



Information Retrieval Meets Large Language Models: A Strategic Report from Chinese IR Community

Qingyao AI^a, Ting BAI^b, Zhao CAO^c, Yi CHANG^d, Jiawei CHEN(✉)^e, Zhumin CHEN^f, Zhiyong CHENG^g, Shoubin DONG^h, Zhicheng DOUⁱ, Fuli FENG^j, Shen GAO^j, Jiafeng GUO^k, Xiangnan HE(✉)^j, Yanyan LAN^a, Chenliang LI^l, Yiqun LIU^a, Ziyu LYU^m, Weizhi MA^a, Jun MA^f, Zhaochun REN^j, Pengjie REN^j, Zhiqiang WANGⁿ, Mingwen WANG^o, Ji-Rong WENⁱ, Le WU^p, Xin XIN^f, Jun XUⁱ, Dawei YIN^q, Peng ZHANG(✉)^r, Fan ZHANG^l, Weinan ZHANG^s, Min ZHANG^a, Xiaofei ZHU^t

^aTsinghua University, ^bBeijing University of Posts and Telecommunications, ^cHuawei Technologies Ltd. Co., ^dJilin University, ^eZhejiang University, ^fShandong University, ^gShandong Artificial Intelligence Institute, ^hSouth China University of Technology, ⁱRenmin University of China, ^jUniversity of Science and Technology of China, ^kInstitute of Computing Technology, Chinese Academy of Sciences, ^lWuhan University, ^mShenzhen Institute of Advanced Technology, Chinese Academy of Sciences, ⁿShanxi University, ^oJiangxi Normal University, ^pHefei University of Technology, ^qBaidu Inc., ^rTianjin University, ^sShanghai Jiao Tong University, ^tChongqing University of Technology

Abstract

The research field of Information Retrieval (IR) has evolved significantly, expanding beyond traditional search to meet diverse user information needs. Recently, Large Language Models (LLMs) have demonstrated exceptional capabilities in text understanding, generation, and knowledge inference, opening up exciting avenues for IR research. LLMs not only facilitate generative retrieval but also offer improved solutions for user understanding, model evaluation, and user-system interactions. More importantly, the synergistic relationship among IR models, LLMs, and humans forms a new technical paradigm that is more powerful for information seeking. IR models provide real-time and relevant information, LLMs contribute internal knowledge, and humans play a central role of demanders and evaluators to the reliability of information services. Nevertheless, significant challenges exist, including computational costs, credibility concerns, domain-specific limitations, and ethical considerations. To thoroughly discuss the transformative impact of LLMs on IR research, the Chinese IR community conducted a strategic workshop in April 2023, yielding valuable insights. This paper provides a summary of the workshop's outcomes, including the rethinking of IR's core values, the mutual enhancement of LLMs and IR, the proposal of a novel IR technical paradigm, and open challenges.

© 2011 Published by Elsevier Ltd.

Keywords: Information Retrieval, Language Language Models, Recommendation System

1. Introduction

In the past few decades, Information Retrieval (IR) has experienced significant growth and development in both industry and academia. In early stage, IR research mainly focused on search, which aimed to assist users in finding

¹Email: sleepyhunt@zju.edu.cn; xiangnanhe@gmail.com; pzhang@tju.edu.cn

relevant information [1]. In recent years, the scope of IR research has expanded to encompass a wide range of online applications and scenarios [2]. This diversification is evident in the rise of recommendation systems as a prominent research area, as observed in flagship conferences ACM SIGIR from 2018 to 2023. Additionally, exciting research directions such as conversational systems, user modeling, and knowledge extraction, have also emerged [3, 4]. These advancements reflect an evolution in the core value of IR, moving beyond merely retrieving relevant documents to meeting the information needs of users [5].

Large Language Models (LLMs) have demonstrated remarkable capabilities in various aspects of text understanding, generation, knowledge inference, and compositional generalization [6, 7]. As a result, LLMs have the potential to open up new research directions in the field of IR. These include, but are not limited to:

- Enabling IR systems to generate content directly that satisfies user information needs, known as generative retrieval [8] and generative recommendation [9]. An example of this is the New Bing¹, which allows for content generation to meet user queries.
- Enhancing the understanding of user intents and behaviors by incorporating rich contextual information into IR system.
- Creating opportunities for the development of superior indexing systems that can handle dynamic, semantically-aware, and multi-modal data.
- Providing better methods for evaluating model accuracy and interpretability in IR tasks.
- Facilitating enhanced interactive experiences between users and IR systems. These advancements highlight the potential of LLMs to revolutionize various aspects of IR research and practice.

Reciprocally, IR technology plays a crucial role in supporting the development of LLMs [10, 11]. The interaction among humans, IR models, and LLMs forms a new triangular IR paradigm, as depicted in Figure 1. In this paradigm, the IR model acts as the means to acquire external knowledge, offering real-time, up-to-date, and relevant information to both LLMs and humans. LLMs contribute precise internal knowledge and information, leveraging their powerful reasoning abilities to provide high-quality responses. Humans, in their role as demanders and evaluators, play a central role in the retrieval process. It is important to emphasize that LLMs, without the inclusion of IR models, possess limitations in terms of short-term memory and reasoning. These limitations could result in the lack of effective integration of new and existing knowledge, hindering LLMs' ability to provide the information that humans need. The IR model, functioning as a “knowledge base” equipped with matching and retrieval capabilities, addresses these limitations by compensating for the lack of factual consistency and long-term memory in LLMs. This combination of “knowledge+reasoning” establishes a dual-wheel drive, enabling a more intelligent and reliable information service system for humans [12, 10].

While LLMs for IR hold promise, they also face significant challenges that cannot be ignored:

- High computational costs: LLMs for IR require substantial computational resources, which can limit their use and research by small and medium-sized companies and institutions. Furthermore, deploying models on terminal devices and ensuring real-time performance can be challenging [13].
- Credibility concerns: The credibility of LLMs for IR is relatively low, as they have been found to provide misleading explanations, incorrect answers, and unreliable information sources. This can undermine people's trust in LLMs [14].
- Performance limitations in specific fields: LLMs may struggle to perform well in specific domains due to the lack of high-quality datasets, differences in data representation, and other related issues. The limited domain-specific knowledge hampers the rapid implementation and application of LLMs [15, 16].

¹<https://www.bing.com/new>

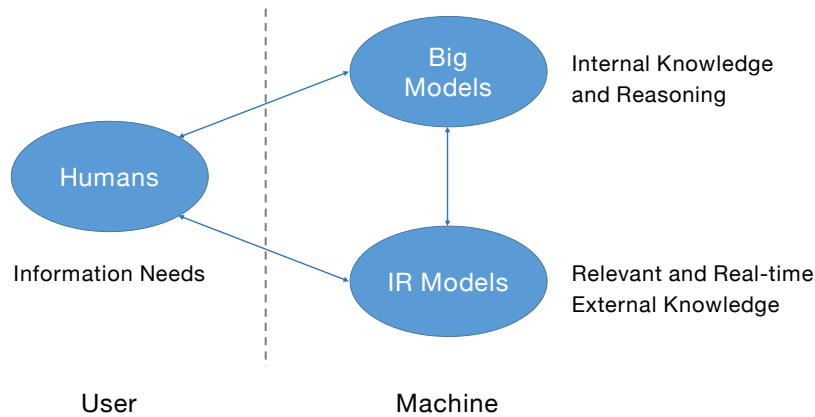


Figure 1. Our proposed new IR technical paradigm introduces three key elements: Humans, IR models, and LLMs. The synergistic relationship among them not only facilitates mutual enhancement but also enables the fulfillment of new tasks as an organic whole.

- **Ethical and moral considerations:** The emergence of generative IR systems raises the requirements for ethical and moral evaluation standards. Ensuring fairness, impartiality, and ethics in the generated content is a significant challenge. It is inherently difficult to guarantee that LLMs meet these regulatory and ethical requirements [17, 18].

Given the flourishing research in the IR field and the promising opportunities brought about by LLMs, it is a favorable moment to envision the future advancements in IR. To this end, the Chinese IR community organized a strategic workshop on April 14th and 15th, 2023, aiming to explore future opportunities and challenges. This paper presents the key findings and conclusions drawn from the workshop, including the rethinking of the core values of IR in Section 2, a discussion of how LLMs and IR can mutually enhance each other in Section 3 and Section 4, the new IR paradigm of fusing humans, LLMs, and IR in Section 5, and a discussion of open challenges and future directions in Section 6.

2. Rethinking Core Values of IR

Before delving into the effects of LLMs on IR research, it is crucial to rethink the IR discipline itself. This contemplation leads us to ponder the following questions: Firstly, what is the fundamental contribution of IR, as a scientific discipline, to other domains? Secondly, considering its extensive development over the past decades, what are the boundaries and extensions of IR? These questions are intended to promote comprehension of the fundamental principles of IR, and foster the integration of LLMs into IR while preserving its essence.

2.1. Scientific Connotation of IR

Classical IR concerns the retrieval of the information that satisfies a user's information need from the target corpus. Therefore, user, ranking, and corpus are the fundamental concepts in the IR discipline. In the light of this, it is important to re-examine the scientific connotation of IR from the three aspects.

User Perspective. Users are the core part of IR [19, 20], since an IR system is supposed to serve users by satisfying their information needs. Therefore, user understanding is a fundamental problem in IR, and the community has made decades-long progress in user understanding. The studies of user understanding in IR mainly concentrate on two sides, respectively user intent understanding and user behavior modeling.

- User intent understanding plays a central role in IR. Commercial search engines like Google² and Bing³, have

²<https://www.google.com>

³<https://www.bing.com>

evolved from keyword semantic matching to user intent optimization in recent decades. Understanding user intent is the first step for better user experiences, which has been broadly studied in various IR applications, e.g., web search [21, 22], product search in e-commerce [23], community question answering [24], etc.

- User behavior modeling is to comprehensively understand and model user behaviors when interacting with IR systems. IR community has studied diverse and multi-dimensional user behaviors including user interaction behaviors [25, 26], sequential behaviors [27, 28], mobility behaviors [29, 30] and social behaviors [31]. User behavior modeling helps to inject personalization into IR systems.

The techniques of user understanding in IR have provided valuable insights for other research domains. In natural language processing (NLP), user intent understanding is the preliminary step to make further actions in task-oriented dialogue systems [32]. User profiling and personalization are also utilized to develop persona-based dialogue systems [33, 34]. In social computing, user modeling is an important area in the design of personalized incentive mechanisms for encouraging participation [35].

Ranking Perspective. Ranking lies at the core part of IR research by providing an ordering of items towards user’s information need [36, 37]. The formats of items are various, including but not limited to textual documents, images, and videos. The problem is usually formulated as ranking a collection of items that match the given query or context, according to some criterion like relevance and recency. Most IR systems are supported by a ranking module.

Different techniques have been used for constructing ranking models, from traditional ranking models to machine learning methods. The traditional ranking models include vector space models and probabilistic models (e.g. BM25 [38]). The machine learning methods include the earliest learning to rank models (e.g. Rankboost [39] and LambdaMart [40]) and the recent neural ranking models such as Deep Structured Semantic Models (DSSM) [41] and DeepRank [42]. Ranking models have been extensively exploited in many IR applications such as ad-hoc retrieval [43, 44, 45], recommender systems [46, 47], community question answering [48, 49], etc.

It’s worth noting that ranking is not restricted to IR alone, as it also has applications in a diverse range of other research fields. For example, ChatGPT [50] collects comparison samples and utilizes a ranking labeler to train the reward model when aligning with user intents. In computational biology, learning-to-rank algorithms have been exploited to rank the candidate 3-D structures in protein structure prediction [51].

Corpus Perspective. For the corpus perspective, the connotation of IR lies in conducting effective search efficiently from a large volume of multi-modal information corpus [1]. Such connotation mainly lies in the following two folds:

- Effective corpus representation. Conventional IR systems represent the documents with bag-of-words or TF-IDF methods [1]. Deep learning methods (e.g., dense retrieval) empowers the IR system to represent multi-modal corpus such as images and videos with dense latent vectors, which has largely improved the retrieval results and broadened the application scope of IR [52, 53]. Recently, contrastive learning and self-supervised learning further improve the learning of corpus representation [54].
- Efficient retrieval infrastructure. The infrastructure of IR involves fundamental facilities to construct an IR system, including but not limited to building dictionaries, indexing, scoring, distributed search, and caching [19]. Such infrastructure determines how to organize the information source and thus largely affect the efficiency of IR [19]. LLMs and generative IR [55] shed new lights to reform the retrieval infrastructure.

2.2. Boundaries and Extensions of IR

Although the growth of Web and mobile technology has greatly advanced the development of IR, it also results in some limitations. A common misconception is that IR simply involves ranking a set of objects in cyberspace, which is an inaccurate and narrow view. As such, it is crucial to reconsider the boundaries and extensions of IR from various perspectives, including the input, process, output, and environment.

Ranking or Generation?. The primary focus of classical IR pertains to the ordering of items based on the information need, such as the query in search engines or the context of recent interactions in recommender systems. The advancement of large-scale generative models such as GPT [56] and Diffusion Models [57] has expanded the content space, allowing exploration beyond existing sets of content through generation. Integrating the generation channel for information seeking opens up new paradigms in IR systems, such as generative search and recommendation [9]. However, this extension presents challenges to building IR systems in infrastructure, algorithm, evaluation, *etc.* For instance, it calls for the infrastructure to accommodate generative models at a large scale and distribute the content generated on the fly in an efficient manner, especially for multi-media content such as micro-videos. Additionally, it calls for algorithm innovation on the alignment of generative models to IR tasks, the aggregation of ranking and generation results, and the joint optimization of ranking and generation models. Furthermore, evaluation techniques must also be extended to accommodate unlimited content and emphasize new perspectives such as fidelity checking.

For Human or For Machine?. IR systems have been designed to assist *human* users to access information resources. However, with the increasing prevalence of AI models, particularly LLMs, the user base of IR systems could potentially be expanded to include intelligent *machines*. This has led to the emergence of a growing body of research referred to as IR-augmented methods or retrieval-enhanced machine learning (REML) [58] in various communities such as computer vision[59], natural language processing [60], and machine learning[61]. From a conceptual perspective, IR systems can be regarded as a memory and access system for external data, information, or knowledge, capable of serving as a generic supporting technique for either humans or machines. Nevertheless, the shift from human to machine poses a multitude of research challenges throughout the entire process of the IR system, including indexing, representation, retrieval, ranking, and feedback. Additionally, there are many research issues to be addressed regarding the learning, evaluation, and deployment of IR technology in conjunction with AI models. Despite these challenges, it is encouraging to anticipate that IR could become a ubiquitous component of the AI paradigm in the future, providing valuable support to society at large.

For Finding or For Decision?. IR systems have long benefited people in finding desired information. For the extension of IR, the system should also support complicated, explainable, and long-term information seeking to help users with reasonable decision-making suggestions [62]. To accomplish this goal, a retrieval system should be aware of the user context in which the information need is to be placed [63]. For complex decision-making tasks, the system is expected to scaffold the task step-by-step [64]. The system should also provide transparent retrieval process and explainable results to provide reasonable decision-making supports [65]. Besides, understanding how people make decisions is also important. Collaboration with other communities will likely be necessary to make progress in decision understanding. These communities may include human-computer interaction, behavioural economics, psychology, cognitive science, as well as specific domain communities like the clinical communities [66, 67]. Finally, placing IR into LLMs to provide AI-Generated Actions (AIGA) is also an emerging topic for the extension of IR [68].

Going beyond Cyberspace?. IR systems have conventionally emphasized digital content within the realm of cyberspace. However, advancements in human-computer interaction [69], unmanned vehicles [70], and robotics [71] enable us to envision the retrieval of contents in the physical world. For instance, the extended IR system may explore information in the physical world such as the wind speed around and the source of noise. Further extensions may incorporate the search and delivery of physical items. These extensions build on the aforementioned ranking-generation hybrids, human-machine hybrids, and finding-decision hybrids. Moreover, new concepts around interaction interfaces, system architecture, and feedback mechanisms will arise. Towards these ends, we may open up more cross-domain research directions and entangle IR with various emerging techniques in other fields such as Metaverse [72] and Embodied Artificial Intelligence [73]. As a result, there will be many research opportunities regarding the development, operation, maintenance, and regulation of such IR systems, increasing the value of the IR industry remarkably.

3. Large Language Models for IR

IR systems have long been divided into five major components: user modeling, indexing, matching/ranking, evaluation, and user interaction. By pre-training on large-scale corpora and fine-tuning to follow human instructions,

LLMs have demonstrated superior capability in language understanding, generation, interaction, and knowledge reasoning [74]. Therefore, it is reasonable to anticipate that LLMs will considerably augment the major IR components from various perspectives. In this section, we discuss the possibilities of LLMs for each major IR component.

3.1. User Modeling

User modeling aims to accurately represent users and their information needs by understanding their characteristics, preferences, and behaviors [75, 76, 77, 78]. In view of the acknowledged capabilities of LLMs, we enumerate several potential directions that LLMs can enhance user modeling.

- **Language Understanding.** Most basically, LLMs can enhance the understanding of user queries, enabling more accurate analysis and leading to more relevant search results.
- **Behavior Understanding.** LLMs allow the analyses of user behaviors in the semantic level of item content, offering a comprehensive understanding of preferences and behaviors. By analyzing data like click-stream, search log, and interaction history, models can identify patterns and relationships, resulting in more accurate user models.
- **Personalization.** LLMs can build comprehensive user models, incorporating various characteristics and preferences. By analyzing social media activity, online behavior, and integrating with IoT devices, models can consider the factors like physical environment and emotional states, leading to more personalized recommendations.
- **Conversational Interfaces.** LLMs facilitate more natural and seamless interactions, generating context-aware responses. Conversational interfaces can incorporate sentiment analysis and emotional response generation, thereby enhancing personalized and engaging user experiences.
- **Hybrid Modeling.** Combining LLMs with rule-based or collaborative filtering models can yield stronger hybrid models. Such hybrid modeling unifies the strengths of different approaches, improving the personalization and accuracy of search and recommendation [79, 80].

The advancement of LLMs is expected to significantly enhance the user modeling in IR systems. Nonetheless, certain challenges, such as data privacy, biases in data and models, and the need for extensive training data, need to be tackled. Addressing these challenges is imperative to fully exploit the potential benefits of LLMs for user modeling in IR, while mitigating any potential negative effects.

3.2. Indexing

The emergence of LLMs provides the basis for generative retrieval, a new retrieval paradigm where the indexing component is dramatically changed. Recently, dense retrieval has been extensively studied and the approaches based on the standard MIPS index and nearest neighbor search are common [52]. Given the success of Transformers as a good associative memory store or search index, Tay et al. [55] propose a novel architecture called differentiable search index (DSI) where the index is stored in the model parameters. DSI uses LLMs to directly learn the mapping of queries to relevant document IDs. It enables the model's internal memory to act as an index, thus greatly simplifying the entire retrieval process. Inspired by the work of DSI, we identify some directions that LLMs will change the technology of indexing:

- **From Static to Dynamic.** Indexing systems based on LLMs need be dynamic, as all corpus information is encoded within the LLM parameters. Incremental index updates are a specific instance of model updates, as noted in the study by [81].
- **From Keyword-based to Semantics-oriented.** Indexing systems based on LLMs should be semantic. Thanks to the powerful contextual modeling capabilities of LLMs, indexing systems based on LLMs are able to find the documents that are related to the query in a more nuanced way.
- **From Uni-modal to Multi-modal.** Indexing systems utilizing LLMs have the potential to be multi-modal. The development of multi-modal LLMs will facilitate the indexing systems capable of indexing various modalities of data in a unified manner, including but not limited to texts, images, and videos.

3.3. Matching/Ranking

LLMs have demonstrated remarkable capability to understand and rank complex content, including both single-modal and multi-modal data. In this part, we focus on two topics: (1) If generative models can already provide exact answers, is ranking still necessary? (2) What are the future research directions in the ranking, and what problems remain to be solved?

While generative LLMs can provide coherent answers to user queries, users still need to know which document, image, or webpage is most relevant to their query, so as to verify the answer generated by the LLM. The ranking results from a retrieval system can improve the interpretability of LLM, and provide the trustworthy information to support the answer generation. The retrieval system typically acts as a plugin to LLMs [82, 83], providing knowledge acquisition abilities.

In future research, LLMs offer many interesting directions to explore in ranking. Under the generative paradigm, we should prioritize ranking and improve the ranking of results to ensure a good user experience and high satisfaction. This is also a more essential goal of ranking in IR. When ranking search results, we should focus on the integrity of the returned results. Instead of simply returning a ranking list based on relevance, the model should return the information that is more integrated and relevant to users' real needs.

In terms of ranking evaluation, existing methods are designed for the form of returning a list of documents, which is no longer suitable under the generative paradigm. There are promising directions for future research in the field of ranking. An example is to leverage the LLMs as the human simulator to measure the user satisfaction and experience of a ranking method [84].

3.4. Evaluation

For validate the effectiveness of LLM-enhanced IR approaches, it is essential to develop proper metrics and datasets. Conventional IR metrics, such as Precision and Recall, Mean Reciprocal Rank (MRR, [85]), Mean Average Precision (MAP, [86]), Normalized discounted cumulative gain (nDCG, [87, 88]), still play a critical role in IR model evaluation. However, with the wide application of LLMs in IR, tailored evaluation strategies become essential to justify the effectiveness of LLMs. To this end, some characteristics needed to be emphasized including:

- **Robustness.** Many models are sensitive to distributional differences between training and test data [89, 90]. It is eager to see the generalization of LLM-enhanced IR models in out-of-distribution (OOD) scenarios.
- **Interpretability.** Neural IR models rely on dense document representations, and suffer from poor interpretability, which is different from previous sparse IR models (e.g., BM25) with explicit term matching [91].
- **Efficiency.** LLMs may require additional training or fine-tuning to adapt to IR tasks, and storage is also a severe bottleneck in the training scenarios [92, 93].
- **Reliability.** LLMs are capable of directly retrieving knowledge from the model itself and generating results for users. However, they are vulnerable to adversarial inputs and some small character changes would negatively affect its reliability [94].

3.5. Interaction

The interaction between users and traditional search engines typically involves three steps [95]. First, a user submits a query generally expressed as keywords. Then, the search engine processes the query, retrieves its index and presents the results in the form of a list. Finally, the user browses the results and clicks on the most relevant links to the query.

With the development of conversational technology, conversational search has become a new retrieval paradigm [96, 97]. The retrieval process has become an iterative form of multiple rounds of dialogue. A user initiates the conversation by asking a question in natural language. The search engine processes the user's input, retrieves its index and generates a response. Then, the user may provide additional information or ask follow-up questions, and the search engine continues to process the user's input and generate responses until the user is satisfied.

LLMs [7] have the potential to overturn the interaction between users and IR systems. For example, search engines powered by LLMs can serve as an AI assistant, such as Windows Copilot, Bing Chat. It can generate responses based

on the conversation’s context and the user’s needs in real-time, without relying on pre-programmed responses. As long as the user inputs a question, it can help retrieve information, search knowledge, use Apps, and call plugins in the simplest way [98]. The IR system powered by LLMs possesses the capability to comprehend and address complex queries, thereby enhancing the interaction process in terms of intuitiveness, personalization, efficiency, effectiveness, responsiveness, and friendliness.

4. IR for Large Language Models

LLMs exhibit certain inherent limitations that are commonly encountered in generative language models. Firstly, they may occasionally generate erroneous or nonsensical responses, a phenomenon referred to as “hallucinations”⁴. Secondly, they are trained on fixed corpora, which restricts their ability to answer questions that require newly emerged knowledge or information after the training date [99, 100]; the stored knowledge within the model can be inevitably incomplete, outdated, or even incorrect [93]. Thirdly, they can hardly respond to “private” questions that require access to confidential data sources, which are not available during the training of LLMs.

Utilizing the retrieval capabilities of IR models presents a viable approach for addressing the above limitations of LLMs. By incorporating retrieval, we can leverage relevant knowledge from external knowledge bases during the generation process, thereby reducing the occurrence of hallucinations. Retrieval models possess strong ability to access up-to-date information, enabling LLMs to respond with fresh and relevant information. Moreover, the combination of retrieval and LLMs is particularly valuable in data-sensitive scenarios, where internal data cannot be utilized for language model training, factual accuracy is of utmost importance, or the retrieval pool may vary over time [101, 102].

In this section, we discuss three directions that augment LLMs with retrieval in different phrases: pre-training, fine-tuning, and leveraging black-box LLMs such as ChatGPT that only offer APIs.

4.1. Pre-training LLMs with Retrieval

There is limited work on pre-training LLMs that incorporate a retrieval module. One notable example is Atlas [103], which pre-trains an encoder-decoder T5 with an extra retrieval module and demonstrates its ability on knowledge-intensive tasks with very few training samples. REALM enhances encoder-only language model training by incorporating a neural knowledge retriever that extracts information from a text-based knowledge database [104]. RETRO is a retrieval-augmented decoder-only language model, where a chunked cross-attention module is employed to aggregate retrieved text from a retrieval pool containing trillions of tokens [105].

In a recent study, based on the auto-regressive language model RETRO, Wang et al. [106] studies the question whether it is useful to pre-train LLMs with retrieval. As a result, the pre-trained retrieval-augmented language model RETRO outperforms the vanilla GPT on text generation tasks with higher factual accuracy, and on knowledge-intensive tasks. In addition, in open-domain question answering tasks, RETRO largely outperforms GPT which just incorporates retrieval at the fine-tuning stage. Note that there is a trade-off between the computation overhead from using retrieval and the performance. Wang et al. [106] suggest a flexible implementation that specifies the number of tokens to generate with the current retrieval result before the next retrieval.

In addition to incorporating a retrieval module into the pre-training of LLMs, it is also valuable to investigate how retrieval functionalities can enhance the pre-training process itself. Generally, the retrieved evidence serves as the *privileged information* within the input context. Such privileged information is often costly or impractical to obtain during the inference stage, though it is available during the training phase [107]. By having both the retrieved evidence and the processed output answer (or subsequent tokens) accessible during training, it is possible to pre-train (or simply fine-tune) LLMs to improve their performance on knowledge-intensive tasks [108, 98].

4.2. Fine-Tuning LLMs with Retrieval Adapters

Pre-training or fine-tuning the entire LLM with retrieval could enhance its retrieval capabilities. However, the considerable cost makes it prohibitive to use. Considering that LLMs serve as a foundation for downstream tasks,

⁴Introducing ChatGPT, <https://openai.com/blog/chatgpt>

updating all parameters should be approached cautiously and infrequently. Therefore, an alternative is to fine-tune LLMs with search adapters, which is more efficient and cost-effective. Adapters are plug-and-play modules for LLM with only a small number of parameters. They enable LLMs to acquire task-specific capabilities without affecting the original parameters, while still achieving comparable or superior performance compared to updating the entire models.

There have been several studies on LLM adapters, which can be roughly classified into two categories: token-based methods and layer-based methods [109, 110]. The former seeks to insert task-related anchor tokens into the input sequence for fine-tuning [111, 112, 113]. For example, Prompt tuning prepends several additional task-specific tunable tokens into the input sequence [114]. The latter seeks to insert additional layers into the models and fine-tune these layers only [115, 116, 117]. For example, Hu et al. [118] propose LoRA which adds a trainable low-rank dense layer before the self-attention in transformers. LLMs with adapters have been demonstrated effective on some downstream tasks such as machine reading comprehension.

Nevertheless, the following research questions remain to be explored to empower LLMs with retrieval capabilities through adapters.

- What adapters are best suited for retrieval? What neural network modules are best suited for retrieval adapters?
- How to fine-tune the retrieval adapters? Which fine-tuning techniques can promote the retrieval capability in highest degree without hurting the inherent capabilities of LLM?
- Why do retrieval adapters work? How do we know the capability boundary of retrieval adapters, and what can we do to avoid potential failures?

Fine-tuning LLMs with retrieval adapters can provide benefits in various aspects. On one hand, it offers new opportunities for researchers with limited computation resources. Progress in scientific research cannot be achieved without the participation of researchers around the world. On the other hand, the lower cost of search adapters allows for wider applications, especially when data privacy is a top priority. Small-size institutions can have LLMs equipped with their own retrieval systems without exposing the raw data.

4.3. Augmenting Black-box LLMs with Retrieval

In many scenarios, LLMs can only be accessed through remote APIs, and the possibility of fine-tuning these models is restricted. The most prevalent approach to leveraging LLMs is treating them as black-box systems and using customized prompts that integrate evidence from external data sources. This method enables users to obtain desired outputs from the LLMs by providing tailored prompts that incorporate relevant external evidence.

There are some preliminary studies on this interesting topic. Shuster et al. [119] proposed a modular system SeeKeR which searches for and chooses knowledge with internet search as a module during language generation. Similarly, Komeili et al. [99] and Lazaridou et al. [100] also introduced approaches which condition on the results from internet search engines (such as Google Search) to generate a response. He et al. [93] proposed a post-processing approach named “rethinking with retrieval (RR)” to solve the same problem. RR uses the chain-of-thought (COT) prompting to generate multiple reasoning paths, then retrieves external evidences based on the steps in these paths. Experiments on three complex reasoning tasks and different datasets demonstrate that RR outperforms all baselines without additional pre-training or fine-tuning. Ram et al. [120] introduced an in-context Retrieval-Augmented Language Modeling (RALM) method that simply uses off-the-shelf general purpose retrievers to retrieve documents and prepends retrieved results to the input of language models. Peng et al. [121] proposed a system named LLM-Augmenter, which augments a fixed black-box LLM with a set of plug-and-play modules. Given a query, LLM-Augmenter first retrieves evidences from external data sources, then generates a prompt that contains the retrieved evidences for ChatGPT. Their experiments on the information seeking dialog and open-domain wiki question answering tasks show this method could significantly reduce the hallucination problem of LLMs.

Recently, OpenAI officially released the Retrieval Plugin⁵, which provides semantic search and retrieval functionality of personal or organizational documents. Relevant documents snippets can be retrieved from external sources, such as personal emails or internal organizational files, with this plugin.

⁵ChatGPT Retrieval Plugin, <https://github.com/openai/chatgpt-retrieval-plugin>

There are several key components in the pipeline of augmenting LLMs with retrieval, which are worth to study in the future.

- **Design of Retriever.** Intuitively, a better retriever could return results with higher quality, and thereby helps to generate better response. Should we use a dense retriever or a sparse retriever in generating grounding documents? Which kind of information to index, documents or passages?
- **Context Modeling.** How to generate an explicit keyword query or representation vector that can describe the current information need and retrieve useful results for language generation? ChatGPT-like LLMs are optimized for dialogue rather than search, hence how to derive search intent from the multi-turn chat history remains a challenge.
- **Selection of Grounding Documents.** In addition to pure relevance, other factors of diversity, informativeness, and freshness should also be considered when selecting a small set of documents.
- **Prompting Mechanism.** It remains a challenge to generate good prompts that can help improve generation quality with the given grounding documents.

5. LLMs + IR: New Paradigm and Framework

Traditional IR models are designed to meet human information needs through the interactions between users and systems. IR models do not preserve knowledge, instead, they search for external knowledge or information to meet the information needs efficiently. With the emergence of LLMs, we propose a new technical paradigm of IR: as shown in Figure 1, the key elements are LLMs, IR models and humans.

- LLMs provide valuable (internal) knowledge and information to meet human information needs, and their reasoning capacity makes it easier to provide high-quality responses.
- IR models that search for external knowledge are indispensable, which provide up-to-date and relevant information for both LLMs and humans. While LLMs may suffice in directly generating answers to meet human information needs for simple questions, the significance of IR models becomes more pronounced when dealing with complex and difficult problems.
- Humans raise information needs in the paradigm and are the tutor for both LLMs and IR models. They endow the system with human values and behavioral characteristics, making it serve users better.

The synergistic relationship among them not only facilitates mutual enhancement but also enables the fulfillment of new tasks as an organic whole.

5.1. Importance of the Three Modules

In accordance with the new IR technical paradigm, an important question arises: Are all three elements essential for this new paradigm? In the following subsection, we will provide our insights and responses.

5.1.1. Paradigm without LLMs

Without LLMs, the paradigm degrades to traditional IR, accompanied by inherent challenges:

- **Dependency on the Internet.** As IR models do not retain knowledge or information themselves, they rely on the Internet to acquire external knowledge, potentially limiting their applicability in certain scenarios.
- **Lacking reasoning ability.** Existing IR models mainly provide collected knowledge/information to fulfill human information needs, lacking the ability to assist users in comprehending the information. Better reasoning ability will deliver more user-friendly and valuable results for humans.

5.1.2. Paradigm without IR

In the era of LLMs, the presence of an IR system is critical as it addresses the limitations of generative language models. Without an IR system, LLMs encounter the following challenges:

- **Lacking factual consistency.** One inherent limitation of generative LLMs is the lack of factual consistency, resulting from their training data and the potential generation of false information. In contrast, the IR system frees from this issue by storing and offering factual information through keyword matching, thereby complementing the shortcomings of LLMs.
- **Lacking effective integration of new and existing knowledge.** Large Language models have the capacity to learn knowledge from large-scale data, such as commonsense information. However, relying solely on LLMs may neglect the crucial mechanism of effectively and efficiently integrating new and existing knowledge. In the human brain, memory involves a complex network of interconnected brain regions, including the hippocampus, prefrontal cortex, and other cortical areas. Retrieval processes are closely linked to memory consolidation and the reactivation of neural networks associated with previously encoded information. When we retrieve information from memory, this activity can strengthen the neural connections related to that information, making it more accessible for future retrieval and improving overall memory performance.

5.1.3. Paradigm without Human

The absence of human input and feedback hinders both LLMs and IR models from delivering personalized information services. The associated challenges are:

- **Identical information services.** Large-scale dialogue systems use a generation-based approach to provide information to users. The generated content is not personalized, which can cause the issues that the content does not match the user's real intention.
- **Uncontrollable value of society.** Without human feedback, LLMs are unable to account for user personality and values, resulting in information services that do not adequately align with societal values.

5.2. Summary

The rise of dialogue-based LLMs, exemplified by platforms like ChatGPT from OpenAI and New Bing from Microsoft, has sparked a trend that could potentially replace traditional IR systems. However, as previously discussed, in the era of LLMs, human information needs and feedback, along with the continued relevance of traditional IR models, remain vital components in the development of trustworthy information service systems.

6. Challenges and Future

While LLMs for IR hold promise, they also present numerous challenges and unanswered questions. In the final section of this article, we discuss some selected issues to outline future directions.

- **High Computational Costs.** The primary challenge in using LLMs is their high computational cost. This poses a significant barrier for small and medium-sized research laboratories and companies, hindering their integration of LLMs into daily workflows and products. Even large companies with ample computational resources face cost pressures when deploying LLMs for online search, recommendation, and advertisement services due to the immense volume of user requests. Common solutions include compressing LLMs, reducing their size from hundreds of billions to tens billions or even smaller, especially before online deployment. Additionally, efforts to develop more efficient and cost-effective hardware for training and inference are underway to address the cost challenge.
- **General-purpose v.s. Domain-specific.** LLMs have demonstrated impressive capabilities in general-purpose tasks like text generation and chatting, owing to their pre-training and fine-tuning on large-scale Internet corpora. However, it is widely recognized that LLMs face limitations when it comes to adapting to domain-specific tasks. On one hand, high-quality professional domain knowledge, which is often not abundantly available on

the Internet, makes it prohibitive to pre-train and fine-tune LLMs. On the other hand, domain-specific knowledge is not always expressed in natural language; it may be represented as semi-structured or structured tables, heuristic rules, equations, and more. Enabling LLMs to effectively handle domain-specific tasks is crucial not only for the specific domains themselves but also for enhancing the overall capabilities and applications of LLMs.

- **Trustworthiness.** There is a widely acknowledged concern that LLMs currently lack the ability to provide reliable and trustworthy answers to user queries. While LLMs can generate explanations and cite sources, it has been observed that a significant portion of these explanations and citations are illogical, inappropriate, or even fake. This poses a substantial risk in real-world search and recommendation scenarios, as generating misleading explanations, answers, and information sources can have detrimental effects on the community at large. To address this issue and enhance the trustworthiness of LLMs, it is crucial to enable LLMs to have a clear understanding of their knowledge and limitations. One potential solution is allowing LLMs to decline providing an answer when uncertain.
- **Controllable Generation.** Considering the public nature of search engines and recommendation systems, it is important to address regulatory and ethical considerations such as fairness, impartiality, and human values when presenting content to users. While LLMs demonstrate proficiency in generating text, they often lack a deep understanding of the meaning behind the generated words. Ensuring that the generated content meets the necessary regulatory and ethical requirements remains a significant challenge, and has no effective solutions for now.
- **High-quality Data:** High-quality data plays a vital role in the development and improvement of LLMs. The success of LLMs heavily relies on the continuous provision of human-labeled data. It is crucial that the labeled data not only meets a certain quantity threshold but also maintains high quality. Obtaining high-quality data in real-world application scenarios involves multiple steps such as data cleaning, data labeling, and data quality evaluation. Professional data annotation providers play a crucial role in supporting these processes. Additionally, it is essential to develop advanced, professional, and sustainable data annotation approaches to meet the growing demand for high-quality data in LLM applications.
- **Long-context Dependency.** Existing LLMs have limited capacity to handle long contexts, while IR tasks rely on long-term context to effectively capture and understand user intent. It is crucial to enable LLM-enhanced IR systems to model users' long-term intent that spans a large range of period.
- **Serving Time Requirements.** The latency in serving LLM results significantly lags behind the time requirements of information retrieval (IR) systems. This presents efficiency challenges when integrating LLMs into IR, thereby impacting the online user experience.
- **Presentation Format.** Traditional IR systems present ranked lists of content, while LLMs excel at generating new information. How to design a new presentation format that effectively fulfills user needs in LLM-enhanced IR remains an open question.
- **Integrating Structural Information.** LLMs primarily rely on textual sequence information, whereas IR systems require the integration of structural information such as user-item interactions and web linkage data. Effectively leveraging this structural information in LLM-based IR systems is an unresolved issue.
- **Balance between Generative and Retrieved Data.** LLMs leverage deep learning and reinforcement learning to generate content at scale, but their generated content may have limitations in terms of freshness and credibility. In contrast, retrieval can provide the content from the Web, ensuring the latest information. Balancing these two types of data in real-life applications is a significant challenge for improving overall performance. One approach is to refine user needs and classify them into different groups, allowing for appropriate data generation methods or balancing ratios. In addition, retrieval can be used to provide additional information to enhance content generation, or assist in screening and filtering the generated content for information grounding.

- **Content Quality and Credibility.** While LLMs are effective in generating content, they can also produce low-quality or even misinformation-filled content. The proliferation of such content on the Internet can disrupt the existing data ecology and impact applications like search engine and recommendation system. Traditional quality assessment techniques like PageRank may not be effective in this context. It is difficult for existing techniques to identify low-quality or misleading content. In the era of AI-generated content, new mechanisms are needed to assess data quality and distinguish between generated and reliable content. One approach is manual or community review to ensure content accuracy, but it is time-consuming and does not scale well. Another approach is leveraging LLM-powered machine learning techniques to train models that can recognize AI-generated content and evaluate its quality. These models can then be applied to select, label, or filter content for different applications.
- **Content Creation Environment.** The proliferation of generated content introduces challenges for content creators, and may even reshape the content ecosystem. The presence of generated content intensifies competition within the content market, pushing content creators to continuously enhance their writing quality, foster innovation, and proactively adapt to industry changes. Moreover, it may influence users' perception of value and demand for content, prompting content creators to constantly adjust their writing approach and strategies. Despite the challenges, LLMs also provide new opportunities for content creators to collaborate and develop effective platforms that facilitate more productive and higher-quality content generation.

References

- [1] M. Kobayashi, K. Takeda, Information retrieval on the web, *ACM computing surveys (CSUR)* 32 (2) (2000) 144–173.
- [2] M. Manoj, J. Elizabeth, Information retrieval on internet using meta-search engines: A review, *Journal of Scientific and Industrial Research* 67 (10) (2008).
- [3] G. Faggioli, M. Ferrante, N. Ferro, R. Perego, N. Tonello, Hierarchical dependence-aware evaluation measures for conversational search, in: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 1935–1939.
- [4] S. Li, R. Xie, Y. Zhu, X. Ao, F. Zhuang, Q. He, User-centric conversational recommendation with multi-aspect user modeling, in: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, pp. 223–233.
- [5] Z. Chen, X. Cheng, S. Dong, Z. Dou, J. Guo, X. Huang, Y. Lan, C. Li, R. Li, T.-Y. Liu, Y. Liu, J. Ma, B. Qin, M. Wang, J. rong Wen, J. Xu, M. Zhang, P. Zhang, Q. Zhang, Information retrieval: a view from the chinese ir community, *Frontiers of Computer Science* 15 (2020) 1–15.
- [6] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, et al., Sparks of artificial general intelligence: Early experiments with gpt-4, *arXiv preprint arXiv:2303.12712* (2023).
- [7] Y. Liu, T. Han, S. Ma, J. Zhang, Y. Yang, J. Tian, H. He, A. Li, M. He, Z. Liu, et al., Summary of chatgpt/gpt-4 research and perspective towards the future of large language models, *arXiv preprint arXiv:2304.01852* (2023).
- [8] H. Lee, S. Yang, H. Oh, M. Seo, Generative retrieval for long sequences, *arXiv preprint arXiv:2204.13596* (2022).
- [9] W. Wang, X. Lin, F. Feng, X. He, T.-S. Chua, Generative recommendation: Towards next-generation recommender paradigm, *arXiv preprint arXiv:2304.03516* (2023).
- [10] J. Zheng, M. Fischer, Bim-gpt: a prompt-based virtual assistant framework for bim information retrieval, *arXiv preprint arXiv:2304.09333* (2023).
- [11] V. Jeronymo, L. Bonifacio, H. Abonizio, M. Fadaee, R. Lotufo, J. Zavrel, R. Nogueira, Inpars-v2: Large language models as efficient dataset generators for information retrieval, *arXiv preprint arXiv:2301.01820* (2023).
- [12] S. Y. Kim, H. Park, K. Shin, K.-M. Kim, Ask me what you need: Product retrieval using knowledge from gpt-3, *arXiv preprint arXiv:2207.02516* (2022).
- [13] A. Edalati, M. Tahaei, A. Rashid, V. P. Nia, J. J. Clark, M. Rezagholizadeh, Kronecker decomposition for gpt compression, *arXiv preprint arXiv:2110.08152* (2021).
- [14] H. Liu, R. Ning, Z. Teng, J. Liu, Q. Zhou, Y. Zhang, Evaluating the logical reasoning ability of chatgpt and gpt-4, *arXiv preprint arXiv:2304.03439* (2023).
- [15] H. Nori, N. King, S. M. McKinney, D. Carignan, E. Horvitz, Capabilities of gpt-4 on medical challenge problems, *arXiv preprint arXiv:2303.13375* (2023).
- [16] Z. Liu, X. Yu, L. Zhang, Z. Wu, C. Cao, H. Dai, L. Zhao, W. Liu, D. Shen, Q. Li, T. Liu, D. Zhu, X. Li, Deid-gpt: Zero-shot medical text de-identification by gpt-4, *arXiv preprint arXiv:2303.11032* (2023).
- [17] J. Zhang, K. Bao, Y. Zhang, W. Wang, F. Feng, X. He, Is chatgpt fair for recommendation? evaluating fairness in large language model recommendation, in: *Recsys, ACM*, 2023.
- [18] A. Blair-Stanek, N. Holzenberger, B. V. Durme, Can gpt-3 perform statutory reasoning?, *arXiv preprint arXiv:2302.06100* (2023).
- [19] C. D. Manning, P. Raghavan, H. Schütze, *Introduction to Information Retrieval*, Cambridge University Press, Cambridge, UK, 2008.
- [20] M. Kobayashi, K. Takeda, Information retrieval on the web, *ACM Comput. Surv.* 32 (2) (2000) 144–173.
- [21] B. J. Jansen, D. L. Booth, A. Spink, Determining the user intent of web search engine queries, in: *Proceedings of the 16th International Conference on World Wide Web, WWW '07, Association for Computing Machinery, New York, NY, USA, 2007*, p. 1149–1150.

- [22] J. Teevan, S. T. Dumais, D. J. Liebling, To personalize or not to personalize: Modeling queries with variation in user intent, in: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Association for Computing Machinery, 2008, p. 163–170.
- [23] N. Su, J. He, Y. Liu, M. Zhang, S. Ma, User intent, behaviour, and perceived satisfaction in product search, WSDM '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 547–555.
- [24] L. Chen, D. Zhang, L. Mark, Understanding user intent in community question answering, in: Proceedings of the 21st International Conference on World Wide Web, WWW '12 Companion, Association for Computing Machinery, New York, NY, USA, 2012, p. 823–828.
- [25] E. Agichtein, E. Brill, S. Dumais, Improving web search ranking by incorporating user behavior information, in: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '06, Association for Computing Machinery, New York, NY, USA, 2006, p. 19–26.
- [26] E. Manavoglu, D. Pavlov, C. Giles, Probabilistic user behavior models, in: Third IEEE International Conference on Data Mining, 2003, pp. 203–210.
- [27] Q. Pi, W. Bian, G. Zhou, X. Zhu, K. Gai, Practice on long sequential user behavior modeling for click-through rate prediction, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 2671–2679.
- [28] F. Yuan, X. He, A. Karatzoglou, L. Zhang, Parameter-efficient transfer from sequential behaviors for user modeling and recommendation, SIGIR '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 1469–1478.
- [29] J. Yuan, Y. Zheng, X. Xie, Discovering regions of different functions in a city using human mobility and pois, KDD '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 186–194.
- [30] Y. Zheng, X. Xie, W.-Y. Ma, Geolife: A collaborative social networking service among user, location and trajectory, IEEE Data(base) Engineering Bulletin (June 2010).
- [31] L. Jin, Y. Chen, T. Wang, P. Hui, A. V. Vasilakos, Understanding user behavior in online social networks: a survey, IEEE Communications Magazine 51 (9) (2013) 144–150. doi:10.1109/MCOM.2013.6588663.
- [32] A. Alfieri, R. Wolter, S. H. Hashemi, Intent disambiguation for task-oriented dialogue systems, CIKM '22, Association for Computing Machinery, New York, NY, USA, 2022, p. 5079–5080.
- [33] J. Li, M. Galley, C. Brockett, G. P. Spithourakis, J. Gao, W. B. Dolan, A persona-based neural conversation model, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics, 2016.
- [34] P. Zhong, C. Zhang, H. Wang, Y. Liu, C. Miao, Towards persona-based empathetic conversational models, in: B. Webber, T. Cohn, Y. He, Y. Liu (Eds.), Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020, Association for Computational Linguistics, 2020, pp. 6556–6566.
- [35] J. Vassileva, Motivating participation in social computing applications: a user modeling perspective, User Model. User Adapt. Interact. 22 (1-2) (2012) 177–201.
- [36] J. Guo, Y. Fan, L. Pang, L. Yang, Q. Ai, H. Zamani, C. Wu, W. B. Croft, X. Cheng, A deep look into neural ranking models for information retrieval, Information Processing & Management 57 (6) (2020) 102067.
- [37] M. Trabelsi, Z. Chen, B. D. Davison, J. Heflin, Neural ranking models for document retrieval, Inf. Retr. 24 (6) (2021) 400–444.
- [38] S. Robertson, H. Zaragoza, The probabilistic relevance framework: Bm25 and beyond 3 (4) (2009) 333–389.
- [39] Y. Freund, R. Iyer, R. E. Schapire, Y. Singer, An efficient boosting algorithm for combining preferences, J. Mach. Learn. Res. 4 (null) (2003) 933–969.
- [40] C. J. C. Burges, From ranknet to lambdarank to lambdamart: An overview, 2010.
- [41] P.-S. Huang, X. He, J. Gao, L. Deng, A. Acero, L. Heck, Learning deep structured semantic models for web search using clickthrough data, CIKM '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 2333–2338.
- [42] L. Pang, Y. Lan, J. Guo, J. Xu, J. Xu, X. Cheng, Deeprank: A new deep architecture for relevance ranking in information retrieval, in: Proceedings of the 2017 ACM Conference on Information and Knowledge Management, CIKM '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 257–266.
- [43] J. Guo, Y. Fan, Q. Ai, W. B. Croft, A deep relevance matching model for ad-hoc retrieval, in: Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM '16, Association for Computing Machinery, New York, NY, USA, 2016, p. 55–64.
- [44] Y. Jiang, P. Zhang, H. Gao, D. Song, A quantum interference inspired neural matching model for ad-hoc retrieval, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 19–28.
- [45] Y. Fujiwara, M. Nakatsuji, H. Shiokawa, T. Mishima, M. Onizuka, Efficient ad-hoc search for personalized pagerank, in: Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, SIGMOD '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 445–456.
- [46] A. Karatzoglou, L. Baltrunas, Y. Shi, Learning to rank for recommender systems, in: Proceedings of the 7th ACM Conference on Recommender Systems, RecSys '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 493–494.
- [47] S. Vargas, P. Castells, Rank and relevance in novelty and diversity metrics for recommender systems, in: Proceedings of the Fifth ACM Conference on Recommender Systems, RecSys '11, Association for Computing Machinery, New York, NY, USA, 2011, p. 109–116.
- [48] K. Zhang, W. Wu, H. Wu, Z. Li, M. Zhou, Question retrieval with high quality answers in community question answering, in: Proceedings of the 23rd ACM International Conference on Information and Knowledge Management, CIKM '14, Association for Computing Machinery, New York, NY, USA, 2014, p. 371–380.
- [49] L. Yang, M. Qiu, S. Gottipati, F. Zhu, J. Jiang, H. Sun, Z. Chen, Cqrank: Jointly model topics and expertise in community question answering, in: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 99–108.
- [50] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton,

- F. Kelton, L. Miller, M. Simens, A. Askill, P. Welinder, P. F. Christiano, J. Leike, R. Lowe, Training language models to follow instructions with human feedback, in: *Advances in Neural Information Processing Systems*, Vol. 35, 2022, pp. 27730–27744.
- [51] K. Duh, K. Kirchhoff, Learning to rank with partially-labeled data, in: *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, Association for Computing Machinery, New York, NY, USA, 2008, p. 251–258.
- [52] V. Karpukhin, B. Oğuz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, W.-t. Yih, Dense passage retrieval for open-domain question answering, arXiv preprint arXiv:2004.04906 (2020).
- [53] L. Zhen, P. Hu, X. Wang, D. Peng, Deep supervised cross-modal retrieval, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10394–10403.
- [54] G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, E. Grave, Towards unsupervised dense information retrieval with contrastive learning, arXiv preprint arXiv:2112.09118 (2021).
- [55] Y. Tay, V. Tran, M. Dehghani, J. Ni, D. Bahri, H. Mehta, Z. Qin, K. Hui, Z. Zhao, J. Gupta, et al., Transformer memory as a differentiable search index, *Advances in Neural Information Processing Systems* 35 (2022) 21831–21843.
- [56] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askill, et al., Language models are few-shot learners, in: *NeurIPS*, Curran Associates, Inc., 2020, pp. 1877–1901.
- [57] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *CVPR, IEEE*, 2022, pp. 10684–10695.
- [58] H. Zamani, F. Diaz, M. Dehghani, D. Metzler, M. Bendersky, Retrieval-enhanced machine learning, in: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, Association for Computing Machinery, New York, NY, USA, 2022, p. 2875–2886. doi:10.1145/3477495.3531722. URL <https://doi.org/10.1145/3477495.3531722>
- [59] S. Gur, N. Neverova, C. Stauffer, S.-N. Lim, D. Kiela, A. Reiter, Cross-modal retrieval augmentation for multi-modal classification, in: *Findings of the Association for Computational Linguistics: EMNLP 2021*, Association for Computational Linguistics, Punta Cana, Dominican Republic, 2021, pp. 111–123. doi:10.18653/v1/2021.findings-emnlp.11. URL <https://aclanthology.org/2021.findings-emnlp.11>
- [60] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. van den Driessche, J. Lespiau, B. Damoc, A. Clark, D. de Las Casas, A. Guy, J. Menick, R. Ring, T. Hennigan, S. Huang, L. Maggiore, C. Jones, A. Cassirer, A. Brock, M. Paganini, G. Irving, O. Vinyals, S. Osindero, K. Simonyan, J. W. Rae, E. Elsen, L. Sifre, Improving language models by retrieving from trillions of tokens, arXiv preprint arXiv:2112.04426 (2021).
- [61] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, T. Lillicrap, Meta-learning with memory-augmented neural networks, in: *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16, JMLR.org*, 2016, p. 1842–1850.
- [62] N. Ford, D. Miller, A. O'rourke, J. Ralph, E. Turnock, A. Booth, Information retrieval for evidence-based decision making, *Journal of documentation* 55 (4) (1999) 385–401.
- [63] D. Xu, T. Schnabel, X. Cui, S. Dean, A. Deshmukh, B. Yang, S. Yu, Foreword for workshop on decision making for information retrieval and recommender systems, in: *Companion Proceedings of the ACM Web Conference 2023*, 2023, pp. 920–920.
- [64] P. Borlund, Interactive information retrieval: An introduction, *Journal of Information Science Theory and Practice* 1 (3) (2013) 12–32.
- [65] H. Yin, Y. Sun, G. Xu, E. Kanoulas, Trustworthy recommendation and search: Introduction to the special issue-part 1 (2023).
- [66] V. Y. Tsvetkov, Cognitive science of information retrieval, *European Journal of Psychological Studies* (1) (2015) 37–44.
- [67] P. Ingwersen, Psychological aspects of information retrieval, *Social Science Information Studies* 4 (2-3) (1984) 83–95.
- [68] M. Janner, Deep generative models for decision-making and control, arXiv preprint arXiv:2306.08810 (2023).
- [69] D. Esposito, J. Centracchio, E. Andreozzi, G. D. Gargiulo, G. R. Naik, P. Bifulco, Biosignal-based human-machine interfaces for assistance and rehabilitation: A survey, *Sensors* 21 (20) (2021) 6863.
- [70] S. A. H. Mohsan, M. A. Khan, F. Noor, I. Ullah, M. H. Alsharif, Towards the unmanned aerial vehicles (uavs): A comprehensive review, *Drones* 6 (6) (2022) 147.
- [71] M. Zhu, S. Biswas, S. I. Dinulescu, N. Kastor, E. W. Hawkes, Y. Visell, Soft, wearable robotics and haptics: Technologies, trends, and emerging applications, *Proceedings of the IEEE* 110 (2) (2022) 246–272.
- [72] S. Mystakidis, *Metaverse*, *Encyclopedia* 2 (1) (2022) 486–497.
- [73] F. Shenavarmasouleh, F. G. Mohammadi, M. H. Amini, H. Reza Arabnia, Embodied ai-driven operation of smart cities: A concise review, *Cyberphysical Smart Cities Infrastructures: Optimal Operation and Intelligent Decision Making* (2022) 29–45.
- [74] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. E. Miller, M. Simens, A. Askill, P. Welinder, P. F. Christiano, J. Leike, R. J. Lowe, Training language models to follow instructions with human feedback, in: *NeurIPS*, 2022.
- [75] R. I. John, G. J. Mooney, Fuzzy user modeling for information retrieval on the world wide web, *Knowl. Inf. Syst.* 3 (1) (2001) 81–95.
- [76] W. R. Hersch, C. Buckley, T. J. Leone, D. H. Hickam, OHSUMED: an interactive retrieval evaluation and new large test collection for research, in: *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, ACM/Springer, 1994, pp. 192–201.
- [77] S. Dennis, P. Bruza, R. McArthur, Web searching: A process-oriented experimental study of three interactive search paradigms, *J. Assoc. Inf. Sci. Technol.* 53 (2) (2002) 120–133.
- [78] Y. Guo, Z. Cheng, L. Nie, Y. Wang, J. Ma, M. S. Kankanhalli, Attentive long short-term preference modeling for personalized product search, *ACM Trans. Inf. Syst.* 37 (2) (2019) 19:1–19:27.
- [79] K. Bao, J. Zhang, Y. Zhang, W. Wang, F. Feng, X. He, Tallrec: An effective and efficient tuning framework to align large language model with recommendation, in: *Recsys, ACM*, 2023.
- [80] Y. Gao, T. Sheng, Y. Xiang, Y. Xiong, H. Wang, J. Zhang, Chat-rec: Towards interactive and explainable llms-augmented recommender system, *CoRR* abs/2303.14524 (2023).
- [81] Y. Sun, X. Wang, Z. Liu, J. Miller, A. Efros, M. Hardt, Test-time training with self-supervision for generalization under distribution shifts,

- in: International conference on machine learning, PMLR, 2020, pp. 9229–9248.
- [82] J. Feng, C. Tao, X. Geng, T. Shen, C. Xu, G. Long, D. Zhao, D. Jiang, Knowledge refinement via interaction between search engines and large language models, arXiv preprint arXiv:2305.07402 (2023).
- [83] Y. Qin, Z. Cai, D. Jin, L. Yan, S. Liang, K. Zhu, Y. Lin, X. Han, N. Ding, H. Wang, R. Xie, F. Qi, Z. Liu, M. Sun, J. Zhou, Webcpm: Interactive web search for chinese long-form question answering, in: Proceedings of ACL 2023, Association for Computational Linguistics, 2023.
- [84] W. Sun, S. Guo, S. Zhang, P. Ren, Z. Chen, M. de Rijke, Z. Ren, Metaphorical user simulators for evaluating task-oriented dialogue systems, ACM Transactions on Information Systems (2022).
- [85] N. Craswell, Mean reciprocal rank, Encyclopedia of Database Systems (2009).
- [86] M. Zhu, Recall, precision and average precision, University of Waterloo, 2004.
- [87] K. Järvelin, J. Kekäläinen, Cumulated gain-based evaluation of ir techniques, ACM Transactions on Information Systems (TOIS) 20 (4) (2002) 422–446.
- [88] W. Sun, L. Yan, X. Ma, P. Ren, D. Yin, Z. Ren, Is chatgpt good at search? investigating large language models as re-ranking agent, arXiv preprint arXiv:2304.09542 (2023).
- [89] B. Mitra, N. Craswell, Neural models for information retrieval, CoRR abs/1705.01509 (2017).
- [90] J. Zhan, X. Xie, J. Mao, Y. Liu, J. Guo, M. Zhang, S. Ma, Evaluating interpolation and extrapolation performance of neural retrieval models, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 2486–2496.
- [91] M. Llorde, D. Ganguly, S. Bhatia, C. Agarwal, Explain like I am BM25: interpreting a dense model’s ranked-list with a sparse approximation, CoRR abs/2304.12631 (2023).
- [92] K. Santhanam, J. Saad-Falcon, M. Franz, O. Khattab, A. Sil, R. Florian, M. A. Sultan, S. Roukos, M. Zaharia, C. Potts, Moving beyond downstream task accuracy for information retrieval benchmarking, CoRR abs/2212.01340 (2022).
- [93] H. He, H. Zhang, D. Roth, Rethinking with retrieval: Faithful large language model inference, arXiv preprint arXiv:2301.00303 (2022).
- [94] X. Shen, Z. Chen, M. Backes, Y. Zhang, In chatgpt we trust? measuring and characterizing the reliability of chatgpt, arXiv preprint arXiv:2304.08979 (2023).
- [95] R. Baeza-Yates, B. Ribeiro-Neto, et al., Modern information retrieval, Vol. 463, 1999.
- [96] F. Radlinski, N. Craswell, A theoretical framework for conversational search, in: Proceedings of the 2017 conference on conference human information interaction and retrieval, 2017, pp. 117–126.
- [97] P. Ren, Z. Chen, Z. Ren, E. Kanoulas, C. Monz, M. De Rijke, Conversations with search engines: Serp-based conversational response generation, ACM Transactions on Information Systems (TOIS) 39 (4) (2021) 1–29.
- [98] T. Schick, J. Dwivedi-Yu, R. Dessi, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, T. Scialom, Toolformer: Language models can teach themselves to use tools, arXiv preprint arXiv:2302.04761 (2023).
- [99] M. Komeili, K. Shuster, J. Weston, Internet-augmented dialogue generation, arXiv preprint arXiv:2107.07566 (2021).
- [100] A. Lazaridou, E. Gribovskaya, W. Stokowiec, N. Grigorev, Internet-augmented language models through few-shot prompting for open-domain question answering, arXiv preprint arXiv:2203.05115 (2022).
- [101] N. Lee, W. Ping, P. Xu, M. Patwary, P. N. Fung, M. Shoenybi, B. Catanzaro, Factuality enhanced language models for open-ended text generation, Advances in Neural Information Processing Systems 35 (2022) 34586–34599.
- [102] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, A. Miller, Language models as knowledge bases?, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 2463–2473.
- [103] G. Izacard, P. Lewis, M. Lomeli, L. Hosseini, F. Petroni, T. Schick, J. Dwivedi-Yu, A. Joulin, S. Riedel, E. Grave, Atlas: Few-shot learning with retrieval augmented language models, arXiv preprint arXiv 2208 (2022).
- [104] K. Guu, K. Lee, Z. Tung, P. Pasupat, M.-w. Chang, Realm: Retrieval-augmented language model pre, ICML (2020).
- [105] S. Borgeaud, A. Mensch, J. Hoffmann, T. Cai, E. Rutherford, K. Millican, G. B. Van Den Driessche, J.-B. Lespiau, B. Damoc, A. Clark, et al., Improving language models by retrieving from trillions of tokens, in: International conference on machine learning, PMLR, 2022, pp. 2206–2240.
- [106] B. Wang, W. Ping, P. Xu, L. McAfee, Z. Liu, M. Shoenybi, Y. Dong, O. Kuchaiev, B. Li, C. Xiao, et al., Shall we pretrain autoregressive language models with retrieval? a comprehensive study, arXiv preprint arXiv:2304.06762 (2023).
- [107] C. Xu, Q. Li, J. Ge, J. Gao, X. Yang, C. Pei, F. Sun, J. Wu, H. Sun, W. Ou, Privileged features distillation at taobao recommendations, in: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2020, pp. 2590–2598.
- [108] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, et al., Webgpt: Browser-assisted question-answering with human feedback, arXiv preprint arXiv:2112.09332 (2021).
- [109] J. He, C. Zhou, X. Ma, T. Berg-Kirkpatrick, G. Neubig, Towards a unified view of parameter-efficient transfer learning, in: International Conference on Learning Representations, 2022.
- [110] Z. Hu, Y. Lan, L. Wang, W. Xu, E.-P. Lim, R. K.-W. Lee, L. Bing, S. Poria, Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models, arXiv preprint arXiv:2304.01933 (2023).
- [111] X. L. Li, P. Liang, Prefix-tuning: Optimizing continuous prompts for 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), 2021, pp. 4582–4597.
- [112] X. Liu, K. Ji, Y. Fu, W. Tam, Z. Du, Z. Yang, J. Tang, P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2022, pp. 61–68.
- [113] X. Liu, Y. Zheng, Z. Du, M. Ding, Y. Qian, Z. Yang, J. Tang, Gpt understands, too, arXiv preprint arXiv:2103.10385 (2021).
- [114] B. Lester, R. Al-Rfou, N. Constant, The power of scale for parameter-efficient prompt tuning, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 3045–3059.
- [115] N. Houslyby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, S. Gelly, Parameter-efficient transfer learning for nlp, in: International Conference on Machine Learning, 2019, pp. 2790–2799.

- [116] J. Pfeiffer, A. Kamath, A. Rücklé, K. Cho, I. Gurevych, AdapterFusion: Non-destructive task composition for transfer learning, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, 2021, pp. 487–503.
- [117] Q. Zhang, M. Chen, A. Bukharin, P. He, Y. Cheng, W. Chen, T. Zhao, Adaptive budget allocation for parameter-efficient fine-tuning, arXiv preprint arXiv:2303.10512 (2023).
- [118] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, Lora: Low-rank adaptation of large language models, arXiv preprint arXiv:2106.09685 (2021).
- [119] K. Shuster, M. Komeili, L. Adolphs, S. Roller, A. Szlam, J. Weston, Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion, arXiv preprint arXiv:2203.13224 (2022).
- [120] O. Ram, Y. Levine, I. Dalmedigos, D. Muhlgay, A. Shashua, K. Leyton-Brown, Y. Shoham, In-context retrieval-augmented language models, arXiv preprint arXiv:2302.00083 (2023).
- [121] B. Peng, M. Galley, P. He, H. Cheng, Y. Xie, Y. Hu, Q. Huang, L. Liden, Z. Yu, W. Chen, J. Gao, Check your facts and try again: Improving large language models with external knowledge and automated feedback, arXiv preprint arXiv:2302.12813 (2023).