

Deep Research Agents: A Systematic Examination And Roadmap

Anonymous authors
Paper under double-blind review

Abstract

The rapid progress of Large Language Models (LLMs) has given rise to a new category of autonomous AI systems, referred to as Deep Research (DR) agents. These agents are designed to tackle complex, multi-turn informational research tasks by leveraging a combination of dynamic reasoning, adaptive long-horizon planning, multi-hop information retrieval, iterative tool use, and the generation of structured analytical reports. In this paper, we conduct a detailed analysis of the foundational technologies and architectural components that constitute Deep Research agents. We begin by reviewing information acquisition strategies, contrasting API-based retrieval methods with browser-based exploration. We then examine modular tool-use frameworks, including code execution, multimodal input processing, and the integration of Model Context Protocols (MCPs) to support extensibility and ecosystem development. To systematise existing approaches, we propose a taxonomy that differentiates between static and dynamic workflows, and we classify agent architectures based on planning strategies and agent composition, including single-agent and multi-agent configurations. We also provide a critical evaluation of current benchmarks, highlighting key limitations such as restricted access to external knowledge, sequential execution inefficiencies, and misalignment between evaluation metrics and the practical objectives of DR agents. Finally, we outline open challenges and promising directions for future research.

1 Introduction

Recent advances in large language models (LLMs) have led to the rapid emergence of sophisticated AI agents capable of autonomous research. Early models such as GPT-3 (Brown et al., 2020) primarily addressed isolated tasks, including question answering and machine translation. Subsequently, integration with external tools enabled models such as WebGPT (Nakano et al., 2021) to navigate the web and synthesise information from diverse sources autonomously. Most recently, a new class of advanced autonomous systems, termed Deep Research (DR) agents, has emerged, exemplified by industry-leading solutions such as OpenAI DR (OpenAI, 2025b), Gemini DR (Google Team, 2025), Grok DeepSearch (xAI Team, 2025), and Perplexity DR (Perplexity Team, 2025). These deep research agents significantly extend LLMs by incorporating advanced reasoning, dynamic task planning, and adaptive interaction with web resources and analytical tools.

Formally, we define “**Deep Research Agents**” as:

AI agents powered by LLMs, integrating dynamic reasoning, adaptive planning, and iterative tool use to acquire, aggregate, and analyse external information, culminating in comprehensive outputs for accomplishing open-ended informational research tasks.

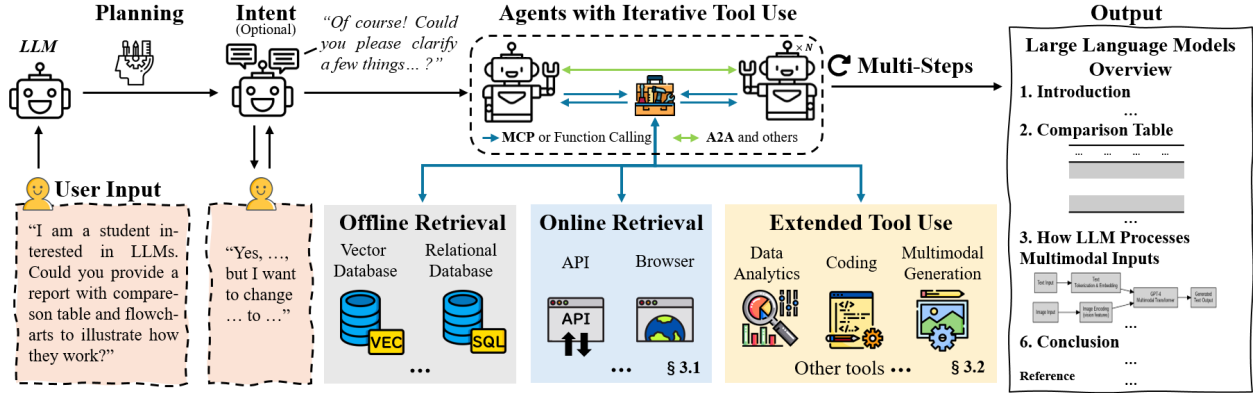


Figure 1: A structural overview of a DR agent in a multi-agent architecture for ease of illustration.

Specifically, DR agents leverage LLMs as their cognitive core, retrieving external knowledge in real-time through web browsers and structured APIs, and dynamically invoking analytical tools via customised toolkits or standardised interfaces such as the Model Context Protocol (MCP). This architecture enables DR agents to autonomously manage complex, end-to-end research workflows by seamlessly integrating reasoning processes with multimodal resources.

Compared with traditional Retrieval-Augmented Generation (RAG) methods (Singh et al., 2025), which primarily enhance factual accuracy but lack sustained reasoning capabilities (Chen et al., 2025c), and conventional Tool Use (TU) systems (Qu et al., 2025) that heavily depend on pre-defined workflows (Wang et al., 2025b), DR agents offer significantly greater autonomy, continual and deep reasoning abilities, dynamic task planning, and adaptive real-time interaction. These advanced capabilities uniquely position DR agents to handle complex, evolving, and knowledge-intensive research scenarios. A representative example of such a DR agent architecture is illustrated in Figure 1, which demonstrates the complete workflow from user input through optional planning and intent clarification, to iterative tool utilization encompassing offline retrieval (vector and relational databases), online retrieval (APIs and browsers), and extended capabilities including data analytics, coding (etc.), and multimodal generation, ultimately producing comprehensive structured report.

Contribution. This survey systematically reviews recent advancements in DR agents, providing a comprehensive analysis of core technologies, methodologies, optimisation pipelines, and representative implementations. Specifically, the contributions of this survey include:

- A thorough analysis of representative DR systems, explicitly examining their system architectures, retrieval mechanisms, tool invocation methods, and performance characteristics, alongside optimisation and tuning paradigms.
- A unified classification framework (Figure 4) that systematically categorises DR systems based on workflow characteristics (static versus dynamic), planning strategies, and agent-based architectures (single-agent versus multi-agent), bridging diverse technical methodologies and current industrial solutions.

- A systematic review and categorisation of existing benchmarks utilised to evaluate DR systems, highlighting how these benchmarks assess critical capabilities, such as retrieval accuracy, reasoning depth, and adaptive tool invocation proficiency.
- A systematic analysis of critical open challenges and research directions, focusing on expanding retrieval scope beyond traditional methods, enabling asynchronous parallel execution, developing comprehensive multi-modal benchmarks, and optimising multi-agent architectures for enhanced robustness and efficiency.

Survey Organization. This survey methodically explores recent advancements in DR agents, organised as follows: Section 2 provides foundational concepts, examining recent progress in reasoning, retrieval-augmented generation, and agent communication protocols. Section 3 comprehensively analyses key DR agent components, including search engine integration (Section 3.1), tool invocation strategies (Section 3.2), architectural workflows (Section 3.3), and optimisation methodologies (Section 3.4). Section 4 reviews major industrial applications and practical implementations of DR agents by leading organisations. Section 5 surveys benchmarks used for evaluating DR systems, categorising them into question-answering and task execution scenarios. Section 6 highlights critical challenges and outlines promising directions for future research, focusing on enhancing information acquisition, asynchronous parallel execution, benchmark alignment, and optimising multi-agent architectures. Finally, Section 7 concludes with a summary and provides insights into the broader implications and opportunities within DR agent research.

2 Background and Preliminaries

2.1 Advances in Reasoning and Tool Integration

Recent advancements in large reasoning models (LRMs) have greatly enhanced the ability of language models to tackle complex and abstract tasks. These models have shown significant improvements in tasks such as arithmetic, common-sense reasoning, and symbolic problem-solving, largely due to innovations in model architectures and training techniques. One such advancement is Chain-of-Thought (CoT) prompting, introduced by Wei et al. (Wei et al., 2023), which explicitly guides models to articulate intermediate logical steps, decomposing complex problems into simpler, sequential stages. This has led to notable improvements in both the interpretability and accuracy of LLMs on various reasoning benchmarks. Building upon CoT, subsequent research has introduced methods to further enhance LLM reasoning, particularly in handling lengthy textual contexts. Approaches such as positional interpolation and sparse attention mechanisms (Bai et al., 2024a; Wang et al., 2024c) have been proposed to extend the effective context window. Furthermore, specialised benchmarks like LongBench (Bai et al., 2024b) and LongFinanceQA (Lin et al., 2025) have been developed to rigorously evaluate and improve the performance of these models in extended-context reasoning.

To address reasoning tasks that require real-time or specialised external knowledge, frameworks like Toolformer (Schick et al., 2023) and MultiTool-CoT (Inaba et al., 2023) have been proposed, enabling LLMs to autonomously incorporate external computational resources and APIs directly within reasoning workflows. These approaches effectively enhance performance in tasks dependent on precise numerical calculations and dynamic information retrieval. Maintaining reasoning coherence across multiple conversational turns also poses distinct challenges. Techniques such as Dialogue CoT (Chae et al., 2023) and Structured CoT (SCoT) (Sultan et al., 2024) explicitly integrate dialogue states and conversational contexts within reasoning chains, significantly improving coherence, context-awareness, and the ability to manage iterative interactions and clarify complex user queries. However, despite substantial improvements, existing reasoning frameworks still

encounter critical issues, including hallucinations, static or outdated internal knowledge, and insufficient responsiveness to rapidly changing information needs. These limitations highlight the necessity of integrating external information sources, real-time retrieval mechanisms, and adaptive reasoning strategies, which are core motivations driving recent advances toward more comprehensive and robust reasoning frameworks suitable for DR Agent applications.

2.2 Advances in Retrieval-Augmented Generation and Agentic Retrieval

Retrieval-augmented Generation (RAG), leveraging external knowledge bases (e.g., webs, APIs), has emerged as an effective strategy to mitigate hallucination problems and enhance the accuracy of web information search (Fan et al., 2024; Gao et al., 2023; Singh et al., 2025). Early RAG architectures typically involved a static pipeline, where retrievers fetched relevant documents from external sources such as Wikipedia or search engines, and generators (e.g., LLMs) produced answers based solely on these retrieved passages. However, static approaches were limited in handling complex or multi-step queries, motivating recent advances toward iterative and interactive retrieval mechanisms to generate richer and more relevant responses, including FLARE (Zhang et al., 2024), Self-RAG (Asai et al., 2023), IAG (Zhang et al., 2023), and ToC (Kim et al., 2023). In addition, studies (Izacard et al., 2023; Lin et al., 2023) expanded retrieval sources from structured databases (e.g., Wikipedia) to large-scale, diverse web corpora such as the Common Crawl dump preprocessed via the CCNet pipeline (Fu et al., 2022). Further improvements of RAG include hybrid approaches that combine internal LLM knowledge and external retrievals for better accuracy and coherence (Aliannejadi et al., 2024). Recently, Huang et al. (Huang et al., 2025) proposed RAG-RL, introducing reinforcement learning and curriculum learning techniques, enabling reasoning language models (RLMs) to more effectively identify and utilise relevant contexts.

Despite these advancements in retrieval methods and reasoning-enhanced models, RAG approaches still face limitations in effectively managing complex reasoning workflows and dynamically adapting to varied task requirements. To address these challenges, recent research extends RAG into an agentic paradigm, integrating additional reasoning and decision-making layers atop conventional RAG pipelines (Singh et al., 2025). Agentic RAG approaches leverage iterative retrieval, adaptive querying, and dynamic workflow adjustments, significantly enhancing multi-step reasoning capabilities. For example, RL-based query refinement techniques (e.g., Hsu et al. (Hsu et al., 2024)) improve retrieval for complex queries, while graph-based retrieval (e.g., GeAR (Shen et al., 2024)) further enhances the processing of multi-hop queries. Despite these advancements, agentic RAG still faces critical challenges, including balancing computational overhead from dynamic reasoning processes (Singh et al., 2025), aligning agent behaviours with user intentions (Zerhoubi & Granitzer, 2024), and ensuring interpretability in adaptive workflows (Hsu et al., 2024; Singh et al., 2025). Moreover, even advanced agentic RAG approaches remain constrained by their reliance on pre-existing or periodically updated corpora, limiting their ability to handle real-time, rapidly changing, or long-tail information needs effectively. Addressing this challenge requires integrating external APIs and web browsing capabilities into RAG architectures, motivating recent DR methods aimed at further enhancing retrieval comprehensiveness and adaptability.

2.3 Model Context Protocol and Agent-to-Agent Policy

Model Context Protocol (MCP) and Agent-to-Agent (A2A) have been proposed to address interoperability challenges in LLM-based agent systems, enabling efficient tool access and effective multi-agent collaboration. **MCP:** Traditional Tool Use (TU) agents face significant challenges, including inconsistent APIs, high maintenance costs, and redundant development efforts, severely limiting interoperability across systems

(Schick et al., 2023). To address these issues, Anthropic introduced the MCP, a unified communication layer allowing LLM-based agents to interact securely and consistently with external services and data sources via standardised interfaces. MCP mitigates data silo problems by providing dynamic service discovery and uniform access patterns. **A2A**: Google’s A2A protocol facilitates decentralised multi-agent collaboration through structured, task-oriented dialogues. Agents from diverse vendors and model architectures can discover peers, delegate responsibilities, and collaboratively manage complex tasks as equal participants (Google, 2025). By abstracting agent discovery into Agent Cards and task coordination into Tasks and Artefacts, A2A supports flexible, incremental, multi-modal workflows, ideally suited to sophisticated collaborative scenarios.

MCP and A2A complement each other by clearly separating responsibilities: MCP serves as a standardised interface for accessing external tools, while A2A orchestrates collaborative agent interactions. Together, they establish a modular and scalable foundation for open, interoperable agent ecosystems, significantly enhancing the practical capabilities of AI systems in tackling complex real-world challenges.

3 Deep Research: Search Engine, Tool Use, Workflow, Tuning, Non-parametric Continual Learning

Comparison with Conventional RAG-based Approaches. DR agents expand the capabilities of traditional RAG methods by integrating dynamic retrieval, real-time TU, and adaptive reasoning into a unified system. RAG-based approaches typically rely on fixed pipelines, limiting their flexibility in handling complex, multi-step queries or rapidly changing contexts. In contrast, DR agents provide greater autonomy, context-awareness, and accuracy by dynamically engaging with external tools and managing multi-stage research tasks in real time.

In this section, we explore five core components essential for the development and optimization of DR agents: (3.1) **search engine integration**, which compares API-based interfaces with browser-based exploration to enhance dynamic knowledge acquisition; (3.2) **Tool Use capabilities**, which investigate the integration of code execution, mathematical computation, file manipulation, and multimodal processing modules within the agent’s inference pipeline; (3.3) **workflow architecture**, analysing foundational designs, the balance between multi-agent and single-agent paradigms, memory mechanisms, and auxiliary components that facilitate the orchestration of complex research workflows; (3.4) **tuning methodologies**, which examine prompt-driven structured generation, LLM-driven prompting, fine-tuning strategies, and reinforcement learning approaches aimed at optimizing agent performance, and (3.5) **Non-parametric continual learning**, which enables LLM agents to self-evolve by dynamically adapting external tools, memory, and workflows without updating internal model weights, offering scalable optimization for complex tasks.

3.1 Search Engine: API vs. Browser

To enhance reasoning depth and accuracy for handling evolving tasks, DR agents employ search engines (SE) to update their knowledge through interaction with the external environment. In Table 1, we present a comparative overview of SEs, base models, and evaluation benchmarks employed by existing DR agents. The SEs can be broadly categorised into two types:

- 1) **API-Based SEs**, which interact with structured data sources, such as search-engine APIs or scientific database APIs, enabling efficient retrieval of organised information.
- 2) **Browser-Based SEs**, which simulate human-like interactions with web pages, facilitating real-time extraction of dynamic or unstructured content, improving the comprehensiveness of the external knowledge.

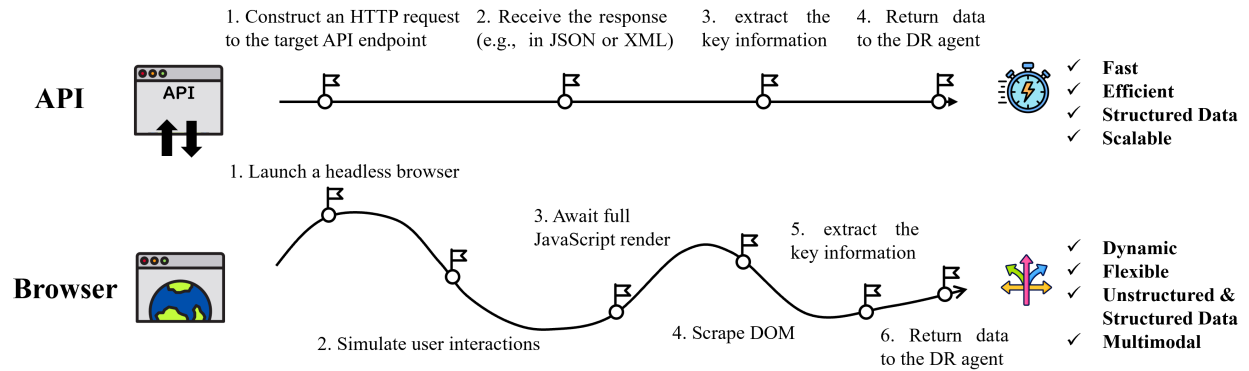


Figure 2: General Comparison of API-Based and Browser-Based Retrieval Workflow.

Table 1: Comparison of DR Agents with Search Engine Details

■ = Primary focus, ■ = Secondary/minor focus, □ = Not present

DR Agent	Search Engine		Benchmark		Base Model	Release
	API	Browser	GAIA	HLE		
Avatar Wu et al. (2024)	■	□	□	□	Stark	Feb-2024
CoSearch-Agent Gong et al. (2024)	■	□	□	□	□	GPT-3.5-turbo Feb-2024
MMAC-Copilot Song et al. (2024)	■	□	■	□	□	GPT-3.5, GPT-4 Mar-2024
Storm Shao et al. (2024a)	■	□	□	□	FreshWiki	GPT-3.5-turbo Jul-2024
OpenResearcher Zheng et al. (2024)	■	□	□	□	Privately Collected QA Data	DeepSeek-V2-Chat Aug-2024
The AI Scientist Lu et al. (2024)	■	□	□	□	MLE-Bench	GPT-4o, o1-mini, o1-preview Aug-2024
Gemini DR Google Team (2025)	■	■	□	■	GPQA	Gemini-2.0-Flash Dec-2024
Agent Laboratory Schmidgall et al. (2025)	■	□	□	□	MLE-Bench	GPT-4o, o1-preview Jan-2025
Search-o1 Li et al. (2025c)	■	□	□	□	GPQA, NQ, TriviaQA	QwQ-32B-preview Jan-2025
WebWalker Wu et al. (2025b)	□	□	□	□	WebWalkerQA	GPT-4o, Qwen-2.5 Jan-2025
Agentic Reasoning Wu et al. (2025c)	■	□	□	□	GPQA	DeepSeek-R1, Qwen2.5 Feb-2025
AutoAgent Tang et al. (2025a)	□	■	■	□	□	Claude-Sonnet-3.5 Feb-2025
Grok DeepSearch xAI Team (2025)	■	■	□	□	GPQA	Grok3 Feb-2025
OpenAI DR OpenAI (2025b)	□	■	■	■	■	GPT-o3 Feb-2025
Perplexity DR Perplexity Team (2025)	■	■	□	■	SimoleQA	Flexible Feb-2025
Towards an AI co-scientist Gottweis et al. (2025)	■	□	□	□	GPQA	Gemini 2.0 Feb-2025
Nouswise Nouswise (2025)	□	□	□	□	□	— Mar-2025
AgentRxiv Schmidgall & Moor (2025)	■	□	□	□	GPQA, MedQA	GPT-4o-mini Mar-2025

Continued on next page

Table 1 continued from previous page

DR Agent		Search Engine		Benchmark		Other QA	Base Model	Release
		API	Browser	GAIA	HLE			
Agent-R1	Ouyang et al. (2025)	■	□	□	□	HotpotQA	Qwen2.5-1.5B-Inst	Mar-2025
AutoGLM	RuminationZhipu AI (2025)	□	■	□	□	GPQA	GLM-Z1-Air	Mar-2025
Copilot	Researcher Microsoft (2025)	□	■	□	□	□	o3-mini	Mar-2025
H2O.ai	DR H2O.ai (2025)	■	■	■	□	□	h2ogpt-oasst1-512-12b	Mar-2025
Manus	Manus AI (2025)	■	■	□	□	□	Claude3.5, GPT-4o	Mar-2025
Openmanus	Liang et al. (2025)	■	■	□	□	□	Claude3.5, GPT-4o	Mar-2025
OWL CAMEL-AI.org	(2025)	■	■	■	□	□	Deepeek-R1, Gemini2.5-Pro, GPT-4o	Mar-2025
R1-Searcher	Song et al. (2025)	■	□	□	□	2WikiMultiHopQA, HotpotQA	Llama3.1-8B-Inst, Qwen2.5-7B	Mar-2025
ReSearch	Chen et al. (2025b)	■	□	□	□	2WikiMultiHopQA, HotpotQA	Qwen2.5-7B, Qwen2.5-7B-Inst	Mar-2025
Search-R1	Jin et al. (2025)	■	■	□	□	2WikiMultiHopQA, HotpotQA, NQ, TriviaQA	Llama3.2-3B, Qwen2.5-3B/7B	Mar-2025
DeepResarcher	Zheng et al. (2025)	□	■	■	□	HotpotQA, NQ, TriviaQA	Qwen2.5-7B-Inst	Apr-2025
Genspark Super Agent Team	(2025b)	■	■	■	□	□	Mixture of Agents*	Apr-2025
WebThinker	Li et al. (2025d)	■	■	■	■	GPQA, WebWalkerQA	QwQ-32B	Apr-2025
SWIRL	Goldie et al. (2025)	■	□	□	□	HotQA, BeerQA	Gemma 2-27b	Apr-2025
SimpleDeepSearcher Sun*	et al. (2025)	□	■	■	□	2WikiMultiHopQA	Qwen-2.5-7B-In, Qwen-2.5-32B-In, DeepSeek-Distilled-Qwen-2.5-32B, QwQ-32B	Apr-2025
Suna AI AI	(2025b)	■	■	□	□	□	GPT-4o, Claude	Apr-2025
Tool-Star	Dong et al. (2025)	■	■	■	■	WebWalker, HotpotQA, 2WikiMultiHopQA	Qwen-2.5	May-2025
WebDancer	Wu et al. (2025a)	■	■	■	□	WebWalkerQA	Qwen-2.5, QwQ-32B, DeepSeek-R1, GPT-4o	May-2025
AgenticSeek	Martin (2025)	□	■	□	□	□	GPT-4o, DeepSeek-R1, Claude	May-2025
Alita	Qiu et al. (2025)	■	■	■	□	PathVQA	GPT-4o, Claude-Sonnet-4	May-2025
DeerFlow	DanielWalnut (2025)	■	□	□	□	□	Doubao-1.5-Pro-32k, DeepSeek-R1, GPT-4o, Qwen	May-2025
PANGU DIVER	DEEP-Shi et al. (2025)	■	□	□	□	C-SimpleQA, HotpotQA, ProxyQA	Pangu-7B-Reasoner	May-2025
O-agents	Zhu et al. (2025a)	■	□	■	□	□	GPT-4o, GPT-4.1, Claude-3.7-Sonnet, DeepSeek-R1, Gemini-2.5	Jun-2025
Kimi-Researcher	Moonshot AI (2025)	■	■	□	■	SimpleQA	Kimi k1.5/k2	Jun-2025
WebSailor	Li et al. (2025b)	■	■	■	□	SimpleQA	Qwen-2.5	Jul-2025
Agent-KB	Tang et al. (2025b)	■	□	■	□	SWE-bench	GPT-4o, GPT-4.1, Claude-3.7-Sonnet, o3-mini, Qwen-3, DeepSeek-R1	Jul-2025
WebShaper	Tao et al. (2025)	■	■	■	□	WebWalkerQA	Qwen-2.5, QwQ-32B	Jul-2025

Continued on next page

Table 1 continued from previous page

DR Agent	Search Engine		Benchmark		Base Model	Release
	API	Browser	GAIA	HLE	Other QA	
Deep Researcher with Test-Time Diffusion Han et al. (2025)	■	□	■	■	□	Gemini-2.5-Pro Jul-2025
ChatGPT-Agent OpenAI (2025a)	□	□	□	□	□	— Jul-2025
AWorld at InclusionAI (2025)	■	■	■	□	HotpotQA	Gemini-2.5-Pro, GPT-4o Jul-2025
Cognitive Kernel-Pro Wan et al. (2025)	■	□	■	□	AgentWebQA, WebWalkerQA, Multi-hop URLQA, DocBench, TableBench	Claude-3.7-Sonnet, CK-Pro-8B Aug-2025
WebWatcher Geng et al. (2025)	■	■	□	■	Browsercom-VL, LiveVQA, MMSearch	Qwen-2.5-VL-32B Aug-2025
WideSearch Wong et al. (2025)	■	□	□	□	WideSearch	DeepSeek-R1, Doubao-Seed-1.6, Claude Sonnet 4, Gemini-2.5-Pro Aug-2025
MiroRL Team & Team (2025)	■	■	■	□	□	Qwen3-14B Aug-2025

API-based retrieval is a fast, efficient, and scalable way for deep-research (DR) agents to access external knowledge with relatively low latency and computational overhead. For example, Gemini DR (Google Team, 2025) coordinates multiple interfaces, most notably the Google Search and arXiv APIs, to conduct large-scale retrieval over hundreds to thousands of web pages, substantially broadening coverage. Grok DeepSearch (xAI Team, 2025) maintains a continuously updated index via news-outlet feeds, the Wikipedia API, and X’s native interface, and, on demand, dispatches a query-driven agent to decompose questions into targeted subqueries and fetch relevant pages in real time. Perplexity Team (2025) first crawls hundreds of sources and then aggregates them to produce a final report. Cognitive Kernel-Pro (Wan et al., 2025) leverages the free DuckDuckGo search interface to enable a fully open-source, low-cost DR pipeline. Agentic Reasoning, ReSearch, R1-Search and SWIRL (Wu et al., 2025c; Chen et al., 2025b; Song et al., 2025; Goldie et al., 2025) explicitly teach models when to search, what to search for, and how to incorporate retrieved evidence into the reasoning process. PANGU DeepDiver (Shi et al., 2025) uses reinforcement learning to adapt search intensity to task difficulty. Agent Laboratory (Schmidgall et al., 2025) calls the arXiv API to extract paper metadata and abstracts for automated literature reviews, while AI Scientist (Lu et al., 2024) queries the Semantic Scholar API to validate novelty and citation relations among model-generated ideas. CoSearch-Agent (Gong et al., 2024) integrates SerpApi to deliver Slack-based, real-time search. DeepRetrieval (Jiang, 2025) operates within a reinforcement-learning framework to optimise queries against the PubMed and ClinicalTrials.gov APIs for high-recall biomedical retrieval, and Search-o1 (Li et al., 2025c) combines the Bing Search API with the Jina Reader API to dynamically extract and refine passages for downstream reasoning. While these API-driven approaches excel at structured, high-throughput acquisition, they can struggle with deeply nested, client-side JavaScript-rendered content, interactive components, or authentication barriers, motivating complementary browser-based mechanisms capable of comprehensively extracting and analysing dynamic or unstructured information.

Browser-based retrieval provides DR agents with dynamic, flexible, and interactive access to multimodal and unstructured web content through simulated human-like browser interactions. For example, Manus AI’s browsing agent operates a sandboxed Chromium instance for each research session, programmatically opening new tabs, issuing search queries, clicking through result links, scrolling pages until content thresholds are met, filling out form elements when necessary, executing in-page JavaScript to reveal lazily loaded sections, and downloading files or PDFs for local analysis (Manus AI, 2025). Although OpenAI DR, Grok DeepSearch,

and Gemini 2.5 DR do not publicly disclose the implementation details of their browsing capabilities, their ability to handle interactive widgets, dynamically rendered content, and multi-step navigation strongly suggests that they too employ comparable headless-browser frameworks behind the scenes. Among open-source studies, AutoAgent (Yu et al., 2024b) operates within a BrowserGym environment to scroll, interact with page components, and download files when APIs are unavailable; DeepResearcher (Zheng et al., 2025) employs a dedicated Web Browsing Agent that, upon receiving a browse request, processes each segment of a webpage in turn, decides whether to continue to subsequent segments based on relevance, and incrementally aggregates pertinent information into a short-term memory buffer before returning it for reasoning. Kimi-Researcher (Moonshot AI, 2025) uses an internal search engine with a text-based browser to retrieve information. Search-R1 and MiroRL (Jin et al., 2025; Team & Team, 2025) employ both search and browser tools during training. AutoGLM (Zhipu AI, 2025) operationalises browsing through a plan-execute loop that opens and reads web pages, layering “rumination” cycles over browser actions to refine evidence and produce long-form reports. Genspark Super Agent (Team, 2025b) orchestrates a mixture-of-agents in which a research sub-agent performs web search and page reading, handing structured notes to downstream writing/analysis agents rather than relying on a single monolithic browser. SimpleDeepSearcher (Sun* et al., 2025) follows a lightweight search-fetch-summarise loop using web search APIs plus HTTP fetching in lieu of full browser automation, caching pages and compressing them before further reasoning. Tool-Star (Dong et al., 2025) explicitly separates a Search Engine tool from a Web Browser Agent: after link retrieval, the browser agent opens pages, extracts salient snippets, and returns compressed evidence to the planner. AgenticSeek (Martin, 2025) couples a local metasearch front end with a headless, stealth browser so the agent can click, scroll, and submit forms on live sites, exposing knobs for budget and anti-bot robustness. AWorld (at InclusionAI, 2025) provides a multi-agent runtime with built-in browser automation and tracing, enabling teams of agents to divide research, browsing, and synthesis for deep-research workflows over dynamic sites. WebThinker (Li et al., 2025d) performs information seeking by issuing searches and following links on returned result pages. WebDancer, WebSailor, and WebShaper (Wu et al., 2025a; Li et al., 2025b; Tao et al., 2025) combine web search with on-page navigation as a minimal sufficient toolset, linking structured tool invocations with observations in a ReAct-style closed loop: they first locate candidate sources horizontally, then drill down vertically, trading a compact action space for training stability and stronger generalisation. WebWatcher (Geng et al., 2025) uses the Google SerpApi for multimodal search and applies OCR-based image processing to emulate browser interactions. While browser-based retrieval excels at capturing real-time and deeply nested content that API calls cannot reach, it also incurs greater latency, resource consumption, and complexity in handling page variability and errors, suggesting that DR agents may benefit from hybrid architectures that combine the efficiency of API-based methods with the comprehensiveness of browser-driven exploration.

3.2 Tool Use: Empowering Agents with Extended Functionalities

To expand DR agents’ capacity to interact with external environments in complex research tasks, specifically by actively invoking and handling diverse tools and data sources, various DR agents have introduced three core tool modules: code interpreters, data analytics, multimodal processing, along with the Model Context Protocol.

Code Interpreter. The code interpreter capability enables DR agents to execute scripts during inference, allowing them to perform data processing, algorithm verification and model simulation. Most DR agents,

*Mixture of Agents refers to an ensemble of nine base models comprising GPT-4.1, GPT-o3, GPT-o4-mini-high, Claude-Sonnet-3.7-Thinking, Claude-Sonnet-3.7, Gemini-2.0-Flash, Gemini-2.5-Pro, DeepSeek-V3, DeepSeek-R1

Table 2: Comparison of DR Agents with Tool Use Capabilities

■ = Involved, ■ = Non Disclosure, □ = Not present

DR Agent	Code Interpreter	Data Analytics	Multimodal	Release
CoSearchAgent Gong et al. (2024)	□	■	□	Feb-2024
Storm Shao et al. (2024a)	■	□	□	Jul-2024
The AI Scientist Lu et al. (2024)	■	□	□	Aug-2024
Agent Laboratory Schmidgall et al. (2025)	■	□	□	Jan-2025
Agentic Reasoning Wu et al. (2025c)	■	□	□	Feb-2025
AutoAgent Tang et al. (2025a)	■	□	■	Feb-2025
Genspark DR Team (2025b)	■	■	■	Feb-2025
Grok DeepSearch xAI Team (2025)	■	■	■	Feb-2025
OpenAI DR OpenAI (2025b)	■	■	■	Feb-2025
Perplexity DR Perplexity Team (2025)	■	■	■	Feb-2025
Towards an AI co-scientist (Gottweis et al., 2025)	□	■	■	Feb-2025
Agent-R1 Ouyang et al. (2025)	■	□	□	Mar-2025
AutoGLM Romination Zhipu AI (2025)	■	□	■	Mar-2025
Copilot Researcher Microsoft (2025)	■	■	■	Mar-2025
Manus Manus AI (2025)	■	■	■	Mar-2025
OpenManus Liang et al. (2025)	■	■	□	Mar-2025
OWL CAMEL-AI.org (2025)	■	■	■	Mar-2025
H2O.ai DR H2O.ai (2025)	■	■	■	Mar-2025
Genspark Super Agent Team (2025b)	■	■	■	Apr-2025
WebThinker Li et al. (2025d)	■	■	□	Apr-2025
Suna Ai AI (2025b)	■	■	□	Apr-2025
Tool-Star (Dong et al., 2025)	■	■	□	May-2025
AgenticSeek Martin (2025)	■	■	□	May-2025
Alita Qiu et al. (2025)	■	■	■	May-2025
DeerFlow DanielWalnut (2025)	■	■	□	May-2025
O-agents (Zhu et al., 2025a)	■	■	■	Jun-2025
Kimi-Researcher (Moonshot AI, 2025)	■	■	□	Jun-2025
Agent-KB (Tang et al., 2025b)	■	■	■	Jul-2025
AWorld (at InclusionAI, 2025)	■	■	■	Jul-2025
Cognitive Kernel-Pro (Wan et al., 2025)	■	■	■	Aug-2025
WebWatcher (Geng et al., 2025)	■	■	■	Aug-2025
MiroRL (Team & Team, 2025)	■	■	□	Aug-2025

except CoSearchAgent, embed a script execution environment. They typically rely on Python utilities such as Aider and Java utilities to orchestrate dynamic scripting, conduct literature-driven analysis and carry out real-time computational reasoning.

Data Analytics. By integrating data analytics modules, DR agents transform raw retrievals into structured insights by computing summary statistics, generating interactive visualisations and conducting quantitative model evaluations, thereby accelerating hypothesis testing and decision-making. Many commercial DR agents have implemented analytics features such as charting, table generation and statistical analysis, either locally or via remote services. However, most of these systems have not publicly disclosed technical details of their implementations. In contrast, academic studies often provide concrete examples: CoSearchAgent (Gong et al., 2024) integrates SQL-based queries within team communication platforms to run aggregate analyses and produce reports; AutoGLM (Zhipu AI, 2025) extracts and analyses structured datasets directly from table-based web interfaces.

Multimodal Processing and Generation. Multimodal processing and generation tools enable DR agents to integrate, analyse and generate heterogeneous data such as text, images, audio and video within a unified reasoning pipeline, thereby enriching their contextual understanding and broadening the range of their outputs. Only a subset of mature commercial and open-source projects, for example Manus (Manus AI, 2025), OWL (CAMEL-AI.org, 2025), AutoAgent (Tang et al., 2025a), AutoGLM (Zhipu AI, 2025), OpenAI (OpenAI, 2025b), Gemini (Google Team, 2025), Perplexity (Perplexity Team, 2025) and Grok DeepSearch (xAI Team, 2025), support this capability, whereas most academic prototypes have not implemented it, often due to the high computational cost. As the typical open source studies, OWL and Openmanus extend their pipelines to include interactions with platforms such as GitHub, Notion and Google Maps and to leverage numerical libraries such as Sympy and Excel for combined data analysis and multimodal media processing (CAMEL-AI.org, 2025; Liang et al., 2025).

Deep Research Agent with Computer Use. Most recently, the boundaries of DR agents have been progressively expanded through integrating computer-assisted task execution capabilities (i.e., computer use). For example, Zhipu AI introduced AutoGLM Rumination (Zhipu AI, 2025), an RL-based system incorporating self-reflection and iterative refinement mechanisms, which significantly enhances multi-step reasoning and advanced function-calling abilities. Specifically, AutoGLM Rumination (Zhipu AI, 2025) autonomously interacts with web environments, executes code, invokes external APIs, and effectively accomplishes sophisticated tasks, including data retrieval, analysis, and structured generation of comprehensive reports. Comparison with OpenAI’s DR: While OpenAI DR primarily focus on intricate reasoning and information retrieval, AutoGLM Rumination exhibits superior autonomy in practical execution. This enhanced autonomy allows it to transform abstract analytical insights into concrete operational tasks, such as automated interactions with web interfaces and real-time data processing. Moreover, AutoGLM Rumination addresses and resolves limitations inherent in simulated browsing environments by seamlessly integrating advanced reasoning capabilities with authentic browser-based interactions. Therefore, the agent gains reliable access to user-authenticated resources, including platforms such as CNKI, Xiaohongshu, and WeChat official accounts. Such integration significantly elevates the agent’s autonomy and adaptability in both information acquisition and execution of real-world tasks.

OpenAI DR, Perplexity DR, Grok DR, H2O, Manus, Genspark Super Agent OpenAI (2025b); Perplexity Team (2025); xAI Team (2025); H2O.ai (2025); Manus AI (2025); Team (2025b) orchestrate multi-step web research with browser or tool use, code execution, and analytics to produce cited, structured reports at industrial scale. OWL, OpenManus, Suna, DeerFlow, WebThinker, AgenticSeek, and AWorld (CAMEL-AI.org, 2025;

Liang et al., 2025; AI, 2025b; DanielWalnut, 2025; Li et al., 2025d; Martin, 2025; at InclusionAI, 2025) provide open-source stacks that integrate browser automation, code interpreters, and (often) MCP-style tooling for end-to-end deep research workflows. AutoGLM Ruminator, Tool-Star, Kimi-Researcher, and MiroRL (Zhipu AI, 2025; Dong et al., 2025; Moonshot AI, 2025; Team & Team, 2025) use reinforcement learning or self-reflection to plan searches, invoke tools and code, and improve multi-step reasoning autonomy. The AI Scientist, Storm, Agent Laboratory, Agent-R1, AutoAgent, and CoSearchAgent (Lu et al., 2024; Shao et al., 2024a; Schmidgall et al., 2025; Ouyang et al., 2025; Tang et al., 2025a; Gong et al., 2024) automate literature review, code/experiment execution, and structured drafting through tool-augmented research pipelines. Towards an AI co-scientist and O-agents (Gottweis et al., 2025; Zhu et al., 2025a) outline blueprints and empirical recipes for building, training, and evaluating effective tool-using research agents. Agent-KB (Tang et al., 2025b) introduces a KB-driven framework for cross-domain experience transfer to boost generalisation on complex tasks. Microsoft Copilot Researcher (Microsoft, 2025) embeds multi-step research and analytics in the Microsoft 365 ecosystem, producing reports and charts within enterprise workflows. Alita (Qiu et al., 2025) explores self-evolving agents that generate and wrap MCP tools, alongside code execution, to extend capabilities with minimal predefined schemas.

3.3 Architecture and Workflow

As shown in Figure 4, this section systematically analyses the construction of DR systems, focusing on workflows categorised into **static** and **dynamic** types. We first introduce the static workflows and then discuss planning strategies, which enhance task allocation and execution through **three distinctive user interaction types** to clarify intent: planning-only (direct planning without clarifying user intent), intent-to-planning (clarifying intent before planning to align the task with user goals), and unified intent-planning (generating a plan and requesting user confirmation). The distinction between **single-agent** and **multi-agent** systems is examined in the context of dynamic workflows, emphasising specialisation in task management. Additionally, we examine memory mechanisms for managing and integrating retrieved information, which enhance the performance and adaptability of DR systems.

3.3.1 Static vs. Dynamic Workflows

Static Workflows. Static workflows rely on manually predefined task pipelines, decomposing research processes into sequential subtasks executed by dedicated agents. These workflows follow explicitly structured procedures, making them particularly suitable for well-defined, structured research scenarios. For instance, AI Scientist (Lu et al., 2024) automates scientific discovery through distinct sequential phases, including ideation, experimentation, and reporting. Similarly, Agent Laboratory (Schmidgall et al., 2025) segments research activities into formalised stages, such as literature review, experimentation, and synthesis of findings. Extending this static paradigm further, AgentRxiv (Schmidgall & Moor, 2025) incorporates inter-agent collaboration mechanisms, enabling incremental knowledge reuse through sharing intermediate research outcomes among specialised agents.

Whilst their ease of implementation and structured clarity, static workflows suffer from limited generalisation capabilities, as each distinct task necessitates a specifically tailored pipeline.

Dynamic Workflows. To overcome the limitations in flexibility and generalizability inherent in static workflows, dynamic workflows support adaptive task planning, allowing agents to dynamically reconfigure task structures based on iterative feedback and evolving contexts. Dynamic architectures leverage advanced mechanisms including automated planning, iterative refinement, and interactive task allocation, enabling tasks

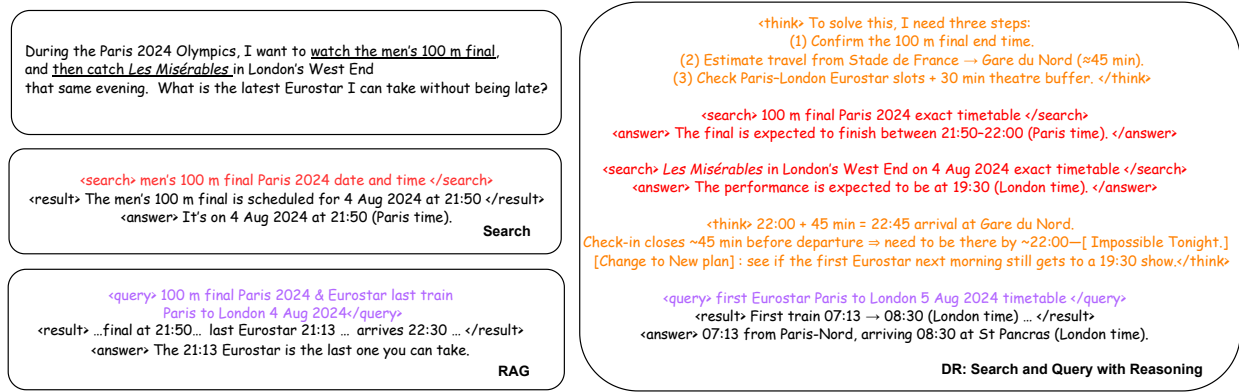


Figure 3: **Comparison of Information Retrieval Methods.** The upper left corner (Search) represents the searching methods, which can use the browser or API; the lower left corner (RAG, Query) represents Retrieval-Augmented Generation, combining retrieval and generative models to output natural language answers; the right side (Deep Research) represents the deep research process, generating complex decisions or analyses through retrieval and explicit reasoning.

to evolve in real-time as new knowledge or external inputs become available. Consequently, dynamic workflows exhibit superior generality and adaptability, making them highly suitable for complex, knowledge-intensive tasks commonly encountered in AI-driven research scenarios.

3.3.2 Dynamic Workflows: Planning Strategies

To enhance DR agents' adaptability in response to evolving user requirements and contexts, existing studies propose three distinctive LLM-based planning strategies, each differing in whether and how they interact with the user to clarify intent:

- 1) The **Planning-Only** approach directly generates task plans based solely on initial user prompts without actively engaging in further clarification, adopted by the majority of existing DR agents, including Grok (xAI Team, 2025), H2O (H2O.ai, 2025) and Manus (Manus AI, 2025).
- 2) The **Intent-to-Planning** strategy actively clarifies user intent prior to planning through targeted questions, subsequently generating tailored task sequences based on clarified user inputs; this method is utilised by OpenAI DR (OpenAI, 2025b).
- 3) The **Unified Intent-Planning** approach synthesises these methods by generating a preliminary plan from the initial prompt, together with interactively engaging the user to confirm or revise the proposed plan. Gemini DR (Google Team, 2025) is representative of this strategy, effectively adopts the strength of user-guided refinement.

3.3.3 Dynamic Workflows: Single-Agent vs. Multi-Agent

Dynamic workflows of DR agents can be differentiated based on agent architectures into single-agent and multi-agent frameworks, each exhibiting distinct characteristics concerning task specialisation, coordination complexity, and scalability of execution.

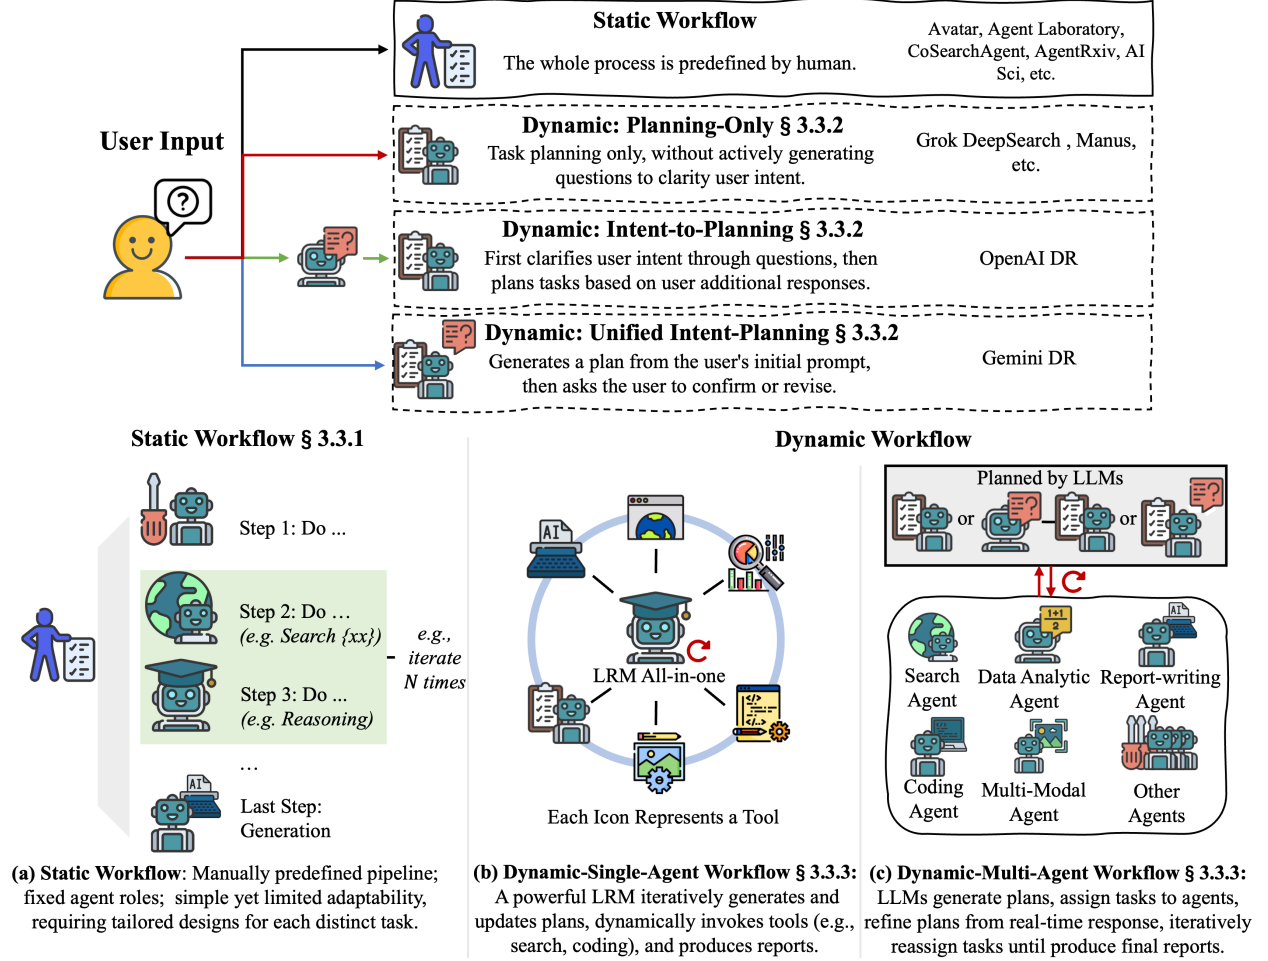


Figure 4: Comparison of DR Workflows: (1) **Static vs. Dynamic Workflows:** Static workflows rely on predefined task sequences, while dynamic workflows allow LLM-based task planning. (2) **Planning Strategies:** Three types include: planning-only (direct planning without clarifying user intent), intent-to-planning (clarifying intent before planning), and unified intent-planning (generating a plan and requesting user confirmation). (3) **Single-Agent vs. Multi-Agent:** Dynamic workflows can be categorised into dynamic-multi-agent systems (tasks distributed across specialised agents) or dynamic-single-agent systems (an LRM autonomously updates and executes tasks).

Dynamic Single-Agent Systems. Dynamic single-agent systems **integrate planning, tool invocation, and execution within a unified LRM**, streamlining task management into a cohesive cognitive loop. Single-agent architectures autonomously refine task plans and invoke appropriate tools based on evolving contexts, typically without explicit inter-agent coordination. Compared to multi-agent architectures, single-agent systems enable direct end-to-end reinforcement learning (RL) optimisation across the entire workflow, facilitating smoother and more coherent integration of reasoning, planning, and tool invocation. Systems such as Search-o1 (Li et al., 2025c), R1-Searcher (Song et al., 2025), DeepResearcher (Zheng et al., 2025), WebDancer (Wu et al., 2025a), WebSailor (Li et al., 2025b), PANGU Deepdiver (Shi et al., 2025), Agent-R1 (Ouyang et al., 2025), ReSearch (Chen et al., 2025b), Search-R1 (Jin et al., 2025), WebWatcher (Wu et al., 2025b), MiroRL (Team & Team, 2025) and Kimi-Researcher (Moonshot AI, 2025) exemplify this paradigm through iterative cycles of explicit reasoning, action, and reflection, aligning with the ReAct framework (Yao et al., 2023). However, this streamlined approach places significant demands on the foundation model’s reasoning capabilities, contextual understanding, and autonomous selection and invocation of tools. Additionally, the tightly integrated nature of single-agent systems may limit modular flexibility, complicating independent scaling or optimisation of individual functional components.

Dynamic Multi-Agent Systems. Dynamic multi-agent systems leverage multiple specialised agents to collaboratively execute subtasks generated and dynamically allocated through adaptive planning strategies. These systems typically employ hierarchical or centralised planning mechanisms, wherein a coordinator agent continuously assigns and redistributes tasks based on real-time feedback and replanning. Representative frameworks include OpenManus (Liang et al., 2025) and Manus (Manus AI, 2025), both adopting hierarchical planner-toolcaller architectures. Similarly, OWL (CAMEL-AI.org, 2025) includes a workforce-oriented model, utilising a central manager agent to orchestrate task distribution among specialised execution agents. Furthermore, Alita (Qiu et al., 2025) incorporates a self-evolution mechanism into DR agents, allowing the agent to online instantiate and configure new MCP servers tailored to specific tasks and environmental conditions. AWorld (at InclusionAI, 2025) is an open-source framework for building, orchestrating, and training tools using agents and larger multi-agent systems, offering memory and context services and MCP tool integration for scalable evaluation and self-improvement. Webwalker (Wu et al., 2025b) mimics human-like web navigation through an explore-critic paradigm. WebThinker (Li et al., 2025d) uses both executed and auxiliary agents to autonomously search, deeply explore web pages, and draft research reports. Such multi-agent configurations effectively handle complex, parallelizable research tasks, thereby enhancing flexibility and scalability in open-ended research scenarios. Nevertheless, a major current challenge of multi-agent systems lies in the inherent complexity of coordinating multiple independent agents, making it difficult to conduct effective end-to-end reinforcement learning optimisation.

3.3.4 Memory Mechanism for Long-Context Optimisation

Memory mechanisms empower DR agents to persistently capture, organise, and recall relevant information across multiple retrieval rounds, thereby reducing redundant queries and improving both the efficiency and coherence of DR tasks. During the DR process, agents typically perform extensive multi-round retrieval, generating hundreds of thousands of tokens (or even millions). Although recent advances in LLMs have significantly expanded context window sizes, current limits still constrain tasks involving extremely long contexts. To address these challenges, DR systems have implemented various optimisations for processing extended contexts. Broadly, these optimisations can be categorised into three main strategies: **(i) Expanding the Context Window Length; (ii) Compressing Intermediate Steps; (iii) Utilising External Structured Storage for Temporary Results.**

Extending the Context Window Length. It is the simple but intuitively effective approach, exemplified by Google’s Gemini model (Google Team, 2025), which supports a context window of up to one million tokens, supplemented by a RAG setup. Despite its straightforwardness, this method often incurs high computational costs and may lead to inefficiencies in resource utilisation during practical deployments.

Compressing Intermediate Step. An alternative strategy involves compressing or summarising intermediate reasoning steps, significantly reducing the number of tokens processed by the model and thereby improving both efficiency and output quality. Representative frameworks such as The AI Scientist (Lu et al., 2024) and CycleResearcher (Weng et al., 2024) pass summarised intermediate results between workflow phases. Further, Search-o1 (Li et al., 2025c) introduced the concept of “Reason-in-Documents”, utilising LRMs to compress documents, substantially reducing token volume and enhancing model decision-making efficiency. Meanwhile, WebThinker (Li et al., 2025d) uses an auxiliary model to compress the external information. However, a potential drawback of this approach is the loss of detailed information, potentially impacting the precision of subsequent reasoning.

Utilising External Structured Storage. This is for preserving and retrieving historical information, enabling DR agents to persistently and efficiently store vast amounts of past context beyond the constraints of the context window, improving memory capacity, retrieval speed, and semantic relevance. Popular open-source frameworks such as Manus (Manus AI, 2025), OWL (CAMEL-AI.org, 2025), Open Manus (Liang et al., 2025), and Avatar (Wu et al., 2024) utilise external file systems to store intermediate outcomes and historical data for subsequent retrieval. Frameworks like AutoAgent (Tang et al., 2025a) have developed self-managing modules that leverage vector databases to support scalable memory storage and fast similarity-based lookup. Beyond plain text or vector stores, some works propose more semantically structured memory frameworks: for instance, Agentic Reasoning (Wu et al., 2025c) employ knowledge graphs to capture intermediate reasoning processes and thereby enhance the precision of information reuse, while Agentrxiv (Schmidgall & Moor, 2025) simulates an academic repository akin to arXiv for storing and retrieving relevant outcomes from other agents. Furthermore, Agent-KB (Tang et al., 2025b) and Alita (Qiu et al., 2025) construct shared knowledge bases and optimised toolsets for agentic problem-solving. Although these structured approaches offer superior semantic retrieval efficiency and accuracy, they typically entail higher development and maintenance costs due to the need for meticulous data structure design and management.

3.4 Tuning: Beyond Prompting toward Capability Enhancement

Table 3: Comparison of DR Agents with Tuning Methods

■ = Yes, ■ = Yes but details unknown, □ = Not present						
DR Agent	SFT	RL	Base Model	Data	Reward Design	Release
Gemini DR Google Team (2025)	■	■	Gemini-2.0-Flash	□	■	Dec-2024
WebWalker	□	□	GPT-4o, Qwen-2.5 7-72B	WebWalkerQA	□	Jan-2025
Grok DeepSearch xAI Team (2025)	□	■	Grok3	□	■	Feb-2025
OpenAI DR OpenAI (2025b)	□	■	GPT-o3	□	■	Feb-2025
Agentic Reasoning Wu et al. (2025c)	■	□	DeepSeek-R1, Qwen2.5	GPQA	Rule-Outcome	Feb-2025
AutoAgent Tang et al. (2025a)	□	■	Claude-Sonnet-3.5	□	□	Feb-2025
Towards an AI co-scientist	□	□	Gemini 2.0	□	□	Feb-2025

Continued on next page

Table 3 continued from previous page

DR Agent	SFT	RL	Base Model	Data	Reward Design	Release
Agent-R1 Ouyang et al. (2025)	□	PPO Schulman et al. (2017), Reinforce++ Hu (2025), GRPO Shao et al. (2024b)	Qwen2.5-1.5B-Inst	HotpotQA	Rule-Outcome	Mar-2025
AutoGLM Romination Zhipu AI (2025)	■	■	GLM-Z1-Air	□	■	Mar-2025
H2O.ai DR H2O.ai (2025)	■	■	h2ogpt-oasst1-512-12b	□	■	Mar-2025
Copilot Researcher Microsoft (2025)	■	■	o3-mini	□	□	Mar-2025
ReSearch Chen et al. (2025b)	□	GRPO Shao et al. (2024b)	Qwen2.5-7B-Inst, Qwen2.5-32B-Inst	2WikiMultiHopQA	Rule-Outcome	Mar-2025
R1-Searcher Song et al. (2025)	■	Reinforce++ Hu (2025), GRPO Shao et al. (2024b)	Qwen2.5-7B-Inst, LLaMA-3.1-8B-Inst	2WikiMultiHopQA HotpotQA	Rule-Outcome	Mar-2025
Search-R1 Jin et al. (2025)	■	PPO Schulman et al. (2017), GRPO Shao et al. (2024b)	Qwen2.5-3B/7B, LLaMA3.2-3B-Inst	NQ, HotpotQA	Rule-Outcome	Mar-2025
Nouswise (Nouswise, 2025)	■	■	Nouswise	□	■	Mar-2025
DeepResearcher Zheng et al. (2025)	□	GRPO Shao et al. (2024b)	Qwen2.5-7B-Inst	NQ, HotpotQA	Rule-Outcome	Apr-2025
Genspark Super Agent Team (2025b)	□	■	Mixture of Agents	□	■	Apr-2025
WebThinker Li et al. (2025d)	■	Iterative Online DPO (Rafailov et al., 2023)	QwQ-32B	Expert Dataset	Rule-Outcome	Apr-2025
SWIRL Goldie et al. (2025)	□	Offline-RL	Gemma-2-27B	HotPotQA	□	Apr-2025
SimpleDeepSearcher Sun* et al. (2025)	■	PPO Schulman et al. (2017)	Qwen-2.5-7B-In, Qwen-2.5-32B-In, Deepseek-Distilled-Qwen-32B, QwQ-32B	NQ, HotpotQA, 2WikiMulti-HopQA, Musique, SimpleQA, MultiHop-RAG	Process-based reward	Apr-2025
PANGU DEEP-DIVER Shi et al. (2025)	■	GRPO Shao et al. (2024b)	Pangu-7B-Reasoner	WebPuzzle	Rule-Outcome	May-2025
Tool-Star (Dong et al., 2025)	■	GRPO Shao et al. (2024b)	Qwen-2.5	NuminaMath, HotpotQA, 2WikiMulti-HopQA	Rule-Outcome	May-2025
WebDancer (Wu et al., 2025a)	■	DAPO (Yu et al., 2025)	Qwen-2.5-7B/32B, QwQ-32B, DeepSeek-R1, GPT-4o, GPT-4.1, Claude-3.7-Sonnet, DeepSeek-R1, Gemini-2.5	CRAWLQA, E2HQA	Rule-Outcome	May-2025
O-agents (Zhu et al., 2025a)	□	□	■	□	□	Jun-2025
Kimi-Researcher (Moonshot AI, 2025)	□	REINFORCE	Kimi k1.5/k2	□	Rule-Outcome	Jun-2025
WebSailor (Li et al., 2025b)	■	DUPO	Qwen-2.5-3B/7B/32B/72B, GPT-4o, GPT-4.1, Claude-3.7-Sonnet, o3-mini, Qwen-3 32B, DeepSeek-R1	SailorFog-QA	Rule-Outcome	Jul-2025
Agent-KB (Tang et al., 2025b)	□	□	Qwen-2.5-3B/7B/32B/72B, QwQ-32B	□	□	Jul-2025
WebShaper (Tao et al., 2025)	■	GRPO Shao et al. (2024b)	Qwen-2.5-3B/7B/32B/72B, QwQ-32B	WebShaper	Rule-Outcome	Jul-2025

Continued on next page

Table 3 continued from previous page

DR Agent	SFT	RL	Base Model	Data	Reward Design	Release
Cognitive Kernel-Pro (Wan et al., 2025)	■	□	Claude-3.7-Sonnet, CK-Pro-8B	OpenWebVoyager, Multi-hop URLQA, AgentWebQA, WebWalkerQA, DocBench, TableBench	□	Aug-2025
WebWatcher (Geng et al., 2025)	□	GRPO Shao et al. (2024b)	Qwen-2.5-VL-32B	BrowseComp- VL, Long-tail VQA, Hard VQA	Rule-Outcome	Aug-2025
MiroRL (Team & Team, 2025)	■	GRPO Shao et al. (2024b)	Qwen3-14B	MiroRL-GenQA	Rule-Outcome	Aug-2025

Parametric Approaches. Prompt-based methods directly leverage the capabilities of pre-trained LLMs, enabling complex functionalities without expensive fine-tuning or additional training. However, it remains challenging to systematically optimise prompt structures and workflows. Moreover, since an agent’s performance is inherently limited by its backbone LLM, increasing the complexity of decision-making processes quickly reaches the model’s performance ceiling. To overcome these limitations, it is essential to incorporate advanced optimisation techniques such as fine-tuning, reinforcement learning (RL) or hybrid training paradigms to further extend the model’s inherent capabilities. Below, we discuss the two main tuning paradigms, supervised fine-tuning (SFT) and RL, and highlight how each extends agent capabilities beyond prompt-only methods.

3.4.1 SFT-based Optimization

Prompt-based approaches, while effective for rapid adaptation, are fundamentally constrained by the intrinsic generalisation capacity of backbone LLMs and often exhibit limited robustness in complex task settings. In order to address these limitations, researchers have increasingly explored fine-tuning methodologies aimed at systematically optimising LLMs for critical components of deep research agents. These components include search query formulation, structured report generation, and external tool utilisation. These efforts aim to enhance retrieval quality, mitigate hallucinations, and enable more reliable long-form and evidence-grounded generation.

An early milestone in this research direction is Open-RAG (Islam et al., 2024), which augments data construction with diverse supervisory signals, including retrieval tokens, relevance tokens, grounding tokens, and utility tokens. Through adversarial training, Open-RAG improves the model’s capability to filter irrelevant information, thereby enhancing both retrieval accuracy and the quality of downstream tasks. Building upon this foundation, AUTO-RAG (Yu et al., 2024b) enhances the autonomous iterative retrieval capabilities of LLMs. In contrast to earlier multi-hop retrieval approaches that relied on few-shot prompting or hand-crafted templates (Jiang et al., 2023; Feng et al., 2024; Wang et al., 2024b), AUTO-RAG constructs reasoning-grounded instruction datasets, enabling models to autonomously plan retrieval queries and engage in multi-round interactions with retrievers. The model dynamically refines its retrieval strategy during generation, gathering sufficient evidence before synthesising a final answer. Extending these retrieval-centric innovations, DeepRAG (Guan et al., 2025) proposes a binary tree search mechanism that recursively generates sub-queries and constructs multi-turn retrieval trajectories. This mechanism enables the model to judiciously balance between internal parametric knowledge and external retrieval-based rollouts. Consequently, it enhances search efficiency and mitigates redundant external queries.

In order to further reduce reliance on manually constructed supervised fine-tuning (SFT) datasets, recent work has sought to reduce dependence on manually constructed supervised fine-tuning datasets by developing fine-tuning strategies based on rejection sampling. CoRAG (Wang et al., 2025a) uses rejection sampling to extract intermediate retrieval chains from standard question answering datasets, allowing for stepwise retrieval augmentation and dynamic reformulation of subqueries as context evolves instead of supervising only final outputs. Li et al. (Li et al., 2025a) propose a hint-infer mechanism that monitors token patterns during generation and triggers external computational tools, such as Python executors or hint libraries, when specific cues are detected. After an initial supervised fine-tuning phase, the model undergoes a rejection sampling fine-tuning process that teaches it to generate its own prompts and invoke tools autonomously without reliance on hand-curated demonstrations. ATLAS (Chen et al., 2025d) proposes a novel approach for LLM-based agents that trains exclusively on selected critical steps from expert trajectories, significantly improving generalisation performance.

Although these SFT methods enhance the generalisation of deep research agents by supporting dynamic retrieval planning, structured information synthesis, and integrated tool use, they remain confined to offline, static retrieval pipelines characteristic of retrieval-augmented systems. In contrast, reinforcement learning offers a more adaptive solution for online query generation and tool invocation. By learning from real-time reward signals, reinforcement learning agents acquire the ability to formulate effective search queries and determine the optimal timing for tool calls. This approach addresses the limitations of synthetic demonstration data and distributional shifts, yielding more robust and adaptive performance in open-ended research environments.

3.4.2 Reinforcement Learning-based Optimisation

RL-based methods optimise DR agents by directly enhancing their adaptive capabilities and generalisation across diverse tasks, surpassing conventional instruction-following or pattern learning approaches. Recent advances have demonstrated that end-to-end RL training significantly strengthens iterative information retrieval, dynamic tool invocation, and integrated reasoning capabilities within DR agents. See comparative analysis in Table 3.

Early RL-based approaches such as DeepRetrieval (Jiang, 2025) optimised query generation for improved information retrieval quality, effectively enhancing downstream text generation by producing more relevant search results. Building on query optimisation, ReSearch (Chen et al., 2025b) extended RL to adaptive reasoning over retrieved information. The model dynamically refined search strategies and iteratively updated results based on continuous feedback, significantly improving task-solving accuracy. Subsequently, R1-Searcher (Song et al., 2025) further optimised retrieval interactions, explicitly training models to refine search strategies through carefully designed reward functions. This allowed better exploitation of external information and improved search result relevance.

Search-R1 (Jin et al., 2025) advanced RL-based retrieval by structurally integrating sophisticated search interactions with complex reasoning processes. The method systematically bridged query generation and information reasoning, enabling nuanced responses through refined integration of retrieved content. Finally, this research line culminated in the development of Agent-R1 (Ouyang et al., 2025), a comprehensive DR

framework integrating RL into end-to-end training of LLM agents. Agent-R1 leveraged diverse tools such as APIs, search engines, and databases, achieving autonomous multi-step task execution and dynamic tool coordination. Through RL-driven optimisation across its entire pipeline, Agent-R1 demonstrated advanced capabilities in adaptive planning, iterative execution, and task refinement. Moreover, WebThinker (Li et al., 2025d) integrates a Web Explorer module for dynamic multi-hop web exploration and employs Iterative Online Direct Preference Optimisation (DPO) to seamlessly interleave search, navigation, and report drafting during reasoning, while Pangu DeepDiver (Shi et al., 2025) builds on the 7B Pangu model pretrained on Huawei’s Ascend NPUs by introducing Search Intensity Scaling (SIS) through a two-phase SFT and RL curriculum, enabling adaptive adjustment of search depth and frequency in open-web environments.

Table 3 reveals three key RL implementation patterns in DR systems: 1) Industrial systems like Gemini DR (Google Team, 2025) and Grok DeepSearch (xAI Team, 2025) employ proprietary RL implementations with undisclosed details, 2) Academic approaches (Chen et al., 2025b; Song et al., 2025) favor modular RL optimization using GRPO (Shao et al., 2024b) and Reinforce++ (Hu, 2025) with transparent reward designs, and 3) Emerging hybrid systems like SimpleDeepSearcher (Sun* et al., 2025) combine process-based rewards with multi-task training across 6 QA datasets. The table also highlights the prevalence of Qwen2.5 and LLaMA3 model families as preferred base architectures for RL optimisation.

Reward Model and Policy Model. Most current open-source RL implementations of DR agents, including the methods discussed above, commonly adopt rule-based reward models that explicitly define task-specific objectives such as retrieval relevance, information accuracy, or successful tool invocation. To efficiently perform policy optimisation, recent systems have increasingly utilised Proximal Policy Optimisation (PPO) (Schulman et al., 2017) and Group Relative Policy Optimisation (GRPO) (Shao et al., 2024b). In particular, GRPO fundamentally reconfigures the advantage estimation paradigm by replacing traditional value functions with group-relative advantage computation. It expands reward space through intra-group normalisation, and sparse binary rewards are transformed into continuous advantage values spanning wider ranges. This expanded signal space provides richer gradient information for policy updates, as evidenced higher high-reward response density compared to PPO. In addition, GRPO provides a variance suppression mechanism by constraining advantage estimation within dynamically clustered response groups, such as grouping by reasoning depth or tool usage patterns, reducing policy gradient variance through local standardisation. In contrast to PPO, GRPO eliminates separate value networks, removing conflicting optimisation objectives between policy and value functions. Empirical measurements show GRPO reduces gradient direction conflicts from 12 to 3 per training epoch, significantly accelerating convergence. As a result, GRPO outperforms conventional PPO in wider reward distribution coverage, enhancing exploration capacity and faster KL divergence stabilisation during alignment.

3.5 Non-parametric Continual Learning

DR agents depend heavily on LRMs and often utilise complex hierarchical workflows. Parameter-based learning approaches such as SFT and RL encounter significant obstacles in this context, including the need to scale model parameters, manage extensive volumes of structured experience data, and design increasingly intricate training algorithms. In contrast, non-parametric continual learning approaches offer a scalable alternative: agents refine their capabilities at runtime by optimising external memory, workflows, and tool configurations through continuous interaction with the external environment rather than by updating internal weights. This non-parametric continual learning paradigm enables efficient online adaptation with minimal data and computational overhead, making it well-suited to DR agents with complex architectures.

Non-parametric continual learning approaches, most notably case-based reasoning (CBR), are currently a mainstream method in LLM-driven agent systems. The CBR-based method enables agents to retrieve, adapt, and reuse structured problem-solving trajectories from an external case bank dynamically. Unlike traditional RAG-based methods, which rely on static databases, CBR facilitates online contextual adaptation and effective task-level generalisation. Such flexibility underscores its potential as a scalable and practical optimisation solution for DR agents with complex architecture. DS-Agent (Guo et al., 2024) is a pioneering LLM-driven agent that introduced CBR into automated data science workflows, employing approximate online retrieval from a constructed case bank. Similarly, LAM (Guo et al., 2025) applies CBR techniques to functional test generation, combining trajectory-level retrieval with LLM planning in a modular system design. Although DS-Agent itself does not include a learning phase, Agent K (Grosnit et al., 2024) advances this paradigm with dynamic external case retrieval and reuse guided by a reward-based memory policy, which exemplifies genuine self-evolution enabling continual adaptation and optimisation without updating model parameters. Focusing on DR agents, AgentRxiv (Schmidgall & Moor, 2025) further extends this paradigm by enabling autonomous research agents to collaboratively share and access a centralised repository of prior research outputs. This framework allows LLM agent laboratories to upload and retrieve reports from a shared preprint server, simulating an online-updating arXiv-like platform, which can be seen as a comprehensive case bank. Such a system empowers agents to enhance their capabilities and knowledge through contextual adaptation without modifying their model parameters.

Compared to prompt-based methods, which encode fixed demonstrations or task heuristics into static input templates, Non-parametric methods enable dynamic retrieval and adaptation of structured trajectories, thereby facilitating continual task generalisation without manual prompt engineering. Relative to RAG, which typically retrieves unstructured textual content from static corpora, CBR operates at the trajectory level and emphasises reasoning-centred memory organisation. A notable example is the Kaggle Grandmaster Agent (Grosnit et al., 2024), which demonstrates how LLMs equipped with modular reasoning components and persistent memory can achieve expert-level structured problem solving, aligning closely with the CBR paradigm. These characteristics make CBR particularly well-suited for agents requiring procedural adaptation and context-sensitive optimisation across tasks. Except memory-based method, self-evolution can also arise from dynamic infrastructure adaptation. For example, Alita (Qiu et al., 2025) monitors task requirements and environmental signals to provision and configure new MCP servers at runtime, seamlessly extending and refining its toolset on demand.

In summary, these self-evolution paradigms in LLM-driven DR agent systems offer substantial promise for structured reasoning and dynamic retrieval and open new pathways for efficient knowledge reuse and continual learning. Although these methods have not yet achieved widespread attention, they address the high data and computational demands inherent to parameter-based approaches and therefore represent an attractive direction for future research and practical deployment.

4 Industrial Applications of Deep Research Agents

4.1 Open AI Deep Research

OpenAI recently introduced its DR capability (OpenAI, 2025b), employing a single-agent architecture centred around a reinforcement learning-based, fine-tuned o3 reasoning model. Upon receiving a research query, the system initiates a concise interactive clarification step to accurately define user intent and research objectives. It then autonomously formulates and executes a sophisticated, multi-step research strategy,

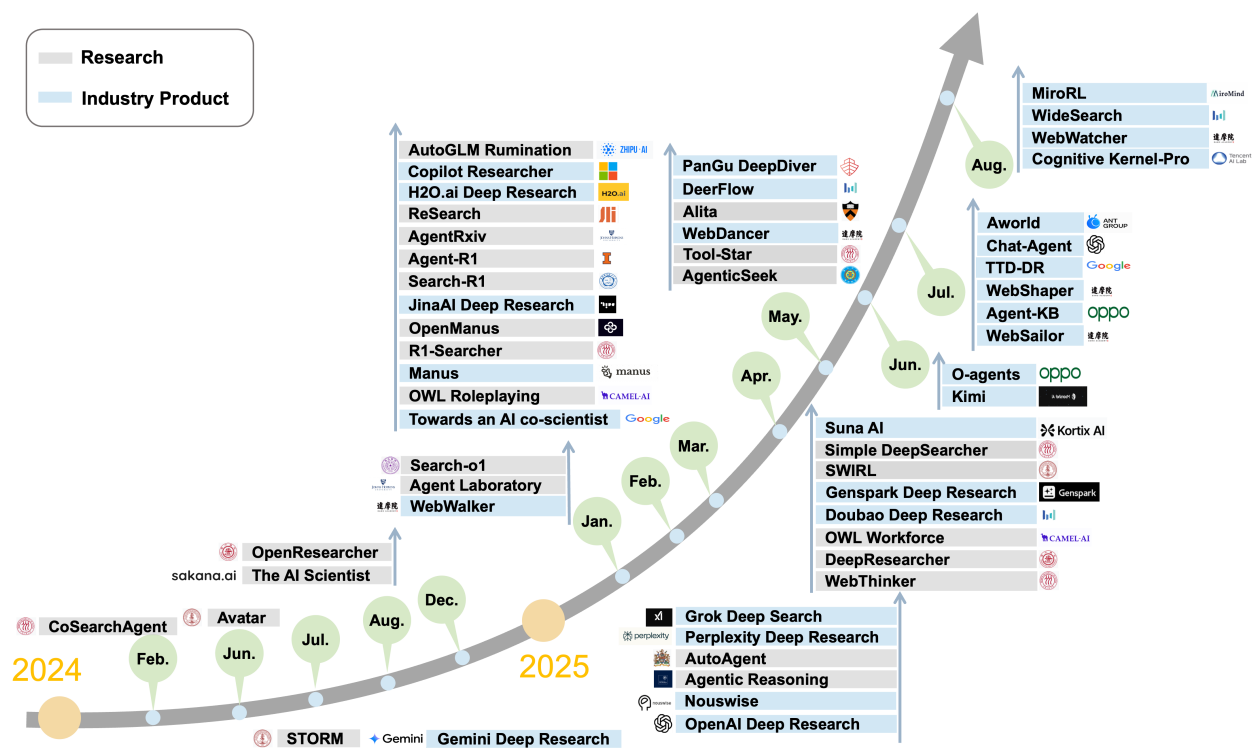


Figure 5: An overview of DR agents’ evolution over years.

encompassing multimodal information retrieval, web browsing, and computational tasks such as data analysis and visualisation through browser tools. Technologically, this solution delivers three significant advancements: (1) **A dynamically adaptive iterative research workflow**: Capable of refining its strategy throughout task execution. (2) **Enhanced context memory and robust multimodal processing capabilities**: Facilitating effective integration of diverse information sources. (3) **Comprehensive toolchain integration**: Combining web browsing capabilities with built-in programming tools to produce structured, authoritative reports supported by precise citations.

4.2 Gemini Deep Research

Google DeepMind recently introduced Gemini DR (Google Team, 2025), an advanced DR agent based on its multimodal Gemini 2.0 Flash Thinking model. Gemini’s reinforcement learning-driven fine-tuning, facilitated by a single-agent architecture, has been shown to enhance planning and adaptive research capabilities, enabling the system to autonomously and expeditiously complete complex tasks. Technologically, this solution delivers four significant advancements: (1) **Interactive Research Planning**: Upon receiving a research query, Gemini autonomously formulates a multi-step investigation plan for interactive user review and modification. (2) **Asynchronous Task Management**: Adopts an asynchronous task management architecture to efficiently handle multiple simultaneous tasks. (3) **Large-scale context windows RAG ensembles**: Enabling effective management and coherent synthesis of multimodal data (eg, text, images) for in-depth professional research analysis. (4) **High speed adaptive retrieval**: Implements fast, multi-round adaptive web search that significantly outperforms other agents in terms of retrieval speed and amount of information per iteration.

4.3 Perplexity Deep Research

Perplexity’s recently developed DR agent (Perplexity Team, 2025) has demonstrated an advanced capability to decompose complex queries into well-defined subtasks. The system is capable of conducting targeted web searches iteratively, critically evaluating authoritative sources, and synthesising structured, comprehensive reports. Technologically, this solution delivers two significant advancements: (1) **Iterative Information Retrieval**: Conducts successive rounds of targeted web searches with dynamic adjustments based on interim insights, ensuring comprehensive information coverage and accuracy. (2) **Dynamic Prompt-Guided Model Selection**: Use a hybrid architecture to autonomously select the optimal combination of specialised models based on the requirements and context of specific tasks, thereby enhancing adaptability and effectiveness in various research scenarios.

4.4 Grok DeepSearch

Grok DeepSearch (xAI Team, 2025), developed by xAI, is a computational framework that combines real-time information retrieval with multimodal reasoning to dynamically solve complex and information-rich problems. Technologically, this solution delivers two significant advancements: (1) **Segment-level module processing pipeline**: Upon receiving a query, Grok3 initiates the credibility assessment module to identify and filter out low-quality information. Subsequently, the system’s real-time data acquisition engine gathers multimodal inputs (e.g. text, images, and code) from various sources. Subsequently, employing the sparse attention mechanism, the system undertakes key reasoning subtasks, including data cleaning, cross-source verification, and multimodal integration, in a concurrent manner. Finally, the iterative optimisation process culminates in the generation of structured outputs, encompassing analysis summaries, advanced visualisations (e.g. 3D

trajectories), and verifiable citations. (2) **Dynamic resource allocation**: Capacity for adaptively alternating between lightweight retrieval and intensive analysis modes is noteworthy, and it is further augmented by the incorporation of a secure sandbox environment for computational verification.

4.5 Microsoft Copilot Researcher and Analyst

Microsoft recently introduced two innovative reasoning agents within Microsoft 365 Copilot: Researcher and Analyst (Spataro, 2025). These agents securely and compliantly access users’ work data (such as emails, meeting notes, documents, and chats) as well as web information, delivering on-demand expert knowledge.

Researcher is designed to assist users in tackling complex, multi-step research tasks, delivering insights with unprecedented quality and accuracy. It combines OpenAI’s advanced research models with Microsoft 365 Copilot’s sophisticated orchestration and deep search capabilities. Users can employ Researcher to craft detailed market entry strategies, identify market opportunities for new products by integrating internal and external data, or prepare comprehensive quarterly reports for client reviews. Additionally, Researcher enhances its insights through connectors to third-party data sources such as Salesforce, ServiceNow, and Confluence.

Analyst is built as an advanced data analytics agent that rapidly transforms raw data into valuable insights within minutes. It leverages OpenAI’s o3-mini inference model, specifically optimised for advanced analytical tasks in professional environments. Analyst uses a chain-of-thought reasoning approach, solving problems step-by-step, generating high-quality responses that closely mirror human analytical thinking.

4.6 Qwen Deep Research

Alibaba Qwen recently launched Qwen Deep Research, an advanced research agent powered by its flagship multimodal model Qwen3-235B-A22B. Through reinforcement learning-optimised task scheduling within a unified agent framework, the system demonstrates enhanced autonomous planning and adaptive execution capabilities, enabling rapid completion of complex research workflows. Key technological advancements include: (1) **Dynamic Research Blueprinting** with interactive plan refinement. (2) **Concurrent Task Orchestration** enabling parallel retrieval validation synthesis.

4.7 Kimi K2 Deep Research

Moonshot AI’s Kimi K2 (Team et al., 2025) advances deep research through token-efficient learning, targeted data engineering, a scalable sparse design, and tool-aligned post-training. Concretely, it delivers four contributions: (1) **Token-efficient pretraining**: stabilised optimisation with selective attention regularisation for reliable large-scale training. (2) **Data rewriting**: synthetic rephrasing for knowledge and learning-note style transformation for mathematics, yielding a vetted corpus of about 15.5 trillion tokens across web text, code, math, and knowledge. (3) **Sparse architecture and systems**: mixture-of-experts with multi-head latent attention and an observed sparsity scaling regularity, supported by flexible parallelism and memory-aware execution on H800 clusters. (4) **Post-training for agents**: supervised tool-use trajectories with sandboxed code execution followed by reinforcement learning that combines verifiable rewards for objective tasks and rubric-based self-critique for subjective tasks.

In addition to the pioneering DR services previously discussed, major technology corporations such as Microsoft and ByteDance, alongside emerging startups including Jina AI (AI, 2025a), H2O (H2O.ai, 2025), and Zhipu AI (Zhipu AI, 2025), have also introduced their proprietary DR platforms. The advent of these solutions has spurred considerable global interest, reflected by their rapid proliferation, thereby underscoring both the technological attractiveness and substantial market potential of DR applications. Looking forward, continuous advancements in LLM reasoning, retrieval integration techniques, and multimodal generation are expected to enable DR agents to transcend traditional information retrieval and basic tool invocation tasks. Consequently, DR systems are anticipated to tackle increasingly sophisticated reasoning and complex knowledge-construction challenges, ultimately positioning DR as a foundational technological pillar for next-generation intelligent collaborative research platforms.

5 Benchmarks for DR Agent

Evaluating DR agents requires benchmarks that capture their full research workflow, including multi-step information retrieval, cross-source synthesis, dynamic tool invocation, and structured evidence-grounded report generation. Existing evaluations fall into two main categories. **Question-Answering (QA)** benchmarks range from single-turn factual queries to complex research-style problems, assessing agents’ factual knowledge, domain-specific reasoning, and ability to locate and integrate relevant information. **Task Execution benchmarks** evaluate broader capabilities such as long-horizon planning, multimodal understanding, tool usage, and environment interaction by measuring how well agents carry out end-to-end research tasks. Although long-form generation datasets such as Qasper (Dasigi et al., 2021) and ELI5 (Fan et al., 2019) provide tests of extended output coherence, their free-form nature does not align with the structured evidence-based reporting expected of DR agents. Consequently, there is a pressing need for specialised benchmarks that reflect the multi-stage, multimodal characteristics of DR workflows and ensure rigorous and relevant assessment of agent performance across all phases of autonomous research.

QA Benchmarks. QA benchmarks span a spectrum of complexity, from simple factual recall to multi-hop reasoning and research-style question answering. At the lower end, datasets such as **SimpleQA** (Wei et al., 2024), **TriviaQA** (Joshi et al., 2017), and **PopQA** (Mallen et al., 2023) focus on parametric or single-hop factual recall, evaluating whether models can retrieve short factual answers from memory or minimal context. **Natural Questions (NQ)** (Kwiatkowski et al., 2019) and **TELEQnA** (Maatouk et al., 2023) add complexity by requiring answer extraction from long documents or domain-specific sources. Benchmarks like **HotpotQA** (Yang et al., 2018), **2WikiMultihopQA** (Ho et al., 2020), and **Bamboogle** (Aksitov et al., 2023) emphasize multi-hop reasoning and supporting evidence selection across documents. At the highest level of difficulty lies **Humanity’s Last Exam (HLE)** (Phan et al., 2025), which targets expert-level, open-domain scientific questions crafted by leading professors in various fields. These questions often require multi-turn retrieval, complex inference, and even multimodal understanding. Additionally, BrowseComp (OpenAI Team, 2025) is another challenging benchmark proposed by OpenAI to measure the ability of AI agents to locate hard-to-find information. It retains the answer verifiability of the Simple QA benchmark while filtering out those that can be easily solved by LLMs with web search, thus testing agents’ information retrieval and synthesis capabilities. Despite recent advancements, leading DR agents still exhibit suboptimal performance on the HLE and BrowseComp benchmark compared to human experts. This highlights these two benchmarks as the most critical and unresolved challenges in the evaluation of DR agents.

Task Execution Benchmarks. Task execution benchmarks evaluate an agent’s integrated capabilities in tool use, environment perception, and information filtering. These can be grouped into two subcategories. The

Table 4: Performance of DR agents on major QA benchmarks. The best performance is highlighted in **bold**, and the second-best is indicated with an underline.□ = **Not present**

DR Agent	Base Model	QA Benchmarks				GPQA	Release
		Hotpot	2Wiki	NQ	TQ		
Search-o1 (Li et al., 2025c)	QwQ-32B-preview	57.3	71.4	<u>49.7</u>	<u>74.1</u>	57.9	Jan-2025
Agentic Reasoning (Wu et al., 2025c)	DeepSeek-R1, Qwen2.5	□	□	□	□	67.0	Feb-2025
Grok DeepSearch (xAI Team, 2025)	Grok3	□	□	□	□	84.6	Feb-2025
AgentRxiv (Schmidgall & Moor, 2025)	GPT-4o-mini	□	□	□	□	41.0	Mar-2025
R1-Searcher (Song et al., 2025)	Qwen2.5-7B-Base	71.9	63.8	□	□	□	Mar-2025
ReSearch (Chen et al., 2025b)	Qwen2.5-32B-Inst	67.7	50.0	□	□	□	Mar-2025
Search-R1 (Jin et al., 2025)	Qwen2.5-7B-Inst	34.5	36.9	40.9	55.2	□	Mar-2025
DeepResearcher (Zheng et al., 2025)	Qwen2.5-7B-Inst	64.3	<u>66.6</u>	61.9	85.0	□	Apr-2025
WebThinker (Li et al., 2025d)	QwQ-32B	□	□	□	□	<u>68.7</u>	Apr-2025
SimpleDeepSearch (Sun* et al., 2025)	QwQ-32B	73.5	□	□	□	□	Apr-2025
SWIRL (Goldie et al., 2025)	Gemma-2-27B	<u>72.0</u>	□	□	□	□	Apr-2025
Tool-Star (Dong et al., 2025)	Qwen2.5-3B	51.9	40.0	□	□	□	May-2025

first category comprises general-purpose assistant tasks such as GAIA (Mialon et al., 2023), AssistantBench (Yoran et al., 2024), and Magentic-One (Fourney et al., 2024). These tasks require agents to plan and execute tool-based workflows (for example, searching, browsing, or form filling) within environments that are open-ended and often web-based. Among them, **GAIA** has emerged as the most important benchmark, offering diverse, realistic tasks that are easily human-solvable but remain highly challenging for current agents. The second subcategory focuses on **research and code-oriented tasks**, including **SWE-bench** (Jimenez et al., 2024), **HumanEvalFix** (Muennighoff et al., 2024), **MLGym** (Nathani et al., 2025), **MLE-bench** (Chan & Others, 2025), **MLBench** (Tang et al., 2024), **MLAgentBench** (Huang et al., 2024), and **ScienceAgentBench** (Chen et al., 2025e), which test agents on completing machine learning pipelines, repairing real-world code, or replicating scientific experiments. These tasks require long-horizon planning, precise tool invocation, and often code generation and validation. Additionally, benchmarks like **RE-Bench** (Wijk et al., 2024) and **RESEARCHTOWN** (Yu et al., 2024a) simulate multi-agent research environments, evaluating how well agents collaborate and iterate in multi-role scientific workflows.

As DR agents continue to integrate more interactive tools, future evaluation may expand into GUI-based manipulation environments. Benchmarks such as **OSWorld** (Xie et al., 2024), **WebArena** (Zhou et al., 2024), and **SpaBench** (Chen et al., 2025a) allow agents to control applications or web interfaces directly, opening new avenues for testing embodied research capabilities in realistic, user-facing scenarios.

Table 5: Performance of DR agents on GAIA test and validation sets. The best performance is highlighted in **bold**, and the second-best is indicated with an underline.□ = **Not present**

DR Agent	Base Model	GAIA			Ave.	Release
		Level-1	Level-2	Level-3		
Test set						
MMAC-Copilot (Song et al., 2024)	GPT-3.5, GPT-4	45.16	20.75	6.12	25.91	Mar-2024
H2O.ai DR (H2O.ai, 2025)	Claude3.7-Sonnet	<u>89.25</u>	79.87	61.22	79.73	Mar-2025
Alita (Qiu et al., 2025)	Claude-Sonnet-4, GPT-4o	92.47	<u>71.7</u>	<u>55.1</u>	<u>75.42</u>	May-2025
Agent-KB (Tang et al., 2025b)	GPT-4.1, Claude-3.7	84.91	74.42	57.69	75.15	Jul-2025
O-agents (Zhu et al., 2025a)	Claude-3.7	83.02	74.42	53.85	73.93	Jun-2025
WebDancer (Wu et al., 2025a)	QwQ-32B	61.5	50.0	25.0	51.5	May-2025
WebShaper (Tao et al., 2025)	Qwen-2.5-72B	69.2	63.4	16.6	60.1	Jul-2025
Deep Researcher with Test-Time Diffusion Han et al. (2025)	Gemini-2.5-Pro	□	□	□	69.1	Jul-2025
Cognitive Kernel-Pro (Wan et al., 2025)	Claude-3-7	83.02	68.60	53.85	70.91	Aug-2025
Dev set						
AutoAgent (Tang et al., 2025a)	Claude-Sonnet-3.5	71.7	53.5	26.9	55.2	Feb-2025
OpenAI DR (OpenAI, 2025b)	GPT-o3-customized	78.7	73.2	58.0	67.4	Feb-2025
Manus (Manus AI, 2025)	Claude3.5, GPT-4o	<u>86.5</u>	70.1	57.7	<u>71.4</u>	Mar-2025
OWL (CAMEL-AI.org, 2025)	Claude-3.7-Sonnet	84.9	68.6	42.3	69.7	Mar-2025
H2O.ai DR (H2O.ai, 2025)	h2ogpt-oasst1-512-12b	67.92	67.44	42.31	63.64	Mar-2025
Genspark Super Agent (Team, 2025b)	Claude 3 Opus	87.8	<u>72.7</u>	<u>58.8</u>	73.1	Apr-2025
WebThinker (Li et al., 2025d)	QwQ-32B	53.8	44.2	16.7	44.7	Apr-2025
SimpleDeepSearch (Sun* et al., 2025)	QwQ-32B	50.5	45.8	13.8	43.9	Apr-2025
Alita (Qiu et al., 2025)	Claude-Sonnet-4, GPT-4o	75.15	□	87.27	□	May-2025

Table 6: Overview of nine widely used QA benchmark datasets employed in recent DR-agent studies. The first group covers single-hop QA tasks, while the second group focuses on multi-hop and multi-turn reasoning.

Benchmark	Release Size		Task & Context	Domain	Multi-hop Nums
TriviaQA (Joshi et al., 2017)	2017	95 k	Single-hop retrieval (Long web/Wiki docs)	Open	1
Natural Questions (Kwiatkowski et al., 2019)	2019	307 k	Document answer extraction (Full Wikipedia article)	Open	1
PopQA (Mallen et al., 2023)	2023	14 k	Single-hop parametric recall (None)	Open	1
TELEQnA (Maatouk et al., 2023)	2023	10 k	Domain factual QA (Telecom standards/articles)	Telecom	1
SimpleQA (Wei et al., 2024)	2024	4.3 k	Single-hop factual recall (None / parametric)	Open	1
HotpotQA (Yang et al., 2018)	2018	113 k	Multi-hop reasoning (2 Wikipedia paragraphs)	Open	2
2WikiMultihopQA (Ho et al., 2020)	2020	192 k	Multi-hop reasoning (Retrieval across Wikipedia)	Open	2+
Bamboogle (Aksitov et al., 2023)	2023	125	Compositional reasoning (Online search)	Open	2-3
Humanity’s Last Exam (Phan et al., 2025)	2025	2.5 k	Expert-level multi-turn (Mixed external sources)	Multi-discipline	2+

6 Challenge and Future Directions

Despite the rapid evolution of DR agents and their demonstrated efficacy in automating multi-step information discovery and synthesis, two overarching challenges persist, defining the roadmap for future innovation. First, the breadth and depth of accessible information remain tightly constrained by reliance on static knowledge repositories or conventional search interfaces. Second, the efficiency and robustness of execution workflows and system architectures are limited by linear planning paradigms and monolithic agent designs. Addressing these challenges will be critical to enabling DR agents to function as truly autonomous, adaptable research assistants capable of navigating complex, heterogeneous data landscapes and orchestrating high-throughput, parallelised reasoning processes.

Broaden Information Source. To meet the information needs of complex tasks, current DR agents adopt static knowledge bases (such as the RAG method) or rely exclusively on search engines and browsers; the former approach is insufficient, while the latter is confined to publicly available web content, thereby significantly constraining their information-acquisition capabilities. This inherent limitation renders them incapable of retrieving information concealed behind applications, proprietary interfaces or specialised databases. For example, conventional browsing and search techniques cannot penetrate enterprise software, mobile applications, or subscription-only services, such as the Bloomberg Terminal, thereby precluding access to critical, real-time market intelligence. In order to surmount this limitation, it is imperative to integrate a more granular and extensive range of modular tools via MCPs. This approach enables agents to dynamically access specialised tools and resources beyond the scope of standard browsers or search engines. Such resources may include proprietary applications, databases, or APIs, thereby facilitating the retrieval of previously inaccessible data. Consequently, DR agents have the capacity to deliver more precise, adaptive, and context-aware interactions, thereby effectively fulfilling diverse and complex user requirements.

Following the integration of proprietary APIs and databases, the rate-limiting factor in the workflow shifts from data acquisition to webpage interaction efficiency. Conventional human-centred browsers create a further

bottleneck for agents. Because they optimise for visual rendering rather than programmatic control, they suffer from sluggish page loads, fragile element locators that shift with every layout change, and aggressive anti-bot defences that often break automated sessions. These shortcomings translate into high latency, unstable scraping and limited parallelism whenever DR agents try to harvest data at scale. To address this bottleneck, researchers have begun to design **AI-native browsers** such as Browserbase (Adam et al., 2024), Browser Use (Müller & Žunič, 2024), Dia, Fellou (Team, 2025a), and the Comet (Team, 2024) from Perplexity. expose a stable, structured DOM view that agents can traverse programmatically (Adam et al., 2024; Müller & Žunič, 2024; Team, 2024). (Adam et al., 2024; Team, 2025a) supply explicit API hooks for clicking elements and filling forms, which removes the need for brittle coordinate-based actions. (Adam et al., 2024) further executes pages asynchronously in a headless container, reducing load-time variance and avoiding the overhead of a visible interface. (Adam et al., 2024) embeds a vision-language model that tracks dynamic page changes and automatically resolves login gates and anti-bot challenges. (Müller & Žunič, 2024; Team, 2024) coordinates dozens of tabs in parallel, allowing DR agents to interact with private dashboards, single-page applications, and interactive visualisations at scale. In combination, these capabilities eliminate the delays and fragility that arise when conventional, human-centred browsers sit between the agent and newly unlocked proprietary data sources.

Fact Checking. To further boost factual accuracy, the latest methods add a structured verification loop and self-reflection abilities on top of multi-step retrieval. Concretely, once an agent has drafted a preliminary answer, it does not rush to deliver a verdict. Instead, it proactively launches cross-checks: it looks for independent sources that confirm the same fact and searches for evidence of contradictions. Grok DeepSearch, for example, follows this strategy—it rates the credibility of every source, inspects consistency through as many as seven layers of depth, and verifies each key claim across multiple origins (xAI Team, 2025). This multi-source cross-validation sharply reduces single-source errors and raises answer reliability. At the same time, agents have begun to reflect on their own reasoning. During inference, they inspect and test intermediate results, much like a human researcher’s reflective thinking. Zhipu’s Rumination model (Zhipu AI, 2025), for instance, pauses after concluding, keeps searching to check whether that conclusion holds, and only then finalises the answer. Such introspection is typically encouraged by adding correctness-oriented rewards in reinforcement learning. If the model detects conflict or uncertainty, it replans its retrieval strategy and, when necessary, backtracks to revise earlier inferences (OpenAI, 2025b). Through this blend of structured verification and self-reflection, research agents now attain an unprecedented level of rigour in fact-checking: they not only supply an answer but also explain why it is trustworthy, dramatically lowering factual errors and hallucinations. In short, modern agents can lay out a search plan, adapt queries as intermediate evidence comes in, and—where needed—rewind prior steps to recover missing information (OpenAI, 2025b).

Asynchronous Parallel Execution. To address the limitation that most existing DR agents rely exclusively on linear task planning, i.e. the sequential execution of subtasks, we introduce two possible methodologies. These methods overcome the inherent efficiency and robustness constraints of purely linear strategies and enable both the exploitation of parallelism and the implementation of dynamic adjustments during task execution. Firstly, an asynchronous, parallel architecture leveraging advanced task-modelling structures, such as directed acyclic graphs (DAGs), presents a promising future direction which could enable parallel execution and dynamic prioritisation of subtasks, effectively managing complex interdependencies among tasks and facilitating potentially sophisticated planning capabilities such as replanning. Secondly, a learned scheduling agent, trained via reinforcement learning to allocate subtasks and adjust execution order based on runtime performance signals (e.g. execution latency), could be proposed. By treating scheduling decisions as

actions in an RL environment, the agent progressively discovers policies that balance parallelism, resource utilisation, and task criticality, yielding more robust and efficient end-to-end research workflows.

Tool-Integrated Reasoning. A fundamental challenge in developing effective DR agents lies in the implementation of Tool-Integrated Reasoning (TIR), a paradigm that extends beyond simple tool usage to encompass complex, multi-step reasoning with dynamic tool integration. TIR requires agents to not only invoke appropriate tools in logical sequence but also to adaptively adjust their reasoning pathways based on intermediate results. Traditional supervised fine-tuning approaches have demonstrated limited generalisation capabilities in tool-based reasoning tasks, often leading to over-reasoning or inappropriate tool selection. Recent research by (Qian et al., 2025) has shown that reinforcement learning frameworks with carefully designed reward structures can significantly enhance models’ tool reasoning abilities. By incorporating fine-grained rewards that evaluate not only final answer correctness but also tool selection appropriateness, parameter specification accuracy, and reasoning efficiency, TIR-optimised agents have demonstrated performance improvements of 15-17% across multiple benchmarks. Furthermore, these agents exhibit superior generalisation to unseen tools and tasks, more rational invocation patterns, and better balance between tool utilisation and self-knowledge. Implementing TIR effectively within DR agents represents a critical step toward achieving truly autonomous research assistants capable of navigating complex information landscapes with minimal human intervention.

Benchmark Misalignment. Most public DR evaluations remain anchored in traditional QA suites whose items are harvested chiefly from static corpora such as Wikipedia. Since a considerable amount of this content is now embedded in backbone model parameters, current competitive agents can often answer directly from memory, bypassing any research procedure and thus inflating their performance. To probe genuine capabilities of retrieval, reasoning and tool usage, the field of DR urgently needs open-web, time-sensitive benchmarks. From this perspective, BrowseComp (OpenAI Team, 2025) constitutes a meaningful step forward by filtering out questions solvable with parametric knowledge and forcing agents to locate hard-to-find information online. Besides, a complementary direction is a continually refreshed leaderboard that updates problems from the latest web environment and events, deterring benchmark hacking through parametric memorisation.

Beyond parametric knowledge hacking of QA benchmark, the metrics of the most existing DR research still collapse open-ended research workflows into narrowly scoped QA prompts or rudimentary GUI-control tasks, overlooking the paradigm’s defining outcome, a structured, multi-modal research report that weaves together textual narrative, tables, figures, and citations. Since the metrics of these benchmarks centre almost exclusively on information retrieval and extraction and tool invocation, they under-assess higher-level competencies such as evidence aggregation across heterogeneous sources, cross-modal synthesis, and discourse-level organisation. Thus, a key research priority is the development of comprehensive benchmarks that evaluate DR agents’ capacity for end-to-end report generation, encompassing long-form narrative, integrated tables and figures, and multimodal coherence, thereby assessing factual accuracy, discourse structure, and cross-modal alignment within a single task.

Parametric Optimisation of Multi-Agent Architectures. End-to-end RL has been demonstrated by OpenAI (OpenAI, 2025b; Ouyang et al., 2025) to significantly enhance the reasoning capabilities of backbone models for DR tasks, a result successfully replicated by several open-source initiatives. However, current implementations predominantly utilise single-agent architectures, requiring the backbone model to simultaneously manage planning, tool invocation, and report generation. This multitasking places excessive computational and cognitive demands on backbone models, thereby reducing their efficiency and robustness.

Distributing workloads across multiple specialised agents has shown promising improvements in system performance (Wang et al., 2024a), yet achieving effective end-to-end training and efficient coordination among multiple agents remains a critical open challenge.

To optimize multi-agent architectures for DR tasks, we propose two promising future directions: (i) adopting **hierarchical reinforcement learning** (HRL), which introduces layered internal reward mechanisms that facilitate efficient feedback propagation and foster cooperative learning among agents; or implementing a post-training optimization pipeline consisting of multiple refinement stages specifically tailored for DR tasks, which could iteratively enhance inter-agent interactions and thus improve overall system stability and adaptability; and (ii) employing an **RL-based dedicated scheduling agent designed to dynamically allocate subtasks and adjust execution order based on real-time performance metrics**. By modelling scheduling decisions as actions within an RL framework, this method progressively learns adaptive policies that optimally balance parallel execution, resource utilisation, and task prioritisation, enhancing both the robustness and efficiency of end-to-end research workflows.

Self-Evolving Language Model Agents. Although initial attempts at self-evolution methods for DR agents have emerged, exemplified by simulated collaborative platforms such as AgentRxiv (Schmidgall & Moor, 2025) that facilitate online sharing and reuse of structured research experiences, the paradigm remains underdeveloped and narrowly focused on only the case-based reasoning paradigm. Similarly, CycleResearcher (Weng et al., 2024) enables the entire research process simulation (research-evaluation-refine) through iterative preference learning with a robust verifier (Zhu et al., 2025b), representing a significant step toward fully automated scientific inquiry and sharing the similar self-evolution concept with AlphaEvolve (Novikov et al., 2025).

To fully realise the potential of self-evolution in DR agents, future research should expand the self-evolution method along two complementary directions. (i) **Comprehensive case-based reasoning framework.** Case-based reasoning approaches (Aamodt & Plaza, 1994) leverage hierarchical experience traces, including planning trajectories and structured tool invocation logs, and employ advanced retrieval and selection mechanisms to enable fine-grained, context-specific adaptation. (ii) **Autonomous workflow evolution** promises enhanced efficiency and flexibility. By representing agent workflows as mutable structures such as trees or graphs, researchers can apply evolutionary algorithms or adaptive graph optimisation to explore, modify and refine execution plans dynamically. Pursuing both directions in tandem will strengthen the robustness of frameworks and reduce the reliance on data and computation resources.

7 Conclusion

LLM-driven Deep Research Agents represent an emerging paradigm for automated research support, integrating advanced techniques such as iterative information retrieval, long-form content generation, autonomous planning, and sophisticated tool utilisation. In this survey, we systematically reviewed recent advancements in DR agents, categorising existing methodologies into prompt-based, fine-tuning-based, and reinforcement learning-based approaches from the perspectives of information retrieval and report generation. Non-parametric methods utilise LLMs and carefully designed prompts to achieve efficient and cost-effective deployment, making them suitable for rapid prototyping. In contrast, fine-tuning and reinforcement learning approaches explicitly optimise model parameters, significantly enhancing the agents’ reasoning and decision-making capabilities. We also examined prominent DR agent systems developed by industry leaders and discussed their technical implementations, strengths, and limitations.

Limitation

Despite notable progress, key challenges remain, including limited generalisation across diverse tasks, inflexible task workflows, difficulty in integrating granular external tools, and substantial computational complexity associated with advanced planning and optimisation. Future research directions thus emphasise broader and more flexible tool integration through modular capability providers (e.g., Operator-based architectures), development of asynchronous and parallel planning frameworks (e.g., Directed Acyclic Graph-based approaches), and sophisticated end-to-end optimisation methods for multi-agent architectures, such as hierarchical reinforcement learning or multi-stage fine-tuning pipelines. With continued advancements in LLM technologies, DR agents have significant potential to transform complex research workflows, enhance human productivity, and drive innovation across academic and industrial domains.

References

- Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI communications*, 7(1):39–59, 1994.
- McQuilkin Adam, Kamath Anirudh, McGuire Sean, and Chance Sophie. Browserbase: A web browser for your ai, 2024. URL <https://github.com/browserbase>.
- Jina AI. Jina ai. <https://jina.ai/deepsearch/>, 2025a. Accessed: 2025-04-28.
- Kortix AI. Suna: Open source generalist ai agent. <https://github.com/kortix-ai/suna>, 2025b. Accessed: 2025-05-28.
- Renat Aksitov, Sobhan Miryoosefi, Zonglin Li, Daliang Li, Sheila Babayan, Kavya Kopparapu, Zachary Fisher, Ruiqi Guo, Sushant Prakash, Pranesh Srinivasan, Manzil Zaheer, Felix Yu, and Sanjiv Kumar. Rest meets react: Self-improvement for multi-step reasoning llm agent, 2023. URL <https://arxiv.org/abs/2312.10003>.
- Mohammad Aliannejadi, Zahra Abbasiantaeb, Shubham Chatterjee, Jeffery Dalton, and Leif Azzopardi. Trec ikat 2023: The interactive knowledge assistance track overview. *arXiv preprint arXiv:2401.01330*, 2024.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*, 2023.
- Agent Team at InclusionAI. Aworld: Enabling agent self-improvement through interactive experience with dynamic runtime, 2025. URL <https://github.com/inclusionAI/AWorld>.
- Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. Longalign: A recipe for long context alignment of large language models. *arXiv preprint arXiv:2401.18058*, 2024a.
- Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, et al. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. *arXiv preprint arXiv:2412.15204*, 2024b.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter,

- Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- CAMEL-AI.org. Owl: Optimized workforce learning for general multi-agent assistance in real-world task automation. <https://github.com/camel-ai/owl>, 2025. Accessed: 2025-03-07.
- Hyungjoo Chae, Yongho Song, Kai Tzu-iunn Ong, Taeyoon Kwon, Minjin Kim, Youngjae Yu, Dongha Lee, Dongyeop Kang, and Jinyoung Yeo. Dialogue chain-of-thought distillation for commonsense-aware conversational agents. *arXiv preprint arXiv:2310.09343*, 2023.
- Jun Shern Chan and Others. Mle-bench: Evaluating machine learning agents on machine learning engineering, 2025.
- Jingxuan Chen, Derek Yuen, Bin Xie, Yuhao Yang, Gongwei Chen, Zhihao Wu, Li Yixing, Xurui Zhou, Weiwen Liu, Shuai Wang, Kaiwen Zhou, Rui Shao, Liqiang Nie, Yasheng Wang, Jianye Hao, Jun Wang, and Kun Shao. Spa-bench: A comprehensive benchmark for smartphone agent evaluation. *arXiv preprint arXiv:2410.15164*, 2025a.
- Mingyang Chen, Tianpeng Li, Haoze Sun, Yijie Zhou, Chenzheng Zhu, Fan Yang, Zenan Zhou, Weipeng Chen, Haofen Wang, Jeff Z Pan, et al. Learning to reason with search for llms via reinforcement learning. *arXiv preprint arXiv:2503.19470*, 2025b.
- Yiqun Chen, Lingyong Yan, Weiwei Sun, Xinyu Ma, Yi Zhang, Shuaiqiang Wang, Dawei Yin, Yiming Yang, and Jiaxin Mao. Improving retrieval-augmented generation through multi-agent reinforcement learning. *arXiv preprint arXiv:2501.15228*, 2025c.
- Zhixun Chen, Ming Li, Yuxuan Huang, Yali Du, Meng Fang, and Tianyi Zhou. Atlas: Agent tuning via learning critical steps. *arXiv preprint arXiv:2503.02197*, 2025d.
- Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao, Chen Wei, Zitong Lu, Vishal Dey, Mingyi Xue, Frazier N. Baker, Benjamin Burns, Daniel Adu-Ampratwum, Xuhui Huang, Xia Ning, Song Gao, Yu Su, and Huan Sun. Scienceagentbench: Toward rigorous assessment of language agents for data-driven scientific discovery. *arXiv preprint arXiv:2410.05080*, 2025e.
- DanielWalnut. Deerflow. <https://github.com/bytedance/deer-flow>, 2025. Accessed: 2025-05-28.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A Smith, and Matt Gardner. A dataset of information-seeking questions and answers anchored in research papers. *arXiv preprint arXiv:2105.03011*, 2021.
- Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui Zhou, Zhicheng Dou, and Ji-Rong Wen. Tool-star: Empowering llm-brained multi-tool reasoner via reinforcement learning. *arXiv preprint arXiv:2505.16410*, 2025.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. Eli5: Long form question answering. *arXiv preprint arXiv:1907.09190*, 2019.
- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6491–6501, 2024.

- Zhangyin Feng, Xiaocheng Feng, Dezhi Zhao, Maojin Yang, and Bing Qin. Retrieval-generation synergy augmented large language models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 11661–11665. IEEE, 2024.
- Adam Fourney, Gagan Bansal, Hussein Mozannar, Cheng Tan, Eduardo Salinas, Erkang, Zhu, Friederike Niedtner, Grace Proebsting, Griffin Bassman, Jack Gerrits, Jacob Alber, Peter Chang, Ricky Loynd, Robert West, Victor Dibia, Ahmed Awadallah, Ece Kamar, Rafah Hosn, and Saleema Amershi. Magentic-one: A generalist multi-agent system for solving complex tasks. *arXiv preprint arXiv:2411.04468*, 2024.
- Tong Fu, Liquan Chen, Zhangjie Fu, Kunliang Yu, and Yu Wang. Ccnet: Cnn model with channel attention and convolutional pooling mechanism for spatial image steganalysis. *Journal of Visual Communication and Image Representation*, 88:103633, 2022.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Haofen Wang, and Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2, 2023.
- Xinyu Geng, Peng Xia, Zhen Zhang, Xinyu Wang, Qiuchen Wang, Ruixue Ding, Chenxi Wang, Jialong Wu, Yida Zhao, Kuan Li, et al. Webwatcher: Breaking new frontiers of vision-language deep research agent. *arXiv preprint arXiv:2508.05748*, 2025.
- Anna Goldie, Azalia Mirhoseini, Hao Zhou, Irene Cai, and Christopher D Manning. Synthetic data generation & multi-step rl for reasoning & tool use. *arXiv preprint arXiv:2504.04736*, 2025.
- Peiyuan Gong, Jiamian Li, and Jiabin Mao. Cosearchagent: a lightweight collaborative search agent with large language models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2729–2733, 2024.
- Google. Announcing the Agent2Agent Protocol (A2A), 2025. URL <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interopability/>. Accessed: 2025-04-22.
- Google Team. Introducing gemini deep research. <https://gemini.google/overview/deep-research/>, 2025. Accessed: 2025-04-06.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, et al. Towards an ai co-scientist. *arXiv preprint arXiv:2502.18864*, 2025.
- Antoine Grosnit, Alexandre Maraval, James Doran, Giuseppe Paolo, Albert Thomas, Refinath Shahul Hameed Nabeezath Beevi, Jonas Gonzalez, Khyati Khandelwal, Ignacio Iacobacci, Abdelhakim Benechehab, Hamza Cherkaoui, Youssef Attia El-Hili, Kun Shao, Jianye Hao, Jun Yao, Balazs Kegl, Haitham Bou-Ammar, and Jun Wang. Large language models orchestrating structured reasoning achieve kaggle grandmaster level, 2024. URL <https://arxiv.org/abs/2411.03562>.
- Xinyan Guan, Jiali Zeng, Fandong Meng, Chunlei Xin, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun, and Jie Zhou. Deeprag: Thinking to retrieval step by step for large language models. *arXiv preprint arXiv:2502.01142*, 2025.
- Siyuan Guo, Cheng Deng, Ying Wen, Hechang Chen, Yi Chang, and Jun Wang. Ds-agent: Automated data science by empowering large language models with case-based reasoning, 2024. URL <https://arxiv.org/abs/2402.17453>.

- Siyuan Guo, Huiwu Liu, Xiaolong Chen, Yuming Xie, Liang Zhang, Tao Han, Hechang Chen, Yi Chang, and Jun Wang. Optimizing case-based reasoning system for functional test script generation with large language models. *arXiv preprint arXiv:2503.20576*, 2025.
- H2O.ai. H2o.ai, 2025. URL <https://h2o.ai/>. Accessed: 2025-04-28.
- Rujun Han, Yanfei Chen, Zoey CuiZhu, Lesly Miculicich, Guan Sun, Yuanjun Bi, Weiming Wen, Hui Wan, Chunfeng Wen, Solène Maître, et al. Deep researcher with test-time diffusion. *arXiv preprint arXiv:2507.16075*, 2025.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060*, 2020.
- Sheryl Hsu, Omar Khattab, Chelsea Finn, and Archit Sharma. Grounding by trying: Llms with reinforcement learning-enhanced retrieval. *arXiv preprint arXiv:2410.23214*, 2024.
- Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv preprint arXiv:2501.03262*, 2025.
- Jerry Huang, Siddarth Madala, Risham Sidhu, Cheng Niu, Julia Hockenmaier, and Tong Zhang. Rag-rl: Advancing retrieval-augmented generation via rl and curriculum learning. *arXiv preprint arXiv:2503.12759*, 2025.
- Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. Mlagentbench: Evaluating language agents on machine learning experimentation. *arXiv preprint arXiv:2310.03302*, 2024.
- Tatsuro Inaba, Hirokazu Kiyomaru, Fei Cheng, and Sadao Kurohashi. Multitool-cot: Gpt-3 can use multiple external tools with chain of thought prompting. *arXiv preprint arXiv:2305.16896*, 2023.
- Shayekh Bin Islam, Md Asib Rahman, KSM Hossain, Enamul Hoque, Shafiq Joty, and Md Rizwan Parvez. Open-rag: Enhanced retrieval-augmented reasoning with open-source large language models. *arXiv preprint arXiv:2410.01782*, 2024.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43, 2023.
- Pengcheng Jiang. Deepretrieval: Powerful query generation for information retrieval with reinforcement learning. *arXiv preprint arXiv:2503.00223*, 2025.
- Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. Active retrieval augmented generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7969–7992, 2023.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2024.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Dong Wang, Hamed Zamani, and Jiawei Han. Search-rl: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*, 2025.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.

- Gangwoo Kim, Sungdong Kim, Byeongguk Jeon, Joonsuk Park, and Jaewoo Kang. Tree of clarifications: Answering ambiguous questions with retrieval-augmented large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 996–1009, 2023.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Chengpeng Li, Mingfeng Xue, Zhenru Zhang, Jiayi Yang, Beichen Zhang, Xiang Wang, Bowen Yu, Binyuan Hui, Junyang Lin, and Dayiheng Liu. Start: Self-taught reasoner with tools. *arXiv preprint arXiv:2503.04625*, 2025a.
- Kuan Li, Zhongwang Zhang, Huifeng Yin, Liwen Zhang, Litu Ou, Jialong Wu, Wenbiao Yin, Baixuan Li, Zhengwei Tao, Xinyu Wang, et al. Websailor: Navigating super-human reasoning for web agent. *arXiv preprint arXiv:2507.02592*, 2025b.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. Search-o1: Agentic search-enhanced large reasoning models. *arXiv preprint arXiv:2501.05366*, 2025c.
- Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yutao Zhu, Yongkang Wu, Ji-Rong Wen, and Zhicheng Dou. Webthinker: Empowering large reasoning models with deep research capability, 2025d. Notion Blog.
- Xinbin Liang, Jinyu Xiang, Zhaoyang Yu, Jiayi Zhang, and Sirui Hong. Openmanus: An open-source framework for building general ai agents. <https://github.com/mannaandpoem/OpenManus>, 2025. Accessed: 2025-04-06.
- Jingyang Lin, Andy Wong, Tian Xia, Shenghua He, Hui Wei, Mei Han, and Jiebo Luo. Facilitating long context understanding via supervised chain-of-thought reasoning. *arXiv preprint arXiv:2502.13127*, 2025.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Richard James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, et al. Ra-dit: Retrieval-augmented dual instruction tuning. In *The Twelfth International Conference on Learning Representations*, 2023.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- Ali Maatouk, Fadhel Ayed, Nicola Piovesan, Antonio De Domenico, Merouane Debbah, and Zhi-Quan Luo. Teleqna: A benchmark dataset to assess large language models telecommunications knowledge, 2023.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories, 2023.
- Manus AI. Leave it to manus. <https://manus.im/>, 2025. Accessed: 2025-04-06.
- Martin. Agenticseek: Private, local manus alternative. <https://github.com/Fosowl/agenticSeek>, 2025. Accessed: 2025-05-28.
- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. Gaia: a benchmark for general ai assistants. In *The Twelfth International Conference on Learning Representations*, 2023.

- Microsoft. Introducing researcher and analyst in microsoft 365 copilot. <https://www.microsoft.com/en-us/microsoft-365/blog/2025/03/25/introducing-researcher-and-analyst-in-microsoft-365-copilot/>, March 2025. Accessed: 2025-04-28.
- Moonshot AI. Kimi-researcher: End-to-end RL training for emerging agentic capabilities. <https://moonshotai.github.io/Kimi-Researcher/>, June 2025. Accessed: 2025-08-18.
- Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. Octopack: Instruction tuning code large language models. *arXiv preprint arXiv:2308.07124*, 2024.
- Magnus Müller and Gregor Žunič. Browser use: Enable ai to control your browser, 2024. URL <https://github.com/browser-use/browser-use>.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vincent Moens, Amar Budhiraja, Despoina Magka, Vladislav Vorotilov, Gaurav Chaurasia, Dieuwke Hupkes, Ricardo Silveira Cabral, Tatiana Shavrina, Jakob Foerster, Yoram Bachrach, William Yang Wang, and Roberta Raileanu. Mlgym: A new framework and benchmark for advancing ai research agents. *arXiv preprint arXiv:2502.14499*, 2025.
- Nouswise. Nouswise — the research agent grounded in the information you trust the most. <https://nouswise.com/homepage>, 2025. Accessed: 2025-08-18.
- Alexander Novikov, Ngân Vu, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian, et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *Google DeepMind*, 2025.
- OpenAI. Introducing chatgpt agent: bridging research and action. <https://openai.com/index/introducing-chatgpt-agent/>, July 2025a. Accessed: 2025-08-18.
- OpenAI. Introducing deep research. <https://openai.com/index/introducing-deep-research/>, 2025b. Accessed: 2025-04-06.
- OpenAI Team. Browsecomp: a benchmark for browsing agents. <https://openai.com/index/browsecomp/>, 2025. Accessed: 2025-04-29.
- Jie Ouyang, Ruirao Yan, Yucong Luo, Mingyue Cheng, Qi Liu, Zirui Liu, Shuo Yu, and Daoyu Wang. Training powerful llm agents with end-to-end reinforcement learning. <https://github.com/0russwest0/Agent-R1>, 2025. Accessed: 2025-04-06.
- Perplexity Team. Introducing perplexity deep research. <https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research>, 2025. Accessed: 2025-04-06.
- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, et al. Humanity’s last exam. *arXiv preprint arXiv:2501.14249*, 2025.

- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiushi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. Toolrl: Reward is all tool learning needs, 2025. URL <https://arxiv.org/abs/2504.13958>.
- Jiahao Qiu, Xuan Qi, Tongcheng Zhang, Xinzhe Juan, Jiacheng Guo, Yifu Lu, Yimin Wang, Zixin Yao, Qihan Ren, Xun Jiang, et al. Alita: Generalist agent enabling scalable agentic reasoning with minimal predefinition and maximal self-evolution. *arXiv preprint arXiv:2505.20286*, 2025.
- Changle Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and Ji-Rong Wen. Tool learning with large language models: A survey. *Frontiers of Computer Science*, 19(8):198343, 2025.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741, 2023.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools, 2023. URL <https://arxiv.org/abs/2302.04761>.
- Samuel Schmidgall and Michael Moor. Agentrxiv: Towards collaborative autonomous research. *arXiv preprint arXiv:2503.18102*, 2025.
- Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Zicheng Liu, and Emad Barsoum. Agent laboratory: Using llm agents as research assistants. *arXiv preprint arXiv:2501.04227*, 2025.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Yijia Shao, Yucheng Jiang, Theodore A. Kanell, Peter Xu, Omar Khattab, and Monica S. Lam. Assisting in writing wikipedia-like articles from scratch with large language models. *arXiv preprint arXiv:2402.14207*, 2024a.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024b.
- Zhili Shen, Chenxin Diao, Pavlos Vougiouklis, Pascual Merita, Shriram Piramanayagam, Damien Graux, Dandan Tu, Zeren Jiang, Ruofei Lai, Yang Ren, et al. Gear: Graph-enhanced agent for retrieval-augmented generation. *arXiv preprint arXiv:2412.18431*, 2024.
- Wenxuan Shi, Haochen Tan, Chuqiao Kuang, Xiaoguang Li, Xiaozhe Ren, Chen Zhang, Hanting Chen, Yasheng Wang, Lifeng Shang, Fisher Yu, et al. Pangu deepdiver: Adaptive search intensity scaling via open-web reinforcement learning. *arXiv preprint arXiv:2505.24332*, 2025.
- Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*, 2025.
- Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. R1-searcher: Incentivizing the search capability in llms via reinforcement learning. *arXiv preprint arXiv:2503.05592*, 2025.

- Zirui Song, Yaohang Li, Meng Fang, Zhenhao Chen, Zecheng Shi, Yuan Huang, and Ling Chen. Mmac-copilot: Multi-modal agent collaboration operating system copilot. *arXiv preprint arXiv:2404.18074*, 2024.
- Jared Spataro. Introducing Researcher and Analyst in Microsoft 365 Copilot, March 2025. URL <https://www.microsoft.com/en-us/microsoft-365/blog/2025/03/25/introducing-researcher-and-analyst-in-microsoft-365-copilot/>.
- Md Arafat Sultan, Jatin Ganhotra, and Ramón Fernandez Astudillo. Structured chain-of-thought prompting for few-shot generation of content-grounded qa conversations. *arXiv preprint arXiv:2402.11770*, 2024.
- Shuang Sun*, Huatong Song*, Yuhao Wang, Ruiyang Ren, Jinhao Jiang, Junjie Zhang, Lei Fang, Zhongyuan Wang, and Ji-Rong Wen Wayne Xin Zhao. Simpledeepsearcher: Deep information seeking via web-powered reasoning trajectory synthesis. 2025. URL <https://github.com/RUCAIBox/SimpleDeepSearcher>.
- Jiabin Tang, Tianyu Fan, and Chao Huang. Autoagent: A fully-automated and zero-code framework for llm agents, 2025a. URL <https://arxiv.org/abs/2502.05957>.
- Xiangru Tang, Yuliang Liu, Zefan Cai, Yanjun Shao, Junjie Lu, Yichi Zhang, Zexuan Deng, Helan Hu, Kaikai An, Ruijun Huang, Shuzheng Si, Sheng Chen, Haozhe Zhao, Liang Chen, Yan Wang, Tianyu Liu, Zhiwei Jiang, Baobao Chang, Yin Fang, Yujia Qin, Wangchunshu Zhou, Yilun Zhao, Arman Cohan, and Mark Gerstein. ML-bench: Evaluating large language models and agents for machine learning tasks on repository-level code. *arXiv preprint arXiv:2311.09835*, 2024.
- Xiangru Tang, Tianrui Qin, Tianhao Peng, Ziyang Zhou, Daniel Shao, Tingting Du, Xinming Wei, Peng Xia, Fang Wu, He Zhu, et al. Agent kb: Leveraging cross-domain experience for agentic problem solving. *arXiv preprint arXiv:2507.06229*, 2025b.
- Zhengwei Tao, Jialong Wu, Wenbiao Yin, Junkai Zhang, Baixuan Li, Haiyang Shen, Kuan Li, Liwen Zhang, Xinyu Wang, Yong Jiang, et al. Webshaper: Agentic data synthesizing via information-seeking formalization. *arXiv preprint arXiv:2507.15061*, 2025.
- Fellou AI Team. Fellou: The world’s first agentic browser, 2025a. URL <https://github.com/FellouAI>.
- Genspark Team. Genspark super agent with enhancements in mixture of agents, 2025b. URL <https://www.genspark.ai/>.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- MiroMind Foundation Model Team and MiroMind AI Infra Team. Mirorl: An mcp-first reinforcement learning framework for deep research agent. <https://github.com/MiroMindAI/MiroRL>, 2025.
- Perplexity Team. Comet: A browser for agentic search by perplexity, 2024. URL <https://www.perplexity.ai/comet>.
- Kaiyang Wan, Honglin Mu, Rui Hao, Haoran Luo, Tianle Gu, and Xiuying Chen. A cognitive writing perspective for constrained long-form text generation. *arXiv preprint arXiv:2502.12568*, 2025.
- Junyang Wang, Haiyang Xu, Haitao Jia, Xi Zhang, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. Mobile-agent-v2: Mobile device operation assistant with effective navigation via multi-agent collaboration. *arXiv preprint arXiv:2406.01014*, 2024a.

- Keheng Wang, Feiyu Duan, Peiguang Li, Sirui Wang, and Xunliang Cai. Llms know what they need: Leveraging a missing information guided framework to empower retrieval-augmented generation. *arXiv preprint arXiv:2404.14043*, 2024b.
- Liang Wang, Haonan Chen, Nan Yang, Xiaolong Huang, Zhicheng Dou, and Furu Wei. Chain-of-retrieval augmented generation. *arXiv preprint arXiv:2501.14342*, 2025a.
- Xindi Wang, Mahsa Salmani, Parsa Omid, Xiangyu Ren, Mehdi Rezagholizadeh, and Armaghan Eshaghi. Beyond the limits: A survey of techniques to extend the context length in large language models. *arXiv preprint arXiv:2402.02244*, 2024c.
- Yaoxiang Wang, Zhiyong Wu, Junfeng Yao, and Jinsong Su. Tdag: A multi-agent framework based on dynamic task decomposition and agent generation. *Neural Networks*, pp. 107200, 2025b.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL <https://arxiv.org/abs/2201.11903>.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. Measuring short-form factuality in large language models. *arXiv preprint arXiv:2411.04368*, 2024.
- Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo Zhang, Jindong Wang, Yue Zhang, and Linyi Yang. Cyclere searcher: Improving automated research via automated review. *arXiv preprint arXiv:2411.00816*, 2024.
- Hjalmar Wijk, Tao Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence Chan, Michael Chen, Josh Clymer, Jai Dhyani, Elena Ericeva, Katharyn Garcia, Brian Goodrich, Nikola Jurkovic, Megan Kinniment, Aron Lajko, Seraphina Nix, Lucas Sato, William Saunders, Maksym Taran, Ben West, and Elizabeth Barnes. Re-bench: Evaluating frontier ai r&d capabilities of language model agents against human experts. *arXiv preprint arXiv:2411.15114*, 2024.
- Ryan Wong, Jiawei Wang, Junjie Zhao, Li Chen, Yan Gao, Long Zhang, Xuan Zhou, Zuo Wang, Kai Xiang, Ge Zhang, et al. Widesearch: Benchmarking agentic broad info-seeking. *arXiv preprint arXiv:2508.07999*, 2025.
- Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, et al. Webdancer: Towards autonomous information seeking agency. *arXiv preprint arXiv:2505.22648*, 2025a.
- Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, et al. Webwalker: Benchmarking llms in web traversal. *arXiv preprint arXiv:2501.07572*, 2025b.
- Junde Wu, Jiayuan Zhu, and Yuyuan Liu. Agentic reasoning: Reasoning llms with tools for the deep research. *arXiv preprint arXiv:2502.04644*, 2025c.
- Shirley Wu, Shiyu Zhao, Qian Huang, Kexin Huang, Michihiro Yasunaga, Kaidi Cao, Vassilis N. Ioannidis, Karthik Subbian, Jure Leskovec, and James Zou. Avatar: Optimizing llm agents for tool usage via contrastive reasoning. *arXiv preprint arXiv:2406.11200*, 2024.

- xAI Team. Introducing grok deepsearch. <https://x.ai/news/grok-3>, 2025. Accessed: 2025-04-06.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*, 2024.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*, 2018.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2023.
- Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Bogin, Ofir Press, and Jonathan Berant. Assistantbench: Can web agents solve realistic and time-consuming tasks? *arXiv preprint arXiv:2407.15711*, 2024.
- Haofei Yu, Zhaochen Hong, Zirui Cheng, Kunlun Zhu, Keyang Xuan, Jinwei Yao, Tao Feng, and Jiaxuan You. Researchtown: Simulator of human research community. *arXiv preprint arXiv:2412.17767*, 2024a.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
- Tian Yu, Shaolei Zhang, and Yang Feng. Auto-rag: Autonomous retrieval-augmented generation for large language models. *arXiv preprint arXiv:2411.19443*, 2024b.
- Saber Zerhoubi and Michael Granitzer. Personarag: Enhancing retrieval-augmented generation systems with user-centric agents. *arXiv preprint arXiv:2407.09394*, 2024.
- Liang Zhang, Katherine Jijo, Spurthi Setty, Eden Chung, Fatima Javid, Natan Vidra, and Tommy Clifford. Enhancing large language model performance to answer questions and extract information more accurately. *arXiv preprint arXiv:2402.01722*, 2024.
- Zhebin Zhang, Xinyu Zhang, Yuanhang Ren, Saijiang Shi, Meng Han, Yongkang Wu, Ruofei Lai, and Zhao Cao. Iag: Induction-augmented generation framework for answering reasoning questions. *arXiv preprint arXiv:2311.18397*, 2023.
- Yuxiang Zheng, Shichao Sun, Lin Qiu, Dongyu Ru, Cheng Jiayang, Xuefeng Li, Jifan Lin, Binjie Wang, Yun Luo, Renjie Pan, et al. Openresearcher: Unleashing ai for accelerated scientific research. *arXiv preprint arXiv:2408.06941*, 2024.
- Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei Liu. Deepresearcher: Scaling deep research via reinforcement learning in real-world environments. *arXiv preprint arXiv:2504.03160*, 2025.
- Zhipu AI. Autoglm rumination. <https://autoglm-research.zhipuai.cn/>, 2025. Accessed: 2025-04-06.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. *arXiv preprint arXiv:2307.13854*, 2024.

He Zhu, Tianrui Qin, King Zhu, Heyuan Huang, Yeyi Guan, Jinxiang Xia, Yi Yao, Hanhao Li, Ningning Wang, Pai Liu, et al. Oagents: An empirical study of building effective agents. *arXiv preprint arXiv:2506.15741*, 2025a.

Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang. Deepreview: Improving llm-based paper review with human-like deep thinking process. *arXiv preprint arXiv:2503.08569*, 2025b.