

# Advancing the Search Frontier with AI Agents

Ryen W. White  
Microsoft Research  
Redmond, WA, USA  
ryenw@microsoft.com

## ABSTRACT

As many of us in the information retrieval (IR) research community know and appreciate, search is far from being a solved problem. Millions of people struggle with tasks on search engines every day. Often, their struggles relate to the intrinsic complexity of their task and the failure of search systems to fully understand the task and serve relevant results [62]. The task motivates the search, creating the gap/problematic situation that searchers attempt to bridge/resolve and drives search behavior as they work through different task facets. Complex search tasks require more than support for rudimentary fact finding or re-finding. Research on methods to support complex tasks includes work on generating query and website suggestions [23, 66], personalizing and contextualizing search [4], and developing new search experiences, including those that span time and space [1, 68]. The recent emergence of generative artificial intelligence (AI) and the arrival of assistive *agents*, based on this technology, has the potential to offer further assistance to searchers, especially those engaged in complex tasks [45, 65]. There are profound implications from these advances for the design of intelligent systems and for the future of search itself. This article, based on a keynote by the author at the 2023 ACM SIGIR Conference, explores these issues and how AI agents are advancing the frontier of search system capabilities, with a special focus on information interaction and complex task completion.

## ACM Reference Format:

Ryen W. White. 2024. Advancing the Search Frontier with AI Agents. Accepted in *Communications of the ACM*, March, 2024

## 1 TAKING SEARCH TO TASK

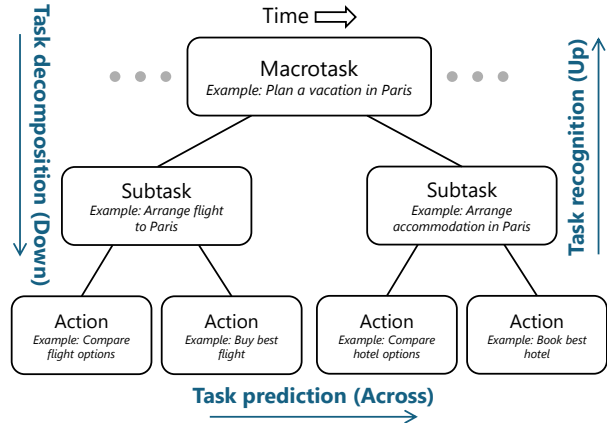
Tasks are a critical part of people’s daily lives. The market for dedicated task applications that help people with their “to do” tasks is likely to grow significantly (effectively triple in size) over the next few years.<sup>1</sup> There are many examples of such applications that can help both individuals (e.g., Microsoft To Do, Google Tasks, Todoist) and teams (e.g., Asana, Trello, Monday.com) tackle their tasks more effectively. Over time, these systems will increasingly integrate AI to better help their users capture, manage, and complete their tasks [64]. In information access scenarios such as search, tasks play an important role in motivating searching via gaps in knowledge and

<sup>1</sup><https://www.verifiedmarketresearch.com/product/task-management-software-market/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Accepted in *Communications of the ACM*, March, 2024

© 2024 Copyright held by the owner/author(s).



**Figure 1: Task tree representation for a complex task involving planning a vacation to Paris, France. The tree depicts different task granularities (macrotask, subtask, action) and different task applications (decomposition, prediction, recognition) as moves around the tree. Time progresses from left to right via a sequence of searcher actions (queries, result clicks, pagination, etc.). Only actions are observable in traditional search engines. Aspects of subtasks and macrotasks may be observable to AI agents when searchers provide higher-level descriptions of their goals in natural language.**

problematic situations [3, 15]. AI can be central in these search scenarios too, especially in assisting with complex search tasks.

### 1.1 Tasks in Search

Tasks drive the search process. The IR and information science communities have long studied tasks in search [46] and many information seeking models consider the role of task directly [3, 15]. Prior research has explored the different stages of task execution (e.g., pre-focus, focus formation, post-focus) [57], task levels [43], task facets [32], tasks defined on intents (e.g., informational, transactional, and navigational [8]; well-defined or ill-defined [25]; lookup, learn, or investigate [36]), the hierarchical structure of tasks [73], the characteristics of tasks, and the attributes of task-searcher interaction, e.g., task difficulty and, of course, a focus in this article, task complexity [11, 28].

As a useful framing device to help conceptualize tasks and develop system support for them, tasks can be represented as *trees* comprising macrotasks (high level goals), subtasks (specific components of those goals), and actions (specific steps taken by searchers toward the completion of those components) [46]. Figure 1 presents an example of a “task tree” for a task involving an upcoming vacation to Paris, France. Examples of macrotasks, subtasks, and actions

are included. Moves around this tree correspond to different task applications such as task recognition (up), task decomposition (down), and task prediction (across). Only actions (e.g., queries, clicks, and so on) are directly observable to traditional search engines. However, with recent advances in AI agents (primarily more fully supporting natural language interactions to improve alignment between searchers and AI agents, but also a growing system awareness of short- and long-term contexts), more aspects of macrotasks and subtasks are becoming visible to search systems and more fully understood by those systems. Challenges in working with tasks include how to represent them within search systems, how to observe more task-relevant activity and content to develop richer task models, and how to develop task-oriented interfaces that place tasks and their completion at the forefront of user engagement. Task complexity deserves a special focus in this article given the challenges that searchers can still face with complex tasks and the significant potential of AI to help searchers resolve complex tasks.

## 1.2 Complex Search Tasks

Recent estimates suggest that half of all Web searches are not answered.<sup>2</sup> Many of those searches are connected to complex search tasks. These tasks are ill-defined and/or multi-step, span multiple queries, sessions, and/or devices, and require deep engagement with search engines (many queries, backtracking, branching, etc.) to complete them [23]. Complex tasks also often have many facets and cognitive dimensions, and are closely connected to searcher characteristics such as domain expertise and task familiarity [42, 62].

To date, there have been significant attempts to support complex search tasks via humans (e.g., librarians, subject matter experts) and search systems (both general Web search engines and those tailored to specific industry verticals or domains). The main technological progress so far has been in areas such as query suggestion and contextual search, with new experiences also being developed that utilize multiple devices, provide cross-session support, and enable conversational search. We are now also seeing an emerging wave of search-related technologies in the area of generative AI [39].

Before proceeding, let us dive into these different types of existing and emerging search support for complex tasks in more detail.

- *Suggestions, personalization, and contextualization*: Researchers and practitioners have long developed and deployed support such as query suggestion and trail suggestion, e.g., [23, 49], including providing guided tours [55] and suggesting popular trail destinations [66] as ways to find relevant resources. This coincides with work on contextual search and personalized search, e.g., [4, 51, 67], where search systems can use data from the current searcher such as session activity, location, reading level, and so on, and the searcher’s long-term activity history, to provide more relevant results. Search engines may also use cohort activities to help with cold-start problems for new users and augment personal profiles for more established searchers [52, 74].

- *Multi-device, cross-device, and cross-session*: Devices have different capabilities and can be used in different settings. Multi-device experiences, e.g., [68], utilize multiple devices *simultaneously* to better support complex tasks such as recipe preparation, auto repair, and home improvement that have been decomposed into steps manually or automatically [78]. Cross-device and cross-session support [1, 60] can help with ongoing/background searches for complex tasks that persist over space and time. For example, being able to predict task continuation can help with “slow search” applications that focus more on result quality than the near instantaneous retrieval of search results [50].
- *Conversational experiences and generative AI*: Natural language is an expressive and powerful means of communicating intentions and preferences with search systems. The introduction of clarification questions on search engine result pages (SERPs) [76], progress on conversational search [21], and even “conversations” with documents (where searchers can inquire about document content via natural language dialog) [53], enable these systems to engage more fully with searchers to better understand their tasks and goals. There are now many emerging opportunities to better align search systems and their users, and support more tasks, via large-scale foundation models such as OpenAI’s GPT,<sup>3</sup> Google’s Gemini,<sup>4</sup> and Meta’s Llama,<sup>5</sup> including offering conversational task assistance via chatbots such as ChatGPT.<sup>6</sup>

All of these advances, and others, have paved the way for the emergence of a new class of generative-AI-powered assistive agents that can help people make progress in their complex search tasks.

## 2 AI AGENTS

Agents are applications of modern AI (foundation models, etc.) to help people with complex cognitive tasks. At Microsoft, we refer to these as *copilots*, which work alongside humans to empower them and amplify their cognitive capabilities.<sup>7</sup> Copilots have conversational user interfaces and their users engage with them via natural language, they are powered by foundation models such as GPT-4, they are extensible with skills/tools/plugins, and they are scoped to specialized domains or applications (including search). Copilots are designed to keep humans at the center of the task completion process and augment human capabilities to help them complete a broader range of tasks in less time and with less effort.

The general AI agent stack (Figure 2) contains four layers: (1) The *frontend*, covering the user experience and extensibility with plugins, enabling developers to provide additional visible tools to the agent; (2) The *AI orchestration* layer that handles the internal information flows, prompting, grounding, executing the plugins and processing their responses, among other things; (3) Agents leverage the power of large *foundation models* that can be provided to the developer as is, specialized to specific tasks, domains, or applications, or developers can bring their own models to use to power agent functionality, and; (4) This all runs on top of massive scale *AI infrastructure* hosted in the cloud on platforms such as Microsoft Azure, Google Cloud, and Amazon Web Services. Underpinning all

<sup>2</sup><https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/>

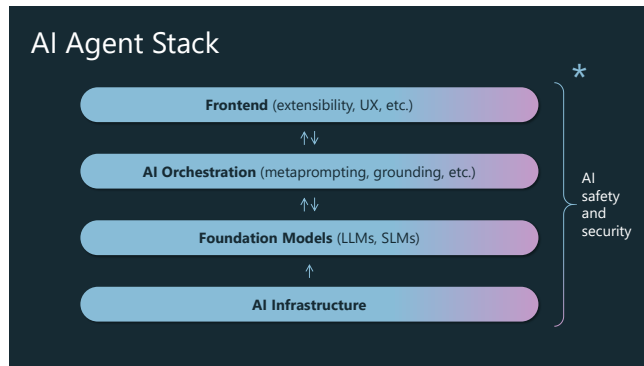
<sup>3</sup><https://openai.com/gpt-4>

<sup>4</sup><https://gemini.google.com>

<sup>5</sup><https://llama.meta.com>

<sup>6</sup><https://openai.com/chatgpt>

<sup>7</sup><https://copilot.microsoft.com>



**Figure 2: AI agent stack depicting the various layers and the important role of AI safety and security across the stack. Foundation models can be either large language models (LLMs), with trillions of parameters, or small language models (SLMs), with just a few billion parameters. The star (\*) symbolizes that there can be multiple cooperating agents as discussed later in Section 4.1.7.**

of this is a need for a strong commitment to responsible AI, which ensures that agents are safe, secure, and transparent. We can do this via an iterative, layered approach with mitigations spanning the model, prompts, grounding, and the user experience.

AI agents can, among other things, help users attain goals, maximize utility, and perform automated functions. Examples of these agents include the Apple Siri and Amazon Alexa personal digital assistants that can answer questions and assist with task management, GitHub Copilot,<sup>8</sup> an AI pair programmer that has been shown to reduce developer effort, enable more task success, and significantly expedite task completion [5], and Auto-GPT,<sup>9</sup> a fully-autonomous agent that can decompose tasks into sub-tasks and execute them independently on a user’s behalf to support goal attainment.

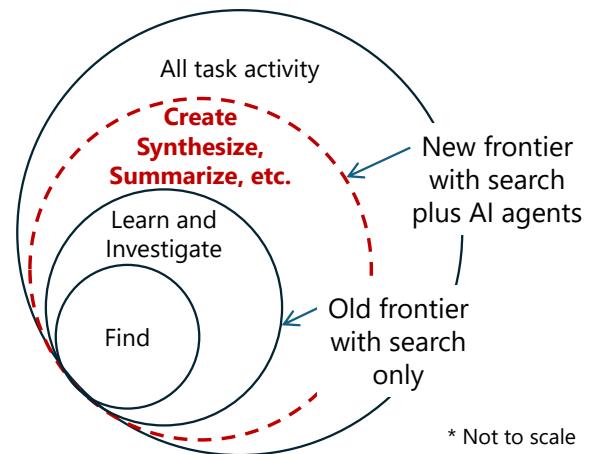
AI agents are also emerging in search systems. Popular Web search engines such as Bing and Google are adding agent functionalities in the form of conversational assistance; Bing has Copilot, mentioned earlier, Google has Gemini, a similar service. In search, agents can help searchers to tackle a broader range of tasks than information finding and go deeper than surface (SERP-level) interactions with content by synthesizing answers on the searcher’s behalf. They also enable searchers to communicate their intents and goals more directly. Returning to the task tree (Figure 1), the focus on engaging agents via natural language interactions allows both searchers and systems to consider higher-level task representations (macrotasks, subtasks) in addition to the more granular actions (queries, result clicks, pagination, and so on) that searchers already perform when engaging with traditional search engines.

### 2.1 Agents in Search

Agents and chat experiences are a complement not a replacement to traditional search engines. Search engines have existed for decades and serve a valuable purpose: providing near instantaneous access

<sup>8</sup><https://github.com/features/copilot>

<sup>9</sup><https://auto-gpt.ai/>



**Figure 3: Advancing the search frontier. Visualizing the set of possible tasks that can be tackled with search only today (i.e., finding, learning, and investigating) plus the expansion in the frontier into support for higher-order task activities with the addition of AI agents (e.g., adding emerging AI support for creative inspiration, synthesis, and summarization).**

to answers and resources for a broad range of search requests. These existing and emerging modalities can and should work well together to help searchers tackle a wider range of tasks.

The ability of agents to better understand intentions and provide assistance beyond fact finding and basic learning/investigation will advance the *search frontier* (i.e., what search systems are capable of and what types of tasks they can support), broadening the range of tasks that searchers can complete, e.g., direct support for tasks requiring creative inspiration, summarizing existing perspectives, or synthesizing those perspectives to generate new insight (Figure 3). This all moves us toward more intelligent search systems that can help with all-task completion, covering the full universe of tasks for which people might need search support, including actuation capabilities to act on tasks in the digital and physical worlds.

One way to define the range of tasks that agents can support is through Bloom’s taxonomy of learning objectives [29]. Creation is at the pinnacle of that taxonomy and we have only scratched the surface in creativity support with next-word prediction through transformer models [16]. We are already seeing expansions into content types beyond text (images, video, audio, and so on) and could consider support for other creative tasks including planning, analysis, and invention. There are also many other layers in Bloom’s taxonomy (e.g., evaluation - help searchers make judgments and decisions, application - help searchers complete new tasks, understanding - explain ideas and concepts to accelerate learning) that could form the basis for future search frontiers.

Beyond greater capabilities, the introduction of AI agents into search will also change how people will engage with search systems. In agents, the mode of interaction is primarily natural language, with some recent expansion toward other content types for both

inputs and outputs via the introduction of image and video generation models, e.g., Stable Diffusion,<sup>10</sup> DALL·E,<sup>11</sup> and Sora.<sup>12</sup> Agents can generate direct answers, with source attribution for provenance, to build trust with users, and to drive traffic back to content creators, which is important to incentivize further content creation that fuels future foundation models. The overall search interaction flow is also different between search engines and AI agents. When using agents, searchers do not need to decompose their goal into sub-goals or sub-queries, examine SERPs and landing pages, and aggregate/synthesize relevant knowledge from retrieved information. Continuing our running example macrotask of vacation planning from earlier, Figure 4 has a comparison of information interaction in the two modalities for some task-related goals. In AI agents, the responsibility for generating answers is delegated by the searcher to the system, which poses challenges in terms of human control and human learning, as we will discuss later in this article.

## 2.2 Adding Agents to Search Engines

It is not practical nor necessary to deploy AI agents for all search tasks. Foundation model inference is expensive at massive scale and search engine algorithms have been honed over decades to provide relevant results for a broad range of tasks (e.g., navigation, fact-finding). Conversational interfaces are less familiar for searchers and it will take time for searchers to adapt to this way of searching. Traditional search engines are sufficient when searchers know exactly what they want. Agents are helpful for more complex search tasks or in situations where searchers may be struggling to find relevant information. Task complexity can be estimated using aggregate metrics such as the amount of engagement with the search engine (e.g., number of query reformulations) for similar tasks historically. As generative AI appears in more applications and searchers better understand agent capabilities, the tasks that searchers bring to agents deployed in search settings will likely evolve and expand, and may well also increase in complexity.

We will also see a growth in search experiences that unify traditional search and copilots. In a step towards this, search engines such as Bing and Google are already integrating dynamic answers from foundation models into their SERPs for some queries (e.g., Google’s so-called “search generative experience”<sup>13</sup>). In these experiences, search results and other answers can be shown together on the SERP with answers from generative AI, allowing searchers to easily engage with them as desired, including asking follow-up questions inline. There are also ways for searchers to move between modalities based on their task and personal preferences. AI agents can also provide searchers with control over other aspects, e.g., Bing offers an ability to adjust conversation style and tone, although it is not clear that searchers are sufficiently familiar with agents at this time to use these more nuanced controls effectively.

Search agents such as Copilot and Gemini use retrieval augmented generation (RAG) [30] to ground copilot responses via timely and relevant results. This has many advantages, including: (1) There is no need to retrain massive foundation models over time; (2) Search results provide relevant and fresh information to

**Table 1: Anecdotal examples of high-level tasks from Bloom’s taxonomy [29] of varying complexities from searcher/copilot perspectives. Tasks such as ‘Find’ and ‘Analyze’ have similar complexities for both humans and machines. It is easier for machines to create content than for humans, but more difficult for machines to verify the correctness of information.**

		Searcher	
		Low	High
AI Agent	Low	Find Recognize List Define	Create Evaluate Compare Predict
	High	Verify Decide Teach Plan	Analyze Investigate Solve Invent

foundation models, and; (3) It provides a provenance signal connecting generated content with online sources. In response to a searcher prompt, the foundation model generates internal queries iteratively that are used to retrieve the results that form context for the agent answers created using generative AI. Displaying these queries to searchers inline in dialog, as in Copilot, creates greater transparency and helps build trust with searchers that the system is understanding their tasks and goals. The orchestrator can also pull in relevant instant answers from the search engine such as weather, stock, and sports, and display those in copilot responses instead of or in addition to the answers generated by the foundation model. Figure 5 shows the high-level search process from query (+ conversation context) to the answer, and the role of various key system components.

AI agents also enable search engines to support more complex search tasks. Using search alone would require more searcher effort to examine search results and manually generate answers or insights (see recent work on the Delphic costs and benefits of search [9]). Of course, there are different perspectives on task complexity, e.g., the agent perspective (denoting the amount of computation, requests, etc. required for the system to complete the task) and the searcher perspective (denoting the amount of manual effort required for the human searcher to generate an answer and complete a task). Table 1 considers the task complexity from these two different perspectives and (again drawing from Bloom’s taxonomy) provides some current, anecdotal examples of the types of tasks that both searcher and systems may find to be more or less complex. Assuming that foundational model costs will drop and sophistication will increase, we focus here on the task complexity for searchers.

## 3 CHALLENGES

Despite the promise of AI agents to dramatically improve information literacy, there are significant challenges in deploying them in search systems at scale and we must find ways to overcome these challenges. Those include issues with the agent output shown in response to searcher requests, the impacts that the agents can have on searchers, and shifts in the degree of control that humans have in the search process that come from introducing agents.

<sup>10</sup><https://stability.ai/stable-image>

<sup>11</sup><https://openai.com/dall-e-3>

<sup>12</sup><https://openai.com/sora>

<sup>13</sup><http://www.google.com/sg>

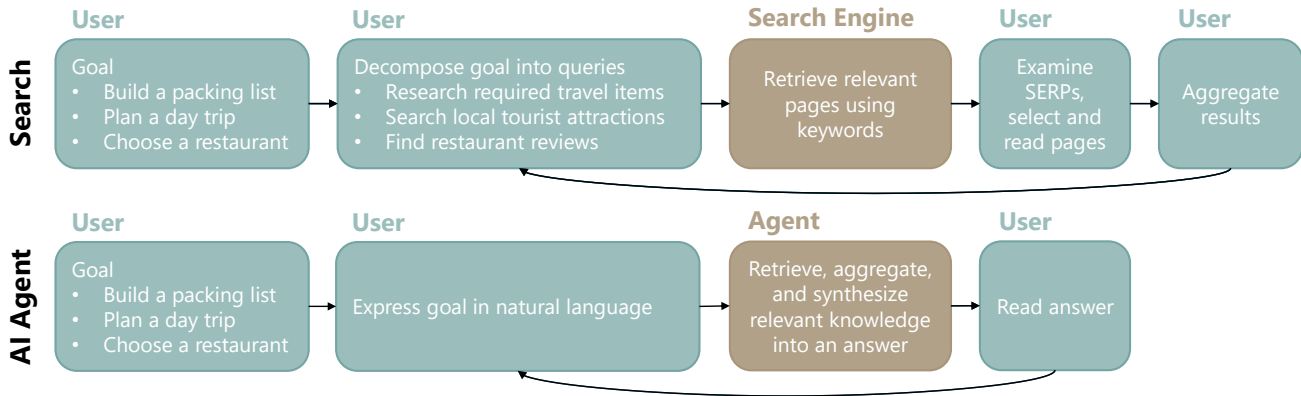


Figure 4: Information interactions in a traditional search engine versus an AI agent.

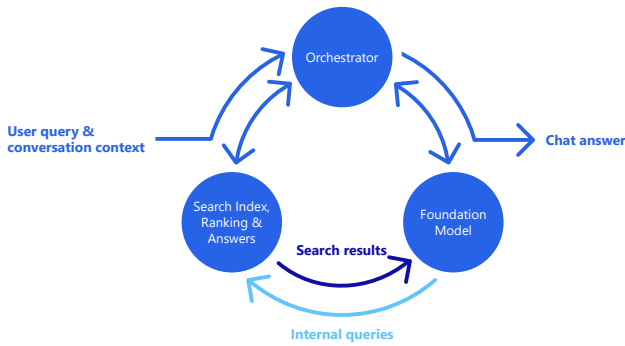


Figure 5: High-level overview of the typical generative AI search process in search engines. The query and the context are passed to the orchestrator, which coordinates with foundation model to create internal queries and generate answers. The orchestrator may also integrate content (e.g., search results, direct answers) from the search engine.

- **Hallucinations:** Searchers rely a lot on the answers from agents, but those answers can be erroneous or non-sensical. So-called “hallucination” is a well-studied problem in foundation models [26]. Agents can hallucinate for many reasons. One of the main reasons being gaps in the training data. RAG, discussed earlier, is a way to help address this by ensuring that the agent has access to up-to-date, relevant information at inference time to help ground its responses. Injection of knowledge from other external sources, such as knowledge graphs and Wikipedia, can also help improve the accuracy of agent responses. An issue related to agents surfacing misinformation is toxicity (i.e., offensive or harmful content), which can also be present in the agent output, and must be mitigated before answers are shown to searchers.
- **Biases:** Biases in the training data, e.g., social biases and stereotypes [34], affect the output of foundation models and hence the answers provided by agents. Synthesis of content from different sources can amplify biases in this data. As with hallucinations,

this is a well studied problem [6]. Agents are also subject to biases from learning from their own or other AI-generated content (via feedback loops); biased historical sequences lead to biased downstream models. Agents may also amplify existing cognitive biases, such as confirmation bias, by favoring responses that are aligned with searchers’ existing beliefs and values, and by providing responses that are optimized to keep searchers engaged with the agent, regardless of the consequences for the searcher.

- **Human learning:** Learning may be affected/interrupted by the use of AI agents since they remove the need for searchers to engage as fully with the search system and the information retrieved. Learning is already a core part of the search process [36, 41, 58]. Both exploratory search and search as learning involve considerable time and effort in finding and examining relevant content. While this could be viewed as a cost, this deep exposure to content also helps people learn. As mentioned earlier, agent users can ask richer questions (allowing them to specify their tasks and goals more fully) but they then receive synthesized answers generated by the agent, creating fewer, new, or simply different learning opportunities for humans that must be understood.
- **Human control:** Supporting search requires considering the degree of searcher involvement in the search process, which varies depending on the search task [2]. Agents enable more strategic, higher-order actions (higher up the “task tree” from Figure 1 than typical interactions with search systems). It is clear that searchers want control over the search process. They want to know what information is/not being included and why. This helps them understand and trust system output. As things stand, searchers delegate full control of answer generation to the AI, but the rest is mixed, i.e., less control of search mechanics (queries, etc.) but more control of task specifications (via natural language and dialog). There is more than just a basic tension between automation and control. In reality, it is not a zero sum game. Agent designers need to ensure human control while increasing automation [48]. New frameworks for multi-agent task completion are moving in this direction [31, 71], with agents and humans working together synergistically to decompose and tackle complex tasks.

Overall, these are just a few of the challenges that affect the viability of AI agents in search settings. There are other challenges,

Model Innovation	Next-Generation Experiences	Measurement	Broader Implications
<ul style="list-style-type: none"> <li>• Task Modeling</li> <li>• Alignment</li> <li>• Augmentation</li> <li>• Grounding</li> <li>• Personalization</li> <li>• Adaptation</li> <li>• Multi-agent</li> </ul>	<ul style="list-style-type: none"> <li>• Multimodality</li> <li>• Learning</li> <li>• Control</li> <li>• Completion</li> </ul>	<ul style="list-style-type: none"> <li>• Understanding</li> <li>• Evaluation</li> </ul>	<ul style="list-style-type: none"> <li>• Responsibility</li> <li>• Economics</li> <li>• Ubiquity</li> </ul>

**Figure 6: Selected opportunities for progress on AI agents in search settings. These opportunities are only a first step in this direction but the figure highlights the many avenues for impact in this area.**

such as searchers’ deeply ingrained search habits that may be a barrier to their adoption of new search functionality, despite the clear benefits to them from embracing agent technologies.

## 4 OPPORTUNITIES

For some time, scholars have argued that the future of information interaction will involve personal search assistants with advanced capabilities, including natural language input, rich sensing, user/task/world models, and reactive and proactive experiences [61]. Technology is catching up with this vision. Opportunities going forward can be grouped into four areas: (1) Model innovation; (2) Next-generation experiences; (3) Measurement, and; (4) Broader implications. The opportunities are summarized in Figure 6. There are likely more such opportunities that are not listed here, but the long list shown in the figure is a reasonable starting point for scientists and practitioners interested in working in this area.

### 4.1 Model Innovation

There are many opportunities to better model search situations and augment and adapt foundational models to better align with searchers’ tasks and goals, and provide more accurate answers. Agents can leverage these model enhancements to improve the support that they provide for complex search tasks. We now present more detail on each opportunity, including a one sentence summary (shown in italics at the beginning of each subsection).

*4.1.1 Task modeling. Build richer task models that more fully represent tasks and task contexts.* This includes how we infer tasks and intent (e.g., from textual content of search process, from user-system interactions, from other situational and contextual information such as location, time, and application usage) and how we represent those tasks internally (e.g., as a hierarchy (Figure 1) or a more abstract representation (semantic vectors, graph embeddings, Markov models, and so on)). We also need to be able to estimate key task characteristics, such as task complexity, which, in one use, can help search systems route requests to the most appropriate interaction modality. In addition, we need to find ways for agents to collect more (and more accurate) user/world knowledge, both in general and specifically related to the task at hand. A better understanding of the short- and long-term task context will help agents more accurately model the tasks themselves.

*4.1.2 Alignment. Develop methods to continuously align agents to tasks/goals/values via feedback,* e.g., conversation content as feedback (e.g., expressed positive sentiment or searchers expressing gratitude to the agent) or explicit feedback on agent answers via likes and dislikes. The performance of agents that are missing alignment will remain fixed over time. Agents need application-aligned feedback loops to better understand searcher goals and tasks and use that feedback to continuously improve answer accuracy and relevance. Beyond research on fine-tuning foundation models from human feedback (e.g., likes/dislikes) [80], we can also build on learnings from research on implicit feedback in IR, including work on improving ranking algorithms via SERP clicks [27] and developing specialized interfaces to capture user feedback [69].

*4.1.3 Augmentation. Enhance agents with relevant external knowledge and enhanced tools and capabilities.* As mentioned earlier, RAG is a common form of knowledge injection for foundation models. Relevance models are tuned to maximize user benefit, not for agent consumption. We need to evaluate whether this difference is meaningful practically and if so, develop new ranking criteria that consider the intended consumer of the search results (human or machine). Despite their incredible capabilities, foundation models still have shortcomings that manifest in the agents that use them. We need to understand these shortcomings through principled evaluation and find ways to leverage external skills/plugins to address them. Agents must find and recommend skills per task demands [63], e.g., invoking Wolfram for computational assistance. We can also integrate tool use directly into tool-augmented models, e.g., *Toolformer* [44], that can teach themselves to use tools. Models of task context may also be incomplete and we should invest in ways to better *ground* agent responses via context, e.g., richer sensing, context filtering, and dynamic prompting.

*4.1.4 Grounding. Apply use-case specific information to reduce hallucinations, build trust, and support content creators.* It is in the interests of agents, searchers, and content creators (and providers and advertisers) to consider the source of the data used in generating answers. Provenance is critical and agents should provide links back to relevant sources (preferably with specific details/URLs not generalities/domains) to help establish and maintain user trust, provide attribution for content creators, and drive engagement for content providers and advertisers. It also important for building trust and for supporting learning for agents to practice faithful reasoning [14], and provide interpretable reasoning traces (e.g., explanations with chain-of-thought) associated with their answers. We should also consider how we integrate search within existing experiences (e.g., in other agents) to ground answers in their context of use and in more places that people seek those answers.

*4.1.5 Personalization. Develop personal agents that can understand searchers and their tasks, using personal data, privately and securely.* Searchers bring their personal tasks to search systems and agents will be no different. Here are some example personal prompts that describe the types of personal tasks that searchers might expect an agent to handle: (1) Write an e-mail to my client in my personal style with a description of the quote in the attached doc. (2) Tell me what’s important for me to know about the company town hall that I missed? (3) Where should I go for lunch today? These tasks

span creation, summarization, and recommendation and quickly illustrate the wide range of expectations that people may have from their *personal* agents. As part of developing such personalized AI support, we need to: (1) Study foundation model *capabilities*, including their ability to identify task-relevant information in personal data and activity histories, and model user knowledge in the current task and topic, and (2) Develop core *technologies*, including infinite memory, using relevant long-term activity (in IR, there has been considerable research on relevant areas such as re-finding [56] and personalization [51]); context compression, to fit more context into finite token limits (e.g., using turn-by-turn summarization rather than raw conversational content); privacy, including mitigations such as differential privacy and federated learning, and research on machine unlearning [7] to intentionally forget irrelevant information over time, including sensitive information that the searcher may have explicitly asked to be removed from trained models, and also remove irrelevant or unwanted data from agent memory.

**4.1.6 Adaptation.** Two main forms of adaptation that we consider here are model specialization and so-called adaptive computation.

- *Model specialization. Develop specialized foundation models for search tasks that are controllable and efficient.* Large foundation models are generalists and have a wide capability surface. Specializing these models for specific tasks and applications discards useless knowledge, making the models more accurate and efficient for the task at hand. Recent advances in this area have yielded strong performance, e.g., the Orca-13B model [38] uses explanation-based tuning (where the model explains the steps used to achieve its output and those explanations are used to train a small language model (SLM)) to outperform state-of-the-art models of a similar size such as Vicuna-13B [13]. Future work could explore guiding specialization via search data, including anonymized large-scale search logs, and as well as algorithmic advances in searcher preference modeling and continual learning.
- *Adaptive computation. Develop methods to adaptively apply different models per task and application demands.* Adaptive compute involves using multiple foundation models (e.g., GPT and a specialized model) each with different inference-time constraints, primarily around speed, capabilities, and cost, and learning which model to apply for a given task. The specialized model can back-off to one or more larger models as needed per task demands. The input can be the task plus the constraints of the application scenario under which the model must operate. Human feedback can also be used to refine the adaptation strategy over time [77].

These adaptation methods will yield more effective and more efficient AI capabilities that agents can use to help searchers across a range of settings, including in offline settings (e.g., on-device only).

**4.1.7 Multi-Agent.** *Utilize multiple specialized agents working together and with humans to help complete a search task.* Multiple agents have been shown to help encourage divergent thinking [35], improve factuality and reasoning [17], and provide guardrails for AI systems [72]. Multi-agent systems such as *AutoGen* [71] and *CAMEL* [31] orchestrate communication between agents to help users complete tasks more effectively. These systems could be used in search settings to, for example, retrieve relevant resources from

diverse sources, improve answer correctness by critiquing and refining AI-generated output, improve human decision making by presenting alternative solutions, and even automate the completion of some tasks or sub-tasks, with humans in the loop throughout.

## 4.2 Next-Generation Experiences

Advancing models is necessary but not sufficient given the central role that interaction plays in the search process [61]. There are many opportunities to develop new search experiences that capitalize on agent capabilities while keeping searchers in control.

**4.2.1 Multimodality.** *Develop experiences bridging (at least) the search and agent (chat) modalities, offering explanations and suggestions.* Given how entrenched and popular traditional search is, it is likely that some form of query-result interaction will remain a core part of how we find information online. Future, agent-enhanced experiences may reflect a more seamless combination of the two interaction modalities in a unified experience. Both Google and Bing are taking a step in that direction by unifying search results and agent answers in a single interface. Explanations on which modality and style (e.g., creative, balanced, and precise) perform best and when will help searchers make decisions about which modalities and settings to use. Modality recommendation given task is also worth exploring: simple tasks may only need traditional search, whereas complex tasks may need agents. Contextualization and personalization will also play an important part in deciding how much information is needed from the searcher (incurring interaction cost but yielding greater control) and how much can be reliably inferred from signals already available to the system. Related to this are opportunities around conversation style suggestion given the current task, e.g., fact-finding task or short reply (needs precision, when generative-AI-powered agents can often be verbose) and generating new content (needs creativity, when agent responses can often be unoriginal or bland). Search providers could also consider offering a single point of entry and an automatic routing mechanism to direct requests to the correct modality given inferences about the underlying task (e.g., from Section 4.1.1) and the appropriateness of each of the modalities for that task. Beyond search and chat, other modalities to help support complex search tasks may include third-party tools and applications, bespoke user interfaces (e.g., tailored dynamically by the agent to the task at hand), interactive visualizations, and proactive recommendations.

**4.2.2 Human Learning.** *Develop agents that can detect learning-related search tasks and support relevant learning activities.* As mentioned earlier, agents can remove or change human learning opportunities by their automated generation and provision of answers. Learning is a core outcome of information seeking [15, 36, 58]. We need to develop agents that can detect learning and sensemaking tasks, and support relevant learning activities via agent experiences that, for example, provide detailed explanations and reasoning, offer links to learning resources (e.g., instructional videos), enable deep engagement with task content (e.g., via relevant sources), and support specifying and attaining learning objectives.

**4.2.3 Human Control.** *Better understand control and develop agents with control while growing automation.* Control is an essential aspect of searcher interaction with agents. Agents should consult

humans to resolve or codify value tensions. Agents should be in collaboration mode by default and must only take control with the permission of stakeholders. Experiences that provide searchers with more agency are critical, e.g., adjust specificity/diversity in agent answers, leading to less generality and less repetition. As mentioned in Section 4.1.4, citations in answers are important. Humans need to be able to verify citation correctness in a lightweight way, ideally without leaving the user experience; Gemini now offers an ability to manually dive deeper and verify answers. We also need a set of user studies to understand the implications of less control of some aspects (e.g., answer generation), more control over other aspects (e.g., macrotask specification), and control over new aspects, such as conversation style and tone, as with Copilot.

**4.2.4 Completion.** *Agents should help searchers complete tasks while keeping searchers in control.* We need to both expand the search frontier by adding/discovering more capabilities of foundation models that can be surfaced through agents *and* deepen task capabilities so that agents can help searchers better complete more tasks. We can view skills and plugins as actuators of the digital world and we should help foundation models fully utilize them. We need to start simple (e.g., reservations), learn and iterate, and increase task complexity as model capabilities improve with time.

The standard mode of engagement with AI agents is reactive; users send requests and the agents respond. Agents should ideally have a *dynamic interaction model* that tailors the interface to the task and the context. They can take initiative, with permission, and provide updates for standing tasks, or they could also offer proactive suggestions or take actions directly when agent uncertainty is low. Agents can also help to support task planning (decomposition, prioritization, and so on) for complex tasks such as travel or events. AI can already help complete repetitive tasks, e.g., action transformers, trained on digital tools<sup>14</sup> and create and apply “tasklets” (user interface scripts) learned from websites [33].

Given the centrality of information interaction in search task completion, it is important to focus sufficient attention on interaction models and experiences in AI agents. In doing so, we must also carefully consider the implications of critical decisions on issues that affect AI in general such as control and automation.

### 4.3 Measurement

Another important direction is in measuring AI agent performance, understanding agent impact and capabilities, and tracking agent evolution over time. Many of the challenges and opportunities in this area also affect the evaluation of foundation models in general (e.g., non-determinism, saturated benchmarks, inadequate metrics).

**4.3.1 Understanding.** *Deeply understand agent capabilities and agent impact on searchers and on their tasks.* We have only scratched the surface in understanding AI agents and their impact. A deeper understanding takes a few forms, including: (1) *User understanding:* Covering mental models of agents and effects of bias (e.g., functional fixedness [18]) on how agents are adopted and used in search settings. It also covers changes in search behavior and information seeking strategies, including measuring changes in effects across modalities, e.g., search *versus* agents and search *plus* agents. There

are also opportunities in using foundation models to understand search interactions via user studies [12] and use foundation models to generate intent taxonomies and classify intents from log data [47]; (2) *Task understanding:* Covering the intents and tasks that agents are used for and most effective for, and; (3) *Agent understanding:* Covering the capabilities and limitations of agents, e.g., similar to the “Sparks of AGI” paper on GPT-4 [10], which examined foundation model capabilities in depth.

**4.3.2 Evaluation.** *Identify and develop metrics for agent evaluation, while considering important factors, and find applications of agent components for IR evaluation.* There are many options for AI agent metrics, including feedback, engagement, precision-recall, generation quality, answer accuracy, and so on. Given the task focus, metrics should likely target the task holistically (e.g., success, effort, satisfaction). In evaluating agents in search settings, it is also important to consider: (1) Repeatability: Non-determinism can make agents difficult to evaluate/debug; (2) Interplay between search and agents (switching, joint task success, and so on); (3) Longer term effects on user capabilities and productivity; (4) Task characteristics: Complexity, and so on, and; (5) New benchmarks: Agents are affected by external data, grounding, queries, and so on. There are also opportunities to consider applications of agent components for IR evaluation. Foundation models can predict searcher preferences [54] and assist with relevance judgments [20], including generating explanations for judges. Also, foundation models can create powerful searcher simulations that can mimic human behavior and values, expanding on early work on searcher simulations in IR [70].

Measuring agent performance is essential in understanding their utility and improving their performance over time. Agents do not function in a vacuum and we must consider the broader implications of their deployment for complex tasks in search settings.

## 4.4 Broader Implications

AI agents must operate in a complex and dynamic world. There are several opportunities beyond advances in technology and in deepening our understanding of agent performance and capabilities.

**4.4.1 Responsibility.** *Understand factors affecting reliability, safety, fairness, and inclusion in agent usage in search.* The broad reach of search engines means that AI agents have a critical obligation to act responsibly. Research is needed on ways to improve answer accuracy via better grounding in more reliable data sources, develop guardrails, understand biases in foundation models, prompts, and the data used for grounding, and understand how well agents work in different contexts, with different tasks, and with different people and cohorts. Red teaming, user testing, and feedback loops are all needed to determine emerging risks in agents and the foundation models that underlie them. This also builds on existing work on responsible AI and responsible IR and FACTS-IR, which has studied biases and harms, and ways to mitigate them [40].

**4.4.2 Economics.** *Understand and expand the economic impact of agents in search.* This includes exploring new business models which agents will create beyond information finding. Advancing the search frontier from information finding deeper into task completion (e.g., into creation and analysis) creates new business opportunity. It also unlocks new opportunities for advertising, including

<sup>14</sup><https://www.adept.ai/blog/act-1>



advertisements that are shown inline with dialog/answers and contextually relevant to the current conversation. There is also a need to more deeply understand the impact of agents on content creation and search engine optimization. Content attribution is vital in such scenarios to ensure that content creators (and advertisers and publishers) can still generate returns. We should avoid the so-called “paradox of reuse” [59] where lower visits to online content leads to less content being created which in turn leads to worse models over time. Another important aspect of economics is the cost-benefit trade-off and is related to work on adaptation (Section 4.1.6). Large model inference is expensive and unnecessary for many applications. This cost will reduce with optimization, for which model specialization and adaptive computation can help, as does the emergence of high-performing SLMs of a few billion parameters, such as Phi, trained on highly-curated data [22].

**4.4.3 Ubiquity. Agent integrations to model and support complex search tasks.** AI agents must co-exist with the other parts of the application ecosystem. Search agents can be integrated into applications such as Web browsers (offering in-browser chat, editing assistance, summarization) and productivity applications (offering support in creating documents, emails, presentations, and so on). These agents can capitalize on application context to do a better job of answering searcher requests. Agents can also span applications through integration with the operating system. This enables richer task modeling and complex task support, since such tasks often involve multiple applications. Critically, we must do this privately and securely to mitigate risks for searchers and earn their trust.

## 4.5 Summary

The directions highlighted in this section are just examples of the opportunities afforded by the emergence of generative AI and agents in search settings. There are other areas for search providers to consider too, such as multilingual agent experiences (i.e., foundation models are powerful and could help with language translation [37, 79]), agent efficiency (i.e., large model inference is expensive and not sustainable at massive scale, so creative solutions are needed [77]), reducing the carbon impact from running foundation models at search engine scale to serve billions of answers in AI agents [19], making agents private by design [75], and government directives (e.g., the 2023 executive order from U.S. President Biden on AI safety and security<sup>15</sup>) and legislation, among many other opportunities.

## 5 THE UNDISCOVERED COUNTRY

AI agents will transform how we search. Tasks are central to people’s lives and more support is needed for complex tasks in search settings. Some limited support for these tasks already exists in search engines, but agents will advance the search frontier to make more tasks actionable and help to make inroads in the “last mile” in search interaction: task completion [62]. Moving forward, search providers should invest in “better together” experiences that utilize agents plus traditional search (plus more modalities going forward), make these joint experiences more seamless for searchers, and add more support for their use in practice, e.g., help people to quickly

understand agent capabilities and potential and/or recommend the best modality for the current task or task stage. This also includes experiences where both modalities are offered separately and can be selected by searchers and those where there is unification and the selection happens automatically based on the task and the conversation context.

The foundation models that power AI agents have other search-related applications, e.g., for generating and applying intent taxonomies [47] or for evaluation [20]. We must retain a continued focus on human-AI cooperation, where searchers stay in control while the degree of system support increases as needed [48], and on AI safety and security. Searchers need to be able to trust agents but also be able to verify their answers with minimal effort.

Overall, the future is bright for IR, and AI research in general, with the advent of generative AI and the agents that build upon it. Agents will help augment, empower, and inspire searchers on their task journeys. Computer science researchers and practitioners should embrace this new era of assistive AI agents and engage across the full spectrum of exciting practical and scientific opportunities, both within search as we focused on in this article, and onwards into other important domains such as personal productivity [5] and scientific discovery [24].

## REFERENCES

- [1] Eugene Agichtein, Ryan W White, Susan T Dumais, and Paul N Bnet. 2012. Search, interrupted: understanding and predicting search task continuation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. 315–324.
- [2] Marcia J Bates. 1990. Where should the person stop and the information search interface start? *Information Processing & Management* 26, 5 (1990), 575–591.
- [3] Nicholas J Belkin. 1980. Anomalous states of knowledge as a basis for information retrieval. *Canadian journal of information science* 5, 1 (1980), 133–143.
- [4] Paul N Bennett, Ryan W White, Wei Chu, Susan T Dumais, Peter Bailey, Fedor Borisjuk, and Xiaoyuan Cui. 2012. Modeling the impact of short-and long-term behavior on search personalization. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. 185–194.
- [5] Christian Bird, Denae Ford, Thomas Zimmermann, Nicole Forsgren, Eirini Kalliamvakou, Travis Lowdermilk, and Idan Gazit. 2022. Taking Flight with Copilot: Early insights and opportunities of AI-powered pair-programming tools. *Queue* 20, 6 (2022), 35–57.
- [6] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [7] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 141–159.
- [8] Andrei Broder. 2002. A taxonomy of web search. In *ACM Sigir forum*, Vol. 36. ACM New York, NY, USA, 3–10.
- [9] Andrei Z Broder and Preston McAfee. 2023. Delphic Costs and Benefits in Web Search: A utilitarian and historical analysis. *arXiv preprint arXiv:2308.07525* (2023).
- [10] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023).
- [11] Katriina Byström and Kalervo Järvelin. 1995. Task complexity affects information seeking and use. *Information processing & management* 31, 2 (1995), 191–213.
- [12] Robert Capra and Jaime Arguello. 2023. How does AI chat change search behaviors? *arXiv preprint arXiv:2307.03826* (2023).
- [13] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Liannin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023) (2023).
- [14] Antonia Creswell and Murray Shanahan. 2022. Faithful reasoning using large language models. *arXiv preprint arXiv:2208.14271* (2022).
- [15] Brenda Dervin. 1998. Sense-making theory and practice: An overview of user interests in knowledge seeking and use. *Journal of knowledge management* 2, 2

<sup>15</sup><https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>

- (1998), 36–46.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325* (2023).
- [18] Karl Duncker and Lynne S Lees. 1945. On problem-solving. *Psychological monographs* 58, 5 (1945), i.
- [19] Brad Everman, Trevor Villwock, Dayuan Chen, Noe Soto, Oliver Zhang, and Ziliang Zong. 2023. Evaluating the Carbon Impact of Large Language Models at the Inference Stage. In *2023 IEEE International Performance, Computing, and Communications Conference (IPCCC)*. IEEE, 150–157.
- [20] Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, et al. 2023. Perspectives on Large Language Models for Relevance Judgment. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*. 39–50.
- [21] Jianfeng Gao, Chenyan Xiong, Paul Bennett, and Nick Craswell. 2023. *Neural Approaches to Conversational Information Retrieval*. Vol. 44. Springer Nature.
- [22] Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644* (2023).
- [23] Ahmed Hassan Awadallah, Ryen W White, Patrick Pantel, Susan T Dumais, and Yi-Min Wang. 2014. Supporting complex search tasks. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management*. 829–838.
- [24] Tom Hope, Doug Downey, Daniel S Weld, Oren Etzioni, and Eric Horvitz. 2023. A computational inflection for scientific discovery. *Commun. ACM* 66, 8 (2023), 62–73.
- [25] Peter Ingwersen and Kalervo Järvelin. 2005. *The turn: Integration of information seeking and retrieval in context*. Vol. 18. Springer Science & Business Media.
- [26] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.
- [27] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. 133–142.
- [28] Jeonghyun Kim. 2006. Task difficulty as a predictor and indicator of web searching interaction. In *CHI'06 extended abstracts on human factors in computing systems*. 959–964.
- [29] David R Krathwohl. 2002. A revision of Bloom’s taxonomy: An overview. *Theory into practice* 41, 4 (2002), 212–218.
- [30] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, 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 (2020), 9459–9474.
- [31] Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2024. Camel: Communicative agents for “mind” exploration of large language model society. *Advances in Neural Information Processing Systems* 36 (2024).
- [32] Yuelin Li and Nicholas J Belkin. 2008. A faceted approach to conceptualizing tasks in information seeking. *Information processing & management* 44, 6 (2008), 1822–1837.
- [33] Yuanchun Li and Oriana Riva. 2021. Glider: A reinforcement learning approach to extract UI scripts from websites. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1420–1430.
- [34] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*. PMLR, 6565–6576.
- [35] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujun Yang, Zhaopeng Tu, and Shuming Shi. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118* (2023).
- [36] Gary Marchionini. 2006. Exploratory search: from finding to understanding. *Commun. ACM* 49, 4 (2006), 41–46.
- [37] James Mayfield, Eugene Yang, Dawn Lawrie, Samuel Barham, Orion Weller, Marc Mason, Suraj Nair, and Scott Miller. 2023. Synthetic Cross-language Information Retrieval Training Data. *arXiv preprint arXiv:2305.00331* (2023).
- [38] Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint arXiv:2306.02707* (2023).
- [39] Marc Najork. 2023. Generative Information Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1–1.
- [40] Alexandra Olteanu, Jean Garcia-Gathright, Maarten de Rijke, Michael D Ekstrand, Adam Roegiest, Aldo Lipani, Alex Beutel, Alexandra Olteanu, Ana Lucic, Ana-Andreea Stoica, et al. 2021. FACTS-IR: fairness, accountability, confidentiality, transparency, and safety in information retrieval. In *ACM SIGIR Forum*, Vol. 53. ACM New York, NY, USA, 20–43.
- [41] Soo Young Rieh, Kevyn Collins-Thompson, Preben Hansen, and Hye-Jung Lee. 2016. Towards searching as a learning process: A review of current perspectives and future directions. *Journal of Information Science* 42, 1 (2016), 19–34.
- [42] Shawon Sarkar and Chirag Shah. 2021. An integrated model of task, information needs, sources and uncertainty to design task-aware search systems. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*. 83–92.
- [43] Reijo Savolainen. 2012. Expectancy-value beliefs and information needs as motivators for task-based information seeking. *Journal of Documentation* 68, 4 (2012), 492–511.
- [44] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, 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* (2023).
- [45] Chirag Shah. 2023. Generative AI and the Future of Information Access. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (Birmingham, United Kingdom) (CIKM '23)*. Association for Computing Machinery, New York, NY, USA, 3. <https://doi.org/10.1145/3583780.3615317>
- [46] Chirag Shah, Ryen White, Paul Thomas, Bhaskar Mitra, Shawon Sarkar, and Nicholas Belkin. 2023. Taking search to task. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*. 1–13.
- [47] Chirag Shah, Ryen W White, Reid Andersen, Georg Buscher, Scott Counts, Sarkar Snigdha Sarathi Das, Ali Montazer, Sathish Manivannan, Jennifer Neville, Xiaochuan Ni, et al. 2023. Using Large Language Models to Generate, Validate, and Apply User Intent Taxonomies. *arXiv preprint arXiv:2309.13063* (2023).
- [48] Ben Shneiderman. 2022. *Human-centered AI*. Oxford University Press.
- [49] Adish Singla, Ryen White, and Jeff Huang. 2010. Studying trailfinding algorithms for enhanced web search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. 443–450.
- [50] Jaime Teevan, Kevyn Collins-Thompson, Ryen W White, and Susan Dumais. 2014. Slow search. *Commun. ACM* 57, 8 (2014), 36–38.
- [51] Jaime Teevan, Susan T Dumais, and Eric Horvitz. 2005. Personalizing search via automated analysis of interests and activities. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. 449–456.
- [52] Jaime Teevan, Meredith Ringel Morris, and Steve Bush. 2009. Discovering and using groups to improve personalized search. In *Proceedings of the second acm international conference on web search and data mining*. 15–24.
- [53] Maartje ter Hoeve, Robert Sim, Elnaz Nouri, Adam Fournay, Maarten de Rijke, and Ryen W White. 2020. Conversations with documents: An exploration of document-centered assistance. In *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*. 43–52.
- [54] Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. 2023. Large language models can accurately predict searcher preferences. *arXiv preprint arXiv:2309.10621* (2023).
- [55] Randall H Trigg. 1988. Guided tours and tabletops: Tools for communicating in a hypertext environment. *ACM Transactions on Information Systems (TOIS)* 6, 4 (1988), 398–414.
- [56] Sarah K Tyler and Jaime Teevan. 2010. Large scale query log analysis of re-finding. In *Proceedings of the third ACM international conference on Web search and data mining*. 191–200.
- [57] Pertti Vakkari. 2001. A theory of the task-based information retrieval process: A summary and generalisation of a longitudinal study. *Journal of documentation* 57, 1 (2001), 44–60.
- [58] Pertti Vakkari. 2016. Searching as learning: A systematization based on literature. *Journal of Information Science* 42, 1 (2016), 7–18.
- [59] Nicholas Vincent. 2022. The Paradox of Reuse, Language Models Edition. <https://mmvg.mataroa.blog/blog/the-paradox-of-reuse-language-models-edition/>. Accessed: 2023-09-12.
- [60] Yu Wang, Xiao Huang, and Ryen W White. 2013. Characterizing and supporting cross-device search tasks. In *Proceedings of the sixth ACM international conference on Web search and data mining*. 707–716.
- [61] Ryen W White. 2016. *Interactions with search systems*. Cambridge University Press.
- [62] Ryen W White. 2018. Opportunities and challenges in search interaction. *Commun. ACM* 61, 12 (2018), 36–38.
- [63] Ryen W White. 2018. Skill discovery in virtual assistants. *Commun. ACM* 61, 11 (2018), 106–113.
- [64] Ryen W White. 2022. Intelligent futures in task assistance. *Commun. ACM* 65, 11 (2022), 35–39.
- [65] Ryen W. White. 2023. Tasks, Copilots, and the Future of Search. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (Taipei, Taiwan) (SIGIR '23)*. Association for Computing Machinery, New York, NY, USA, 5–6. <https://doi.org/10.1145/3539618.3593069>

- [66] Ryen W White, Mikhail Bilenko, and Silviu Cucerzan. 2007. Studying the use of popular destinations to enhance web search interaction. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. 159–166.
- [67] Ryen W White, Wei Chu, Ahmed Hassan, Xiaodong He, Yang Song, and Hongning Wang. 2013. Enhancing personalized search by mining and modeling task behavior. In *Proceedings of the 22nd international conference on World Wide Web*. 1411–1420.
- [68] Ryen W White, Adam Fourney, Allen Herring, Paul N Bennett, Nirupama Chandrasekaran, Robert Sim, Elnaz Nouri, and Mark J Encarnación. 2019. Multi-device digital assistance. *Commun. ACM* 62, 10 (2019), 28–31.
- [69] Ryen W White, Ian Ruthven, and Joemon M Jose. 2005. A study of factors affecting the utility of implicit relevance feedback. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. 35–42.
- [70] Ryen W White, Ian Ruthven, Joemon M Jose, and CJ Van Rijsbergen. 2005. Evaluating implicit feedback models using searcher simulations. *ACM Transactions on Information Systems (TOIS)* 23, 3 (2005), 325–361.
- [71] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023. AutoGen: Enabling next-gen LLM applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155* (2023).
- [72] Yiran Wu, Feiran Jia, Shaokun Zhang, Qingyun Wu, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, and Chi Wang. 2023. An empirical study on challenging math problem solving with gpt-4. *arXiv preprint arXiv:2306.01337* (2023).
- [73] Iris Xie. 2008. *Interactive information retrieval in digital environments*. IGI global.
- [74] Jinyun Yan, Wei Chu, and Ryen W White. 2014. Cohort modeling for enhanced personalized search. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 505–514.
- [75] Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. 2021. Differentially private fine-tuning of language models. *arXiv preprint arXiv:2110.06500* (2021).
- [76] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating clarifying questions for information retrieval. In *Proceedings of the web conference 2020*. 418–428.
- [77] Jieyu Zhang, Ranjay Krishna, Ahmed H Awadallah, and Chi Wang. 2023. EcoAssistant: Using LLM Assistant More Affordably and Accurately. *arXiv preprint arXiv:2310.03046* (2023).
- [78] Yi Zhang, Sujay Kumar Jauhar, Julia Kiseleva, Ryen White, and Dan Roth. 2021. Learning to decompose and organize complex tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2726–2735.
- [79] Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Lingpeng Kong, Jiajun Chen, Lei Li, and Shujian Huang. 2023. Multilingual machine translation with large language models: Empirical results and analysis. *arXiv preprint arXiv:2304.04675* (2023).
- [80] Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593* (2019).