational Linguistics.
- Nyoungwoo Lee, Suwon Shin, Jaegul Choo, Ho-Jin Choi, and Sung-Hyon Myaeng. 2021b. [Constructing multi-modal dialogue dataset by replacing text with semantically relevant images](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 897–906, Online. Association for Computational Linguistics.
- Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. 2021. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7331–7341.
- Jie Lei, Licheng Yu, Tamara Berg, and Mohit Bansal. 2020. [TVQA+: Spatio-temporal grounding for video question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8211–8225, Online. Association for Computational Linguistics.
- Jure Leskovec, Anand Rajaraman, and Jeffrey D. Ullman. 2014. *Mining of Massive Datasets, 2nd Ed.* Cambridge University Press.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Christy Y. Li, Xiaodan Liang, Zhiting Hu, and Eric P. Xing. 2019. [Knowledge-driven encode, retrieve, paraphrase for medical image report generation](#).
- Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemao Liu. 2022a. A survey on retrieval-augmented text generation. *arXiv preprint arXiv:2202.01110*.

- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Miaoran Li, Baolin Peng, Jianfeng Gao, and Zhu Zhang. 2022b. Opera: Harmonizing task-oriented dialogs and information seeking experience. *arXiv preprint arXiv:2206.12449*.
- Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Lidong Bing, Shafiq Joty, and Soujanya Poria. 2023b. Chain of knowledge: A framework for grounding large language models with structured knowledge bases.
- Yu Li, Baolin Peng, Yelong Shen, Yi Mao, Lars Liden, Zhou Yu, and Jianfeng Gao. 2022c. Knowledge-grounded dialogue generation with a unified knowledge representation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 206–218, Seattle, United States. Association for Computational Linguistics.
- Zhiyu Li, Shuai Lu, Daya Guo, Nan Duan, Shailesh Jannu, Grant Jenks, Deep Majumder, Jared Green, Alexey Svyatkovskiy, Shengyu Fu, and Neel Sundaresan. 2022d. Automating code review activities by large-scale pre-training. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022, Singapore, Singapore, November 14-18, 2022*, pages 1035–1047. ACM.
- Zujie Liang, Huang Hu, Can Xu, Chongyang Tao, Xubo Geng, Yining Chen, Fan Liang, and Daxin Jiang. 2021. Maria: A visual experience powered conversational agent. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5596–5611, Online. Association for Computational Linguistics.
- Weizhe Lin and Bill Byrne. 2022a. Retrieval augmented visual question answering with outside knowledge. In *EMNLP*, pages 11238–11254. Association for Computational Linguistics.
- Weizhe Lin and Bill Byrne. 2022b. Retrieval augmented visual question answering with outside knowledge. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11238–11254, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Weizhe Lin, Zhilin Wang, and Bill Byrne. 2023. FVQA 2.0: Introducing adversarial samples into fact-based visual question answering. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 149–157, Dubrovnik, Croatia. Association for Computational Linguistics.
- Qi Liu, Dani Yogatama, and Phil Blunsom. 2022a. Relational memory-augmented language models. *Transactions of the Association for Computational Linguistics*, 10:555–572.
- Shangqing Liu, Yu Chen, Xiaofei Xie, Jing Kai Siow, and Yang Liu. 2021. Retrieval-augmented generation for code summarization via hybrid GNN. In *ICLR*.
- Shih-Hung Liu, Kuan-Yu Chen, Berlin Chen, Hsin-Min Wang, Hsu-Chun Yen, and Wen-Lian Hsu. 2015. Combining relevance language modeling and clarity measure for extractive speech summarization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(6):957–969.
- Xuliang Liu and Hao Zhong. 2018. Mining stackoverflow for program repair. In *SANER*, pages 118–129. IEEE Computer Society.
- Ye Liu, Semih Yavuz, Rui Meng, Dragomir Radev, Caiming Xiong, and Yingbo Zhou. 2022b. Uniparser: Unified semantic parser for question answering on knowledge base and database. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8858–8869, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xiang Long, Chuang Gan, and Gerard De Melo. 2018. Video captioning with multi-faceted attention. *Transactions of the Association for Computational Linguistics*, 6:173–184.
- Siyu Lou, Xuenan Xu, Mengyue Wu, and Kai Yu. 2022. Audio-text retrieval in context. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4793–4797. IEEE.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022a. Learn to explain: Multimodal reasoning via thought chains for science question answering. *arXiv preprint arXiv:2209.09513*.
- Shuai Lu, Nan Duan, Hojae Han, Daya Guo, Seung-won Hwang, and Alexey Svyatkovskiy. 2022b. ReACC: A retrieval-augmented code completion framework. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6227–6240, Dublin, Ireland. Association for Computational Linguistics.
- Xing Han Lu, Siva Reddy, and Harm de Vries. 2023. The StatCan dialogue dataset: Retrieving data tables through conversations with genuine intents. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2799–2829, Dubrovnik, Croatia. Association for Computational Linguistics.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. *arXiv preprint arXiv:2301.13379*.

- Kaixin Ma, Hao Cheng, Xiaodong Liu, Eric Nyberg, and Jianfeng Gao. 2022. [Open-domain question answering via chain of reasoning over heterogeneous knowledge](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5360–5374, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aman Madaan, Shuyan Zhou, Uri Alon, Yiming Yang, and Graham Neubig. 2022. Language models of code are few-shot commonsense learners. *arXiv preprint arXiv:2210.07128*.
- Ilaria Manco, Emmanouil Benetos, Elio Quinton, and György Fazekas. 2021. Muscaps: Generating captions for music audio. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Matias Martinez, Westley Weimer, and Monperrus Martin. 2014. Do the fix ingredients already exist? an empirical inquiry into the redundancy assumptions of program repair approaches. *Companion Proceedings of the 36th International Conference on Software Engineering*.
- Steve McConnell. 2004. *Code complete*. Pearson Education.
- Rafael Mestre, Stuart Middleton, Matt Ryan, Masood Gheasi, Timothy Norman, and Jiatong Zhu. 2023. [Augmenting pre-trained language models with audio feature embedding for argumentation mining in political debates](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 274–288, Dubrovnik, Croatia. Association for Computational Linguistics.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. 2023. Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*.
- Benjamin Milde, Jonas Wacker, Stefan Radomski, Max Mühlhäuser, and Chris Biemann. 2016a. [Ambient search: A document retrieval system for speech streams](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2082–2091, Osaka, Japan. The COLING 2016 Organizing Committee.
- Benjamin Milde, Jonas Wacker, Stefan Radomski, Max Mühlhäuser, and Chris Biemann. 2016b. [Demonstrating ambient search: Implicit document retrieval for speech streams](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations*, pages 233–237, Osaka, Japan. The COLING 2016 Organizing Committee.
- Simon Mille, Spyridon Symeonidis, Maria Rousi, Montserrat Marimon Felipe, Klearchos Stavrothanasopoulos, Petros Alvanitopoulos, Roberto Carlini Salguero, Jens Grivolla, Georgios Meditskos, Stefanos Vrochidis, and Leo Wanner. 2020. [A case study of NLG from multimedia data sources: Generating architectural landmark descriptions](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 2–14, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. 2018. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, pages 19–27.
- Keerthiram Murugesan, Mattia Atzeni, Pavan Kapani-pathi, Kartik Talamadupula, Mrinmaya Sachan, and Murray Campbell. 2021. [Efficient text-based reinforcement learning by jointly leveraging state and commonsense graph representations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 719–725, Online. Association for Computational Linguistics.
- Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. 2021. Attention bottlenecks for multimodal fusion. *Advances in Neural Information Processing Systems*, 34:14200–14213.
- Kai Nakamura, Sharon Levy, Yi-Lin Tuan, Wenhu Chen, and William Yang Wang. 2022. [HybridDialogue: An information-seeking dialogue dataset grounded on tabular and textual data](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 481–492, Dublin, Ireland. Association for Computational Linguistics.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Hailey Schoelkopf, Riley Kong, Xiangru Tang, Mutethia Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, Dragomir Radev, and Dragomir Radev. 2022. [FeTaQA: Free-form table question answering](#). *Transactions of the Association for Computational Linguistics*, 10:35–49.
- Noor Nashid, Mifta Sintaha, and Ali Mesbah. 2023. Retrieval-based prompt selection for code-related few-shot learning. In *Proceedings of the 45th International Conference on Software Engineering (ICSE’23)*.
- Danilo Neves Ribeiro, Shen Wang, Xiaofei Ma, Rui Dong, Xiaokai Wei, Henghui Zhu, Xinchu Chen, Peng Xu, Zhiheng Huang, Andrew Arnold, and Dan

- Roth. 2022. [Entailment tree explanations via iterative retrieval-generation reasoner](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 465–475, Seattle, United States. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Feifei Pan, Mustafa Canim, Michael Glass, Alfio Gliozzo, and Peter Fox. 2021. [CLTR: An end-to-end, transformer-based system for cell-level table retrieval and table question answering](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 202–209, Online. Association for Computational Linguistics.
- Md Rizwan Parvez, Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. [Retrieval augmented code generation and summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2719–2734, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Game-based video-context dialogue. *arXiv preprint arXiv:1809.04560*.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are nlp models really able to solve simple math word problems? *arXiv preprint arXiv:2103.07191*.
- Hao Peng, Ankur Parikh, Manaal Faruqi, Bhuwan Dhingra, and Dipanjan Das. 2019. [Text generation with exemplar-based adaptive decoding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2555–2565, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*.
- Soumajit Pramanik, Jesujoba Alabi, Rishiraj Saha Roy, and Gerhard Weikum. 2021. Uniqorn: unified question answering over rdf knowledge graphs and natural language text. *arXiv preprint arXiv:2108.08614*.
- Yuhua Qi, Xiaoguang Mao, Yan Lei, Ziyang Dai, and Chengsong Wang. 2014. The strength of random search on automated program repair. In *ICSE*, pages 254–265. ACM.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021a. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021b. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021a. Zero-shot text-to-image generation. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. PMLR.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021b. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Kiran Ramnath, Leda Sari, Mark Hasegawa-Johnson, and Chang Yoo. 2021. [Worldly wise \(WoW\) - cross-lingual knowledge fusion for fact-based visual spoken-question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1908–1919, Online. Association for Computational Linguistics.
- Rita Ramos, Desmond Elliott, and Bruno Martins. 2023. [Retrieval-augmented image captioning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3666–3681, Dubrovnik, Croatia. Association for Computational Linguistics.
- Brandon Royal, Kien Hua, and Brenton Zhang. 2020. Deep composer: Deep neural hashing and retrieval approach to automatic music generation. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE.
- Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2022. Retrieval-augmented transformer for image captioning. In *CBMI*, pages 1–7. ACM.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761*.
- Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. [Simple entity-centric questions challenge dense retrievers](#). In *Proceedings of*

- the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626. IEEE Computer Society.
- Lei Shen, Haolan Zhan, Xin Shen, Yonghao Song, and Xiaofang Zhao. 2021. **Text is not enough: Integrating visual impressions into open-domain dialogue generation**. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 4287–4296, New York, NY, USA. Association for Computing Machinery.
- Ensheng Shi, Yanlin Wang, Wei Tao, Lun Du, Hongyu Zhang, Shi Han, Dongmei Zhang, and Hongbin Sun. 2022. **RACE: Retrieval-augmented commit message generation**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5520–5530, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zhan Shi, Hui Liu, Martin Renqiang Min, Christopher Malon, Li Erran Li, and Xiaodan Zhu. 2021. **Retrieval, analogy, and composition: A framework for compositional generalization in image captioning**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1990–2000, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yiheng Shu, Zhiwei Yu, Yuhan Li, Börje Karlsson, Tingting Ma, Yuzhong Qu, and Chin-Yew Lin. 2022. **TIARA: Multi-grained retrieval for robust question answering over large knowledge base**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8108–8121, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yixuan Su, Zaiqiao Meng, Simon Baker, and Nigel Collier. 2021. **Few-shot table-to-text generation with prototype memory**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 910–917, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473.
- Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang, and Nancy X. R. Wang. 2021. **D2S: Document-to-slide generation via query-based text summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1405–1418, Online. Association for Computational Linguistics.
- Chao-Hong Tan, Jia-Chen Gu, Chongyang Tao, Zhen-Hua Ling, Can Xu, Huang Hu, Xiubo Geng, and Daxin Jiang. 2022. **TegTok: Augmenting text generation via task-specific and open-world knowledge**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1597–1609, Dublin, Ireland. Association for Computational Linguistics.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. **Lamda: Language models for dialog applications**. *arXiv preprint arXiv:2201.08239*.
- Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese, and Steven C.H. Hoi. 2022. **Plug-and-play VQA: Zero-shot VQA by conjoining large pre-trained models with zero training**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 951–967, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. **Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions**. *arXiv preprint arXiv:2212.10509*.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. **Multimodal transformer for unaligned multimodal language sequences**. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Xiaohan Wang, Linchao Zhu, and Yi Yang. 2021a. **T2VLAD: global-local sequence alignment for text-video retrieval**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 5079–5088. Computer Vision Foundation / IEEE.
- Yue Wang, Weishi Wang, Shafiq Joty, and Steven C.H. Hoi. 2021b. **CodeT5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8696–8708, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhenhailong Wang, Manling Li, Ruochen Xu, Luwei Zhou, Jie Lei, Xudong Lin, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Derek Hoiem, et al. 2022. **Language models with image descriptors are strong few-shot video-language learners**. *arXiv preprint arXiv:2205.10747*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. **Chain of thought prompting elicits reasoning in large language models**. *arXiv preprint arXiv:2201.11903*.

- Makarius Wenzel, Lawrence C. Paulson, and Tobias Nipkow. 2008. The isabelle framework. In *TPHOLS*, volume 5170 of *LNCS*, pages 33–38. Springer.
- Jason Weston, Emily Dinan, and Alexander Miller. 2018. Retrieve and refine: Improved sequence generation models for dialogue. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*, pages 87–92, Brussels, Belgium. Association for Computational Linguistics.
- Martin White, Michele Tufano, Matias Martinez, Monperrus Martin, and Denys Poshyvanyk. 2019. Sorting and transforming program repair ingredients via deep learning code similarities. *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 479–490.
- Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss. 2018. Incorporating background knowledge into video description generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3992–4001, Brussels, Belgium. Association for Computational Linguistics.
- Sixing Wu, Ying Li, Dawei Zhang, and Zhonghai Wu. 2020a. Improving knowledge-aware dialogue response generation by using human-written prototype dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1402–1411, Online. Association for Computational Linguistics.
- Sixing Wu, Ying Li, Dawei Zhang, Yang Zhou, and Zhonghai Wu. 2020b. Diverse and informative dialogue generation with context-specific commonsense knowledge awareness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5811–5820, Online. Association for Computational Linguistics.
- Xian Wu, Shuxin Yang, Zhaopeng Qiu, Shen Ge, Yangtian Yan, Xingwang Wu, Yefeng Zheng, S. Kevin Zhou, and Li Xiao. 2022a. DeltaNet: Conditional medical report generation for COVID-19 diagnosis. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2952–2961, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yu Wu, Furu Wei, Shaohan Huang, Yunli Wang, Zhoujun Li, and Ming Zhou. 2019. Response generation by context-aware prototype editing. In *AAAI*, volume 33, pages 7281–7288.
- Yuhuai Wu, Albert Qiaochu Jiang, Wenda Li, Markus Norman Rabe, Charles E Staats, Mateja Jamnik, and Christian Szegedy. 2022b. Autoformalization with large language models. In *NeurIPS*.
- Xin Xia, Lingfeng Bao, David Lo, Pavneet Singh Kochhar, Ahmed E. Hassan, and Zhenchang Xing. 2017. What do developers search for on the web? *Empir. Softw. Eng.*, 22(6):3149–3185.
- Jia Xin, Wang Hao, Yin Dawei, and Wu Yunfang. 2021. Enhancing question generation with commonsense knowledge. In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 976–987, Huhhot, China. Chinese Information Processing Society of China.
- Jitao Xu, Josep-Maria Crego, and Jean Senellart. 2020. Boosting neural machine translation with similar translations. In *Annual Meeting of the Association for Computational Linguistics*, pages 1570–1579. Association for Computational Linguistics.
- Ran Xu, Caiming Xiong, Wei Chen, and Jason Corso. 2015. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *AAAI*, volume 29.
- Yichong Xu, Chenguang Zhu, Ruochen Xu, Yang Liu, Michael Zeng, and Xuedong Huang. 2021. Fusing context into knowledge graph for commonsense question answering. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1201–1207, Online. Association for Computational Linguistics.
- Fabian Yamaguchi, Nico Golde, Daniel Arp, and Konrad Rieck. 2014. Modeling and discovering vulnerabilities with code property graphs. In *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*, pages 590–604. IEEE Computer Society.
- Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, Josef Sivic, and Cordelia Schmid. 2023a. Vid2seq: Large-scale pretraining of a visual language model for dense video captioning. *CoRR*, abs/2302.14115.
- Xingyi Yang, Muchao Ye, Quanzeng You, and Fenglong Ma. 2021. Writing by memorizing: Hierarchical retrieval-based medical report generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5000–5009, Online. Association for Computational Linguistics.
- Yue Yang, Wenlin Yao, Hongming Zhang, Xiaoyang Wang, Dong Yu, and Jianshu Chen. 2022a. Z-LaVI: Zero-shot language solver fueled by visual imagination. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1186–1203, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2022b. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *AAAI*, volume 36, pages 3081–3089.
- Zhicheng Yang, Jinghui Qin, Jiaqi Chen, Liang Lin, and Xiaodan Liang. 2022c. LogicSolver: Towards interpretable math word problem solving with logical prompt-enhanced learning. In *Findings of the*

- Association for Computational Linguistics: EMNLP 2022*, pages 1–13, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zhuolin Yang, Wei Ping, Zihan Liu, Vijay Korthikanti, Weili Nie, De-An Huang, Linxi Fan, Zhiding Yu, Shiyi Lan, Bo Li, Ming-Yu Liu, Yuke Zhu, Mohammad Shoeybi, Bryan Catanzaro, Chaowei Xiao, and Anima Anandkumar. 2023b. Re-vilm: Retrieval-augmented visual language model for zero and few-shot image captioning. *CoRR*, abs/2302.04858.
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2022. Retrieval-augmented multimodal language modeling. *CoRR*, abs/2211.12561.
- Xi Ye and Greg Durrett. 2022. [The unreliability of explanations in few-shot prompting for textual reasoning](#). In *Advances in Neural Information Processing Systems*.
- Yunhu Ye, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. 2023. Large language models are versatile decomposers: Decompose evidence and questions for table-based reasoning. *arXiv preprint arXiv:2301.13808*.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. 2022. Scaling autoregressive models for content-rich text-to-image generation. *CoRR*, abs/2206.10789.
- Xiaoqing Yu and Anxiao Jiang. 2021. [Expanding, retrieving and infilling: Diversifying cross-domain question generation with flexible templates](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3202–3212, Online. Association for Computational Linguistics.
- Zheng Yuan, Qiao Jin, Chuanqi Tan, Zhengyun Zhao, Hongyi Yuan, Fei Huang, and Songfang Huang. 2023. RAMM: retrieval-augmented biomedical visual question answering with multi-modal pre-training. *CoRR*, abs/2303.00534.
- Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, et al. 2022. Socratic models: Composing zero-shot multimodal reasoning with language. *arXiv preprint arXiv:2204.00598*.
- Hang Zhang, Xin Li, and Lidong Bing. 2023a. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, and Xudong Liu. 2020. [Retrieval-based neural source code summarization](#). In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering, ICSE '20*, page 1385–1397, New York, NY, USA. Association for Computing Machinery.
- Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2018. [Guiding neural machine translation with retrieved translation pieces](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1325–1335, New Orleans, Louisiana. Association for Computational Linguistics.
- Jun Zhang, Yan Yang, Chencai Chen, Liang He, and Zhou Yu. 2021. [KERS: A knowledge-enhanced framework for recommendation dialog systems with multiple subgoals](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1092–1101, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023b. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Jinming Zhao, Gholamreza Haffari, and Ehsan Shareghi. 2023a. [Generating synthetic speech from SpokenVocab for speech translation](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1975–1981, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ruochen Zhao, Xingxuan Li, Yew Ken Chia, Bosheng Ding, and Lidong Bing. 2023b. [Can chatgpt-like generative models guarantee factual accuracy? on the mistakes of new generation search engines](#).
- Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023c. [Verify-and-edit: A knowledge-enhanced chain-of-thought framework](#).
- Luwei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *AAAI*, volume 34, pages 13041–13049.
- Mingyang Zhou, Grace Luo, Anna Rohrbach, and Zhou Yu. 2022a. [Focus! relevant and sufficient context selection for news image captioning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6078–6088, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Shuyan Zhou, Uri Alon, Frank F Xu, Zhiruo Wang, Zhengbao Jiang, and Graham Neubig. 2022b. Docprompting: Generating code by retrieving the docs. *arXiv preprint arXiv:2207.05987*.
- Yucheng Zhou and Guodong Long. 2023. [Style-aware contrastive learning for multi-style image captioning](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2257–2267, Dubrovnik, Croatia. Association for Computational Linguistics.

Jie Zhu, Junhui Li, Muhua Zhu, Longhua Qian, Min Zhang, and Guodong Zhou. 2019. [Modeling graph structure in transformer for better AMR-to-text generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5459–5468, Hong Kong, China. Association for Computational Linguistics.

Wanrong Zhu, An Yan, Yujie Lu, Wenda Xu, Xin Wang, Miguel Eckstein, and William Yang Wang. 2023. [Visualize before you write: Imagination-guided open-ended text generation](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 78–92, Dubrovnik, Croatia. Association for Computational Linguistics.

Modality	ACL	Google	Total analyzed
Image	(67) 17	6	23
Code	(177) 9	24	33
Structured	(108) 44	11	55
Audio	(17) 6	14	20
Video	(22) 7	7	14
Total	(291) 83	62	145

Table 1: Paper statistics. Number in parenthesis is the number before manual filtering. “Google” represents searching on google scholar and manually filtering. “Total analyzed” represents the number of total papers after manual filtering




Figure 1: Paper trend analysis

A Appendix

A.1 Search Criteria and Results

For searching the ACL anthology articles, we use a keyword search over titles and abstracts. We strictly enforce the keyword “retriev”. Then, we enforce either “generat” or “ground” to appear. For each modality, we then add modality-specific keywords: “image” for the image modality, “code” for the code modality, any one from “structured knowledge/table/database/knowledge graph” for the structured knowledge modality, any one from “audio/speech” for the audio modality, and “video” for the video modality.

For searching on Google Scholar, we add the keyword “language models” to select more NLP-related articles. We then perform manual filtering on the top 3 pages of returned results.

The number of retrieved and analyzed research papers can be found in Table 1.

A trend analysis of how the number of papers change across time is shown in Figure 1 We could observe that the domain of multimodal retrieval-augmented generation has indeed developed a lot

recently, with peaks reached around end of 2022. The observation is consistent with our hypothesis that multimodal RAG is especially important and helpful in the age of large-scale general-purpose models.