

AI for Auto-Research: A Survey

Lingdong Kong*, Xian Sun*, Wei Chow*, Linfeng Li, Kevin Qinghong Lin, Xuan Billy Zhang, Song Wang, Rong Li, Qing Wu, Wei Gao, Yingshuo Wang, Shaoyuan Xie, Jiachen Liu, Leigang Qu, Shijie Li, Lai Xing Ng, Benoit R. Cottureau, Ziwei Liu, Tat-Seng Chua, Wei Tsang Ooi

Awesome AI Auto-Research Team

AI-assisted research is crossing a threshold: fully automated systems can now generate research papers for as little as \$15, while long-horizon agents can execute experiments, draft manuscripts, and simulate critique with minimal human input. Yet this productivity frontier exposes a deeper integrity problem: under scientific pressure, even frontier LLMs still fabricate results, miss hidden errors, and fail to judge novelty reliably. Surveying developments through April 2026, we present an end-to-end analysis of AI across the *complete* research lifecycle, organized into four epistemological phases: ¹**Creation** (idea generation, literature review, coding & experiments, tables & figures), ²**Writing** (paper writing), ³**Validation** (peer review, rebuttal & revision), and ⁴**Dissemination** (posters, slides, videos, social media, project pages, and interactive agents). We identify a sharp, stage-dependent boundary between reliable assistance and unreliable autonomy: AI excels at structured, retrieval-grounded, and tool-mediated tasks, but remains fragile for genuinely novel ideas, research-level experiments, and scientific judgment. Generated ideas often degrade after implementation, research code lags far behind pattern-matching benchmarks, and end-to-end autonomous systems have not yet consistently reached major-venue acceptance standards. We further show that greater automation can obscure rather than eliminate failure modes, making human-governed collaboration the most credible deployment paradigm. Finally, we provide a structured taxonomy, benchmark suite, and tool inventory, cross-stage design principles, and a practitioner-oriented playbook, with resources maintained at our project page.

 **Project Page:** <https://worldbench.github.io/awesome-ai-auto-research>

 **GitHub Repo:** <https://github.com/worldbench/awesome-ai-auto-research>




Figure 1 AI auto-research across the complete lifecycle. We organize AI assistance into four phases and eight stages: ¹**Creation** spans idea generation, literature review, coding & experiments, and tables & figures; ²**Writing** centers on paper writing; ³**Validation** includes peer review and rebuttal & revision; and ⁴**Dissemination** transforms papers into posters, slides, videos, social media, project pages, and interactive paper agents.

Contents

1	Introduction	4
2	Preliminaries	5
2.1	Research Lifecycle	5
2.2	Methodological Families	6
2.3	Scope & Literature Collection	7
2.4	Development Timeline	8
3	Phase 1: Creation	10
3.1	Idea Generation	10
3.1.1	LLM Internal Knowledge-Based Generation	10
3.1.2	External Signal-Driven Generation	11
3.1.3	Multi-Agent Collaborative Generation	11
3.1.4	Assessment: Novelty and Feasibility	11
3.1.5	Findings and Observations	12
3.2	Literature Review	12
3.2.1	Literature Retrieval	12
3.2.2	Survey and Related Work Generation	13
3.2.3	Deep Research Agents	13
3.2.4	Assessment: Retrieval and Synthesis Quality	14
3.2.5	Findings and Observations	14
3.3	Coding and Experiments	14
3.3.1	Code Generation	15
3.3.2	Paper-to-Code	15
3.3.3	Experiment Execution and Orchestration	15
3.3.4	Assessment: Code Correctness and Reproducibility	16
3.3.5	Findings and Observations	17
3.4	Tables and Figures	17
3.4.1	Scientific Figure Generation	17
3.4.2	Table Understanding and Generation	18
3.4.3	Mathematical Formulas and Algorithm Pseudocode	18
3.4.4	Assessment: Visual Fidelity and Scientific Accuracy	18
3.4.5	Findings and Observations	19
3.5	Summary and Transition: Creation	19
4	Phase 2: Writing	20
4.1	Paper Writing	20
4.1.1	Semi-Automated Writing Assistance	20
4.1.2	Fully Automated Paper Generation	21
4.1.3	Assessment: Writing Quality and AI Detection	21
4.1.4	Findings and Observations	22
4.2	Summary and Transition: Writing	22
5	Phase 3: Validation	23
5.1	Peer Review	23
5.1.1	Automated Review Generation	23
5.1.2	Meta-Review Generation	24
5.1.3	Reviewer Matching	24
5.1.4	Assessment: Review Consistency, Bias, and Robustness	24
5.1.5	Findings and Observations	25
5.2	Rebuttal and Revision	25
5.2.1	Reviewer Comment Analysis	25
5.2.2	Automated Rebuttal Generation	26
5.2.3	Assessment: Rebuttal Effectiveness	26
5.2.4	Findings and Observations	27
5.3	Summary and Transition: Validation	27

6	Phase 4: Dissemination	27
6.1	Research Dissemination (Paper2X)	28
6.1.1	Paper to Posters	28
6.1.2	Paper to Slides	28
6.1.3	Paper to Videos and Talks	29
6.1.4	Paper to Social Media	29
6.1.5	Paper to Agents and Tools	29
6.1.6	Assessment: Fidelity, Usability, and Adoption	30
6.1.7	Findings and Observations	31
6.2	Summary and Transition: Dissemination	31
7	Cross-Cutting Analysis	31
7.1	End-to-End Research Systems	32
7.1.1	Sequential and Pipeline-Based Systems	32
7.1.2	Search-Based and Self-Improving Systems	32
7.1.3	Skill-Based and Tool-Integrated Systems	32
7.1.4	Multi-Agent and Community-Scale Systems	33
7.1.5	Assessment: Lifecycle Coverage and Phase Boundaries	33
7.1.6	Findings and Observations	34
7.2	Evaluation Across the Research Lifecycle	34
7.2.1	Stage-Specific Benchmarks	34
7.2.2	Evaluation Methodologies	35
7.2.3	Emerging Evaluation Paradigms	36
7.2.4	Evaluation Gaps	36
7.3	Cross-Cutting Insights	36
7.3.1	Artifact Generation Outpaces Scientific Verification	37
7.3.2	Human-Governed Collaboration Remains the Most Reliable Deployment Mode	37
7.3.3	Capability Boundaries Emerge in Open-Ended Research Tasks	37
7.3.4	Effective Systems Converge on Layered Architectures	38
7.3.5	AI Use Has Become a Governance Problem, Not a Detection Problem	38
7.4	Open Challenges and Future Directions	38
7.4.1	Faithfulness Across Phase Boundaries	38
7.4.2	Scientific Judgment and Novelty Assessment	39
7.4.3	Verification, Reproducibility, and Accountability	39
7.4.4	Citation, Versioning, and Source Provenance	39
7.4.5	Governance, Disclosure, and Research Integrity	40
7.4.6	Cross-Domain Generalization and Infrastructure Access	40
7.4.7	Human Expertise and Cognitive Ownership	40
7.4.8	Toward Reliable AI-Assisted Research	40
8	Conclusion	41
A	Auto-Research Tool Inventory	42
A.1	Phase 1: Creation	42
A.2	Phase 2: Writing	46
A.3	Phase 3: Validation	47
A.4	Phase 4: Dissemination	49
A.5	Cross-Phase: End-to-End Systems	50
B	Survey Coverage Comparison & Taxonomy Analysis	51

1 Introduction

AI-assisted research is crossing a threshold. Large language models (LLMs) and their agentic extensions are no longer limited to local writing or coding support; they are beginning to operate across the research lifecycle itself. Recent systems illustrate the scale of this shift: The AI Scientist generated complete research papers at roughly \$15 per paper [122]; FARS ran continuously for 228 hours, consumed 11.4 billion tokens, and produced 100 papers, averaging one every 2.3 hours [14]; and ARIS reports an overnight workflow that ran 20+ GPU experiments, pruned unsupported claims, and improved a draft score from 5.0 to 7.5 through iterative review and revision [233]. These systems suggest a new paradigm: AI is moving from assisting individual research tasks to orchestrating multi-stage workflows that generate ideas, search literature, execute experiments, draft manuscripts, simulate critique, and prepare dissemination materials.

This rapid progress also exposes the defining tension of the field. AI systems are increasingly capable of producing research-like artifacts, yet remain far less reliable at verifying whether those artifacts are novel, faithful, executable, and scientifically meaningful. Generated ideas can appear promising but weaken after implementation [184]; generated code can run while implementing the wrong algorithm [71]; fluent manuscripts can conceal unsupported claims; automated reviews can be coherent yet lenient or vulnerable to manipulation [267]; rebuttals can promise revisions that are not later fulfilled [21]; and dissemination materials can simplify results beyond the evidence. The core challenge is therefore no longer whether AI can produce the *forms* of research, but whether it can preserve the *substance* of research: evidence, judgment, provenance, and accountability.

A lifecycle view is essential for understanding this challenge. Research is not a collection of independent tasks: ideas become experiments, experiments become claims, claims become manuscripts, reviews become revisions, and papers become public-facing summaries. Errors introduced early can be amplified downstream, especially when AI systems generate plausible outputs without preserving evidence or provenance. Despite the rapid emergence of research agents, writing assistants, scientific coding tools, automated reviewers, rebuttal systems, and Paper2X applications, the field still lacks a unified analysis of AI auto-research across the complete academic lifecycle. Without such a view, it is difficult to determine where AI reliably helps, where it fails systematically, and which deployment modes are scientifically credible.

Surveying developments through April 2026, we present the first end-to-end analysis of AI auto-research across the complete academic research lifecycle. We organize the field into four epistemological phases and eight stages: ¹**Creation**, covering idea generation, literature review, coding & experiments, and tables & figures; ²**Writing**, covering paper writing; ³**Validation**, covering peer review and rebuttal & revision; and ⁴**Dissemination**, covering posters, slides, videos, social media, project pages, and interactive paper agents. This structure follows the temporal sequence of research while making explicit the distinct AI capabilities, risks, and verification requirements introduced by each phase.

Our analysis yields five central findings. First, AI capability is strongest when tasks are structured, grounded, and externally checkable, but drops sharply for open-ended research tasks requiring novelty, implicit domain knowledge, long-horizon reasoning, or scientific judgment. Second, artifact generation consistently outpaces verification: across stages, AI can often produce plausible outputs faster than it can prove that they are correct, faithful, or meaningful. Third, the most reliable deployment mode is human-governed collaboration rather than full autonomy: AI can reduce mechanical friction in retrieval, drafting, coding, visualization, review support, and dissemination, but researchers must retain responsibility for judgment, interpretation, experimental design, argumentation, and accountability. Fourth, effective systems increasingly rely on layered architectures that combine exploration, tool-based execution, and verification, suggesting that orchestration, provenance, and feedback design are as important as model scale. Fifth, AI use in research is becoming a governance problem rather than a detection problem: as AI assistance becomes routine, the key questions are disclosure, attribution, responsibility, and whether scientific integrity is preserved.

This work makes three contributions to the emerging field of AI auto-research:

- We provide a unified taxonomy of AI auto-research across four phases and eight stages, covering both mature areas such as writing and coding, and underexplored areas such as rebuttal, scientific visualization, and research dissemination.

- We synthesize tools, benchmarks, and methodological families across the lifecycle, showing how systems have evolved from prompt-based assistance to retrieval-augmented, agentic, fine-tuned, and hybrid workflows.
- We identify cross-cutting capability boundaries and open challenges, including phase-boundary faithfulness, scientific judgment, reproducibility, citation provenance, governance, cross-domain generalization, and cognitive ownership.

The remainder of this paper is organized as follows. Section 2 introduces the lifecycle framework, methodological families, literature-collection scope, and development timeline. Section 3 to Section 6 build the roadmap of the four phases for AI-assisted research in temporal order. Section 7 synthesizes end-to-end systems, evaluation paradigms, cross-cutting insights, and open challenges. Section 8 concludes the paper.

2 Preliminaries

As AI-assisted research tools expand from isolated single stages (such as writing or coding aids) into multi-stage assistants, the field has become increasingly difficult to compare using a single vocabulary. Existing systems differ not only in their technical designs, but also in the research stages they target, the degree of autonomy they assume, and the forms of scientific risk they introduce.

To support a unified analysis, we first establish **four foundational elements**: (i) the high-level academic research lifecycle framework that organizes this survey (Section 2.1), (ii) the methodological families that recur across each stage (Section 2.2), (iii) the scope and methodology of our literature collection (Section 2.3), and (iv) a brief timeline of key developments (Section 2.4).

2.1 Research Lifecycle

We define the research lifecycle as **eight interconnected stages**, organized into **four phases**. Each phase groups stages that serve a shared function in the production, validation, and communication of scientific knowledge.

Phase 1: Creation. This phase covers the stages through which a research contribution is materially produced, including hypothesis formation, evidence gathering, experimentation, and scientific visualization.



S1 Idea Generation

Generating, refining, and evaluating research hypotheses. Techniques include direct LLM prompting, retrieval-augmented generation, knowledge-graph reasoning, and multi-agent collaboration for structured hypothesis formation.



S2 Literature Review

Retrieving, synthesizing, and organizing prior work into coherent research contexts. Modern systems span semantic retrieval, citation-graph traversal, survey generation, and deep research agents that iteratively explore the literature.



S3 Coding & Experiments

Translating ideas into executable code, running experiments, and analyzing empirical results. This stage includes code generation, paper-to-code translation, autonomous experiment orchestration, and result interpretation.



S4 Tables & Figures

Constructing method diagrams, result plots, comparison tables, mathematical formulas, and algorithmic illustrations. These artifacts transform raw outputs and conceptual designs into structured scientific representations.

Phase 2: Writing. This phase organizes the outputs of *Creation* into a formal scholarly manuscript for communication and external scrutiny.



S5 Paper Writing

Drafting, editing, polishing, and structuring academic manuscripts. AI assistance ranges from grammar correction and citation support to section-level drafting and full-paper generation.

Phase 3: Validation. This phase covers the stages through which the research community scrutinizes, critiques, and iteratively refines a manuscript.



S6 Peer Review

Generating structured reviews, matching reviewers to manuscripts, assessing review quality, and supporting meta-review decisions. These systems aim to assist, rather than replace, the community’s evaluative process.



S7 Rebuttal & Revision

Analyzing reviewer comments, identifying required evidence, drafting responses, and supporting manuscript revision. This stage connects external critique with additional analysis, clarification, and experimental follow-up.

Phase 4: Dissemination. This phase converts the manuscript and its supporting materials into formats accessible to broader research and public audiences.



S8 Paper2X

Converting papers into posters, slides, videos, project pages, demos, and social media content. Each output format targets a different audience and requires distinct design choices, fidelity constraints, and communication strategies.

Although presented in temporal order, the lifecycle is not strictly linear. Reviewer critiques in Phase 3 (*Validation*) may require returning to Phase 1 (*Creation*) for additional experiments, while dissemination outputs in Phase 4 (*Dissemination*) may expose ambiguities or errors that trigger revisions in Phase 2 (*Writing*). These feedback loops are central to research practice and are especially important for AI-assisted workflows, where errors can propagate across stages if not explicitly checked.

This four-phase grouping reflects the functional structure of research. Evidence and artifacts are produced in **P1** *Creation*, organized into a manuscript in **P2** *Writing*, externally scrutinized in **P3** *Validation*, and communicated to broader audiences in **P4** *Dissemination*.

We separate *Writing* from *Creation* because manuscript construction is not merely a formatting step: it is a rhetorical and evidential organization process that requires different AI capabilities from those used to produce code, experiments, or figures. We group Peer Review and Rebuttal under *Validation* because together they form the community-facing mechanism through which claims are challenged, defended, and revised. Finally, we treat *Dissemination* as a full phase because posters, slides, videos, project pages, and social media summaries are increasingly important knowledge artifacts with their own fidelity and trust requirements.

2.2 Methodological Families

Across the research lifecycle, AI-assisted research systems reuse a small set of methodological patterns. We group them into five broad families: ¹*prompt engineering*, ²*retrieval-augmented generation (RAG)*, ³*training-free agentic methods*, ⁴*training-based methods*, and ⁵*hybrid approaches*. These families are not mutually exclusive or strictly chronological; rather, they describe how current systems elicit, ground, specialize, and orchestrate LLM behavior. Many practical systems combine several of them, for example using prompts for decomposition, RAG for grounding, tools for execution, and trained modules for scoring or ranking.

Table 1 Dominant methodological families, representative systems, and research maturity across each stage of the four-phase research lifecycle. Notations: ✓ = “primary approach”, ◦ = “secondary/emerging”, — = “not used”.

Stage	Prompt Eng.	RAG	Agentic	Training	Hybrid	Representative works	Maturity
Phase 1: Creation							
S1: Idea Generation	✓	◦	✓	—	◦	AI Scientist [122], VirSci [193], Spark [168]	★★★★★
S2: Literature Review	◦	✓	✓	—	✓	PaperQA2 [189], AutoSurvey [214], STORM [178]	★★★★★
S3: Coding & Exp.	◦	◦	✓	—	✓	AIDE [81], PaperCoder [174], R&D-Agent [234]	★★★★★
S4: Tables & Fig.	✓	—	✓	◦	—	MatPlotAgent [236], AutoFigure [114], DeTikZify [13]	★★★★★
Phase 2: Writing							
S5: Paper Writing	✓	✓	◦	✓	◦	CycleResearcher [220], ScholarCopilot [215], XtraGPT [25]	★★★★★
Phase 3: Validation							
S6: Peer Review	◦	◦	✓	✓	✓	DeepReviewer [269], MARG [35], ReviewAgents [43]	★★★★★
S7: Rebuttal	◦	✓	✓	—	◦	RebuttalAgent [63], Paper2Rebuttal [129]	★★★★★
Phase 4: Dissemination							
S8: Dissemination	✓	◦	✓	—	—	Paper2Poster [146], PPTAgent [262], SlideGen [111]	★★★★★

Prompt engineering provides the simplest interface for adapting general-purpose LLMs to research tasks [217, 239]. It includes direct prompting, chain-of-thought reasoning, role assignment, structured templates, rubric-based instructions, and output constraints. Because it requires no additional training, it remains widely used for lightweight tasks such as brainstorming, editing, review drafting, rebuttal outlining, and social media generation, but it is sensitive to prompt wording and usually lacks persistent grounding.

Retrieval-augmented generation (RAG) grounds model outputs in external sources, including paper corpora, citation graphs, code repositories, benchmark records, and experimental logs [98]. It is especially important for literature review, citation support, evidence checking, rebuttal generation, and stages where source attribution is required. RAG reduces hallucination by exposing models to evidence at inference time, but does not ensure that selected sources are correct, version-consistent, or faithfully represented.

Training-free agentic methods extend LLMs with planning, tool use, memory, self-reflection, and iterative execution, enabling multi-step workflows without updating model parameters [169, 182, 239]. These methods are central to deep literature exploration, code debugging, experiment orchestration, review-response planning, and Paper2X workflows. Their strength lies in orchestration, while their main risk is error propagation when retrieval, tool use, or self-critique fails.

Training-based methods specialize models for stage-specific distributions, such as peer reviews, scientific manuscripts, code repositories, citation contexts, rebuttal traces, or benchmark trajectories [144, 213]. They include supervised fine-tuning, instruction tuning, preference optimization, reinforcement learning, and domain-specific adaptation. They can improve consistency, format adherence, domain vocabulary, and task-specific judgment, but depend heavily on data quality and may overfit to narrow benchmark or venue distributions.

Hybrid approaches combine multiple families into integrated research systems, for example by coupling RAG with agentic planning, fine-tuning domain-specific submodules, or embedding prompt-based controllers inside larger workflows [9, 94, 122, 178]. Hybrid systems are increasingly dominant because research workflows require generation and grounding, autonomy and verification, and flexible reasoning with stage-specific specialization.

Table 1 maps these methodological families to the eight lifecycle stages, using primary and secondary markers to indicate common design patterns in recent systems.

2.3 Scope & Literature Collection

This survey focuses on AI tools, methods, and benchmarks that support *human-driven academic research*, with an emphasis on computer science and machine learning. We cover work published or publicly released between 2023 and early 2026, while also referencing earlier foundational methods when they define recurring technical paradigms. Cross-disciplinary systems are included when they demonstrate capabilities relevant to the research

lifecycle, such as autonomous experimentation, literature synthesis, scientific coding, or evidence-grounded writing. We exclude general-purpose LLM capabilities that are not explicitly connected to research workflows, as well as closed systems for which insufficient technical or evaluative information is available.

To construct the survey corpus, we combined three complementary collection strategies:

- **Systematic keyword search** across Google Scholar, Semantic Scholar, arXiv, and DBLP, using queries related to AI-assisted research, automated research agents, literature review, scientific coding, paper writing, peer review, rebuttal generation, and research dissemination.
- **Snowball citation tracing** from representative seed papers in each lifecycle stage, including both backward tracing to foundational work and forward tracing to recent systems and benchmarks.
- **Community and repository monitoring**, including open-source projects, curated reading lists, and benchmark leaderboards that document emerging tools not yet covered by formal publications.

A paper, system, or benchmark was included only if it satisfied all three criteria: (i) it targets at least one stage of the research lifecycle defined in Section 2.1; (ii) it is publicly accessible through a publication, preprint, open-source repository, benchmark page, or technical report; and (iii) it provides sufficient methodological or evaluative detail to support critical analysis. When multiple versions of the same system exist, we prioritize the most recent or most technically complete version, while noting earlier versions when they mark important historical milestones.

The resulting corpus spans all four phases of the lifecycle, but the distribution is uneven. Most documented systems concentrate on **P1** (*Creation*), especially literature review, coding, and experiment automation, followed by **P2** (*Writing*), **P3** (*Validation*), and **P4** (*Dissemination*). This imbalance reflects both research maturity and publication availability: creation-stage tools are more frequently benchmarked and open-sourced, whereas dissemination-oriented tools are often commercial, workflow-specific, or evaluated through less standardized criteria. The benchmark landscape across stages is summarized in Table 2.

2.4 Development Timeline

The development of AI-assisted research can be understood as a shift from *stage-specific assistance* toward *multi-stage research automation*. Before 2024, most systems targeted isolated research tasks, such as literature search, scientific question answering, code generation, or domain-specific experiment planning. Early demonstrations, including Coscientist [15], showed that LLM-based agents could plan and execute scientific workflows in constrained laboratory settings, while domain foundation models such as AlphaFold 3 [1] illustrated the broader potential of AI systems to transform specialized scientific discovery.

In 2024, the field began moving from isolated tools toward end-to-end research agents. The AI Scientist [122] provided an early demonstration of an automated pipeline spanning idea generation, experiment execution, paper writing, and review-style evaluation. Around the same period, general coding agents, retrieval-augmented literature systems, and scientific reasoning benchmarks matured rapidly, making it possible to evaluate individual components of the research lifecycle more systematically. This transition marked an important change in emphasis: AI systems were no longer viewed only as assistants for local tasks, but increasingly as orchestrators of multi-step research workflows.

By 2025 and early 2026, the field entered a stage of rapid specialization and benchmarking. Dedicated systems emerged for nearly every lifecycle stage, including literature synthesis, paper-to-code translation, autonomous experiment orchestration, manuscript writing, peer review, rebuttal support, figure generation, and research dissemination. For example, OpenScholar [9] advanced retrieval-augmented scientific synthesis, AI Scientist v2 [229] explored stronger forms of end-to-end automated research, and FARS [14] demonstrated large-scale autonomous paper generation. At the same time, previously underexplored stages began receiving dedicated attention, including rebuttal writing (*e.g.*, RebuttalAgent [63]) and scientific visualization (*e.g.*, AutoFigure-Edit [114]).

These developments suggest that the field is no longer bottlenecked by model capability alone, but also by orchestration, evaluation, reliability, and governance across the full research lifecycle.

Table 2 Summary of datasets and benchmarks for AI-assisted research, organized by phases and stages.

#	Stage	Benchmark	Ref.	Year	GitHub	HF	Evaluation Focus	Scale	Link
Phase 1: Creation									
1	S1 : Idea Gen.	IdeaBench	[59]	2024	-	-	Novelty, feasibility	Multiple LLMs	↗
2	S1 : Idea Gen.	LiveIdeaBench	[162]	2024	-	-	Real-time model comparison	40+ models	↗
3	S1 : Idea Gen.	AI Idea Bench 2025	[154]	2025		-	Multi-dimensional assessment	3,495 papers	↗
4	S1 : Idea Gen.	ResearchBench	[120]	2025	-	-	Inspiration-based task decomp.	-	↗
5	S1 : Idea Gen.	Scientist-Bench	[199]	2025	-	-	Guided & open-ended AI research	Multi-domain	↗
6	S1 : Idea Gen.	HindSight	[78]	2026	-	-	Impact-based idea evaluation	-	↗
7	S1 : Idea Gen.	HeurekaBench	[147]	2026		-	Open-ended data-driven science	Multi-domain	↗
8	S2 : Lit. Rev.	LitSearch	[4]	2024			Literature retrieval	-	↗
9	S2 : Lit. Rev.	DeepScholar-Bench	[149]	2025		-	Research synthesis quality	-	↗
10	S2 : Lit. Rev.	ReportBench	[103]	2025		-	Deep research report quality	100 prompts	↗
11	S2 : Lit. Rev.	ScholarGym	[179]	2026	-	-	Information-gathering evaluation	2,536 queries	↗
12	S2 : Lit. Rev.	SciNetBench	[176]	2026	-	-	Relation-aware retrieval	18M papers	↗
13	S2 : Lit. Rev.	IDRBench	[40]	2026	-	-	Interactive deep research	100 tasks	↗
14	S3 : Coding	SWE-bench	[82]	2024			GitHub issue resolution	500 problems	↗
15	S3 : Coding	MLAgentBench	[73]	2024		-	ML experimentation	13 tasks	↗
16	S3 : Coding	LAB-Bench	[96]	2024			Biology research tasks	Multi-domain	↗
17	S3 : Coding	DiscoveryBench	[130]	2024			Data-driven discovery	-	↗
18	S3 : Coding	DiscoveryWorld	[76]	2024		-	Virtual discovery environment	120 tasks	↗
19	S3 : Coding	MLE-Bench	[20]	2024			Kaggle ML competitions	75 tasks	↗
20	S3 : Coding	ScienceAgentBench	[32]	2024			Scientific data analysis	-	↗
21	S3 : Coding	KernelBench	[143]	2025			GPU kernel generation	-	↗
22	S3 : Coding	TritonBench	[101]	2025		-	Triton operator generation	-	↗
23	S3 : Coding	ResearchCodeBench	[71]	2025	-	-	Novel ML code implementation	212 tasks	↗
24	S3 : Coding	SciReplicate-Bench	[224]	2025		-	Algorithm reproduction	100 tasks	↗
25	S3 : Coding	MLR-Bench	[24]	2025	-		Open-ended ML research	201 tasks	↗
26	S3 : Coding	MLGym	[135]	2025	-	-	AI research agent framework	-	↗
27	S3 : Coding	CURIE	[91]	2025		-	Rigorous experimentation	-	↗
28	S3 : Coding	PaperBench	[192]	2025			Paper replication	20 ICML papers	↗
29	S3 : Coding	AstaBench	[16]	2025			Scientific research suite	2,400+ problems	↗
30	S3 : Coding	ResearchClawBench	[225]	2025		-	Scientist-aligned workflows	Multi-domain	↗
31	S3 : Coding	EXP-Bench	[92]	2026			AI conducting experiments	461 tasks/51 papers	↗
32	S3 : Coding	FrontierScience	[208]	2026	-	-	Expert-level scientific tasks	Olympiad + PhD	↗
33	S3 : Coding	PostTrainBench	[158]	2026			LLM post-training automation	-	↗
34	S4 : Tab. & Fig.	MatPlotBench	[236]	2024	-	-	Data visualization	-	↗
35	S4 : Tab. & Fig.	PlotCraft	[253]	2025	-	-	Complex visualization	1K tasks	↗
36	S4 : Tab. & Fig.	TeXpert	[85]	2025	-	-	LaTeX code generation	3 difficulty levels	↗
37	S4 : Tab. & Fig.	PaperBananaBench	[268]	2026	-	-	Scientific illustration quality	292 test cases	↗
38	S4 : Tab. & Fig.	SciFlow-Bench	[256]	2026	-	-	Framework figure evaluation	500 figures	↗
39	S4 : Tab. & Fig.	Figure-Bench	[270]	2026			Text-to-illustration generation	3,300 pairs	↗
Phase 2: Writing									
40	S5 : Writing	ScholarCopilot	[215]	2025	-	-	Citation accuracy	40.1% top-1 acc.	↗
41	S5 : Writing	SciIG	[46]	2025	-	-	Introduction writing quality	NAACL/ICLR papers	↗
42	S5 : Writing	PaperWritingBench	[191]	2026	-	-	AI paper writing quality	200 papers	↗
Phase 3: Validation									
43	S6 : Peer Rev.	ClaimCheck	[142]	2025	-	-	Grounded LLM critiques	-	↗
44	S6 : Peer Rev.	Review-CoT	[43]	2025	-	-	Review reasoning chains	142K reviews	↗
45	S6 : Peer Rev.	AI Detection Bench	[244]	2025	-	-	AI review detection	788K reviews	↗
46	S7 : Rebuttal	ReviewMT	[197]	2024	-	-	Multi-turn review dialogue	26,841 papers	↗
47	S7 : Rebuttal	Re ²	[250]	2025	-	-	Full-stage review + rebuttal	19,926 papers	↗
48	S7 : Rebuttal	Commitment Checklist	[21]	2026	-	-	Unfulfilled rebuttal commitments	ICLR 2025	↗
Phase 4: Dissemination									
49	S8 : P2Slides	PPTEval	[262]	2025		-	Slide content, design, coherence	10K+ presentations	↗
50	S8 : P2Video	PresentQuiz	[271]	2025		-	Video faithfulness	101 paper-video pairs	↗
Cross-Phase									
51	Cross-Phase	RE-Bench	[221]	2024		-	Open-ended ML R&D	7 environments	↗
52	Cross-Phase	PaperBench	[192]	2025			End-to-end paper replication	20 ICML papers	↗

3 Phase 1: Creation

This phase covers the stages through which a research contribution is materially produced: generating an idea (**s1**), situating it within prior work (**s2**), producing empirical or analytical evidence (**s3**), and constructing visual representations of methods and results (**s4**). Together, these stages address two foundational questions: *what is the contribution, and what evidence supports it?*

Among the four phases, *Creation* currently has the richest tool ecosystem and broadest benchmark coverage, but its maturity remains uneven. **s1** (*Idea Generation*) has attracted extensive tooling, yet suffers from an ideation–execution gap in which seemingly novel ideas often weaken after implementation. **s2** (*Literature Review*) is rapidly improving through retrieval-augmented and agentic synthesis, but citation fidelity, coverage completeness, and multi-paper relational reasoning remain difficult. **s3** (*Coding and Experiments*) has progressed through code generation, paper-to-code translation, and autonomous experiment orchestration, but performance still drops sharply on genuinely novel research code. **s4** (*Tables and Figures*) remains comparatively underdeveloped despite its importance in daily research practice. We discuss these four stages in order below.

3.1 Idea Generation

Idea generation is the entry point of the research lifecycle, where candidate hypotheses, research questions, and experimental directions are proposed and refined. Existing approaches range from direct LLM prompting to externally grounded generation, multi-agent collaboration, and dedicated evaluation of novelty, feasibility, diversity, and downstream impact. Across these directions, the central challenge is that LLMs can produce ideas that appear novel and well-motivated, yet often struggle to generate ideas that remain feasible, distinctive, and impactful after execution.

A comprehensive inventory of ideation systems is provided in Table 3 (Appendix).

3.1.1 LLM Internal Knowledge-Based Generation

The simplest form of AI-assisted ideation prompts an LLM directly with a research domain, problem description, or literature context. Si *et al.* [183] established an influential baseline through a large-scale human study involving 100+ NLP researchers, finding that LLM-generated ideas were rated significantly higher in novelty than human ideas ($p < 0.05$). This result demonstrates the surface-level generative capacity of LLMs, but it also raises a central question for this stage: whether apparent novelty corresponds to executable and impactful research.

Subsequent work has explored three ways to strengthen direct generation. First, *iterative refinement* uses feedback loops to improve idea specificity and reduce shallow novelty. ResearchAgent [10] incorporates academic graph feedback to refine generated ideas, SciMON [209] iteratively compares candidate ideas against prior work to mitigate the tendency of direct LLM prompting toward shallow contributions, and Chain of Ideas [102] organizes literature into progressive reasoning chains that outperform simple prompting baselines.

Second, *learned quality signals* introduce explicit scoring or optimization objectives. Spark [168] combines retrieval-augmented generation with a judge model trained on 600K OpenReview reviews to estimate creativity, DeepInnovator [39] trains a 14B model under a “Next Idea Prediction” paradigm and reports 80–94% win rates against frontier models on ideation tasks, and Goel *et al.* [52] optimize AI co-scientist plans using rubric rewards extracted from existing papers, with RL-optimized plans preferred by human experts 70% of the time.

Third, *adaptive test-time compute* treats reasoning effort as a controllable resource. IRIS [47] uses MCTS in a human-in-the-loop ideation platform to allocate search as ideas converge, while FlowPIE [210] evolves scientific ideas at test time through flow-guided literature exploration.

A recent creativity-centered survey [175] further categorizes these methods into knowledge augmentation, prompt steering, inference-time scaling, multi-agent collaboration, and parameter adaptation.

3.1.2 External Signal-Driven Generation

Direct LLM generation is limited by the model’s parametric knowledge and by its tendency to produce plausible but weakly grounded ideas. External signal-driven methods address this limitation by anchoring ideation in structured knowledge, retrieved literature, or temporal research trends. Three signal sources are especially common, each grounding ideas from a different perspective: relational structure, textual evidence, and temporal opportunity.

Knowledge graphs provide relational structure for hypothesis formation. SciAgents [49] performs multi-agent reasoning over scientific knowledge graphs, while MOOSE-Chem [238] decomposes chemistry hypothesis generation into inspiration retrieval, hypothesis composition, and ranking, rediscovering hypotheses from 51 high-impact papers. MOOSE-Chem2 [237] extends this direction toward fine-grained, experimentally actionable hypotheses. *Paper retrieval* grounds ideas in unstructured literature. SciPIP [211] proposes ideas anchored to retrieved papers, and IdeaSynth [152] represents idea facets as nodes on an interactive canvas for literature-grounded refinement; in a user study with 20 participants, IdeaSynth encouraged users to explore more alternatives than LLM-only baselines. *Trend analysis* targets the temporal dimension of research opportunity. Nova [69] uses iterative planning and search to identify emerging research directions with improved diversity. Together, these methods suggest that external grounding is not merely an auxiliary feature, but a key mechanism for connecting generated ideas to the research frontier.

3.1.3 Multi-Agent Collaborative Generation

Multi-agent ideation systems attempt to improve idea quality by simulating aspects of research-community interaction, such as role specialization, critique, revision, and debate. VirSci [193] constructs a virtual scientific community in which multiple LLM agents participate in structured discussions, reporting higher novelty scores than a single-agent AI Scientist baseline (5.24 vs. 4.94). Its analysis suggests that agent diversity and discussion structure matter, with the best configuration using 8 members over 5 rounds with 50% diversity.

However, multi-agent scaling is not uniformly beneficial. A SIGDIAL 2025 study [206] finds that three critique–revision rounds are often sufficient, while additional rounds produce diminishing returns. Other systems explore richer collaboration mechanisms beyond discussion alone: Gu *et al.* [57] study *combinatorial creativity* by composing ideas across domains, and Deep Ideation [259] designs agents that navigate scientific concept networks through structured graph exploration.

Yet recent evidence also points to a deeper limitation: the “Artificial Hivemind” study [79] reports that LLM-generated ideas tend to cluster in narrow regions of the idea space, suggesting that diversity collapse may be a structural property of current models rather than a problem solved simply by adding more agents.

3.1.4 Assessment: Novelty and Feasibility


Evaluating generated ideas is difficult because strong research ideas must satisfy multiple criteria simultaneously: novelty, feasibility, clarity, significance, and eventual impact. Early benchmarks quantify parts of this space, but the central question is whether an idea remains valuable after it is implemented, tested, and situated against prior work.

IdeaBench [59] evaluates idea generation against 2,374 influential papers across eight research domains, while LiveIdeaBench [162] probes scientific creativity using 1,180 keyword prompts across 22 domains. Both suggest that scientific creativity is not well predicted by general-purpose benchmarks, with reasoning-focused models often performing better. ResearchBench [120] extends evaluation through inspiration-based task decomposition, and AI Idea Bench 2025 [154] scales assessment to 3,495 papers across two evaluation axes.

A recurring pattern across these benchmarks is the gap between apparent novelty and practical feasibility. IdeaBench reports that many LLMs score above 0.6 on novelty but below 0.5 on feasibility [59], indicating that generating plausible-sounding ideas remains easier than generating ideas that can be executed and validated. HindSight [78] sharpens this concern by introducing a time-split, impact-based evaluation, showing that LLM-as-Judge can overvalue novel-sounding ideas that do not later materialize into impactful work. This finding suggests that current evaluation protocols may reward apparent novelty rather than genuine research potential, reinforcing the need for execution-grounded and temporally robust assessment.


3.1.5 Findings and Observations

Stage 1: Idea Generation



State & Progress

Maturity Progression Grounding
Collaboration Training Benchmarks



Gaps & Limitations

Execution Feasibility Diversity
Mis-Evaluation Shallow Closed-Loop

- Idea generation is one of the most tool-rich stages in Phase 1 (*Creation*), with systems spanning prompting, retrieval, multi-agent collaboration, learned scoring, and test-time search.
- Clear capability progression: prompting → RAG → multi-agent → RL-trained, each generation addressing the weaknesses of its predecessor.
- External grounding is increasingly central: retrieval- and knowledge-graph-based methods better connect generated ideas to the research frontier than LLM-only prompting [211, 238].

- Ideas that score well before implementation can degrade substantially after execution ($\Delta = -1.98$ vs. -0.63 for human ideas [184]), exposing a gap between surface novelty and executable substance.
- Persistent novelty-feasibility tradeoff (> 0.6 vs. < 0.5 [59]) remains unresolved, and diversity collapse is structural, not solvable by scaling [79].
- LLM-as-Judge evaluation can reward apparent rather than genuine innovation, with reported novelty judgments negatively correlating with later real-world impact ($\rho = -0.29$ [78]).

3.2 Literature Review

Literature review anchors research in prior knowledge by retrieving relevant work, synthesizing evidence, and organizing existing findings into a coherent intellectual context. Compared with idea generation, this stage is more grounded and externally verifiable, making it one of the fastest-maturing areas in AI-assisted research. Existing systems have moved from semantic paper retrieval to citation-aware synthesis and long-horizon deep research agents. Yet two limitations remain central: systems can retrieve and summarize individual papers increasingly well, but still struggle with faithful citation, coverage completeness, and multi-paper relational reasoning.

A comprehensive inventory of literature review systems is provided in Table 4 (Appendix).

3.2.1 Literature Retrieval

Retrieval is the foundation of AI-assisted literature review: every downstream synthesis depends on whether the system can surface the right papers from scientific corpora that now contain tens of millions of entries. Existing methods can be grouped into three modes. *Semantic retrieval* forms the baseline, using dense representations and LLM-based query understanding to move beyond keyword matching. LitLLM [2] integrates LLMs with academic databases for dense retrieval, while PaperQA2 [189] extends this direction with citation verification and reports strong performance on scientific literature search.

Citation-graph-augmented retrieval adds structural signals beyond embeddings. Instead of treating papers as isolated documents, these methods use citation links, paper relations, and graph traversal to improve contextual coverage. OpenResearcher [265] combines RAG with graph traversal for accelerated literature exploration. *Agentic multi-step retrieval* further shifts retrieval from a one-shot ranking problem to an iterative search process. PaSa [62] deploys an LLM agent that issues follow-up queries and refines candidate sets, approximating how human researchers probe an unfamiliar topic. Alongside these methods, dedicated benchmarks have emerged to audit retrieval quality: LitSearch [4] targets retrieval precision, while CiteME [151] focuses on citation fidelity. Together, these efforts show that finding relevant papers is becoming easier, but ensuring that retrieved papers are used faithfully remains difficult.

3.2.2 Survey and Related Work Generation

Synthesis transforms retrieved papers into structured narratives. This marks a shift from retrieval-oriented systems, which optimize paper ranking and coverage, to generation-oriented systems, which must identify themes, compare methods, expose contradictions, and articulate research gaps. The subfield has developed through several increasingly structured designs.

Single-pass systems established the feasibility of automated survey drafting. AutoSurvey [214] demonstrated that LLMs can generate surveys of reasonable quality end-to-end, while SurveyX [110] improved content quality and approached human-expert performance on selected dimensions. *Structure-aware systems* then elevated outline planning from a formatting step to a core synthesis artifact. STORM [178] introduces multi-perspective question-asking to build comprehensive topic outlines, and SurveyForge [230] learns outline heuristics from human-written surveys together with memory-driven content generation, outperforming AutoSurvey on outline quality.

Multi-agent decomposition separates retrieval, verification, organization, and narrative writing into specialized subtasks. LiRA [51] and Agentic AutoSurvey [119] employ dedicated agents for different roles, while IterSurvey [251] treats outline generation as an iterative planning problem with stability checks. InteractiveSurvey [219] further introduces user customization, allowing researchers to refine reference categorization and outline structure through an interactive interface.

Citation- and editor-aware systems close the loop between synthesis and the writing environment. SurveyG [137] constructs a three-layer citation graph (Foundation/Development/Frontier) with hierarchical traversal, Citegeist [11] builds a dynamic RAG pipeline on the arXiv corpus, and CiteLLM [65] embeds hallucination-free reference discovery directly inside a LaTeX editor. Open-source systems such as GPT Researcher [38], PaperQA [128], and ChatPaper [124] further illustrate the growing practical adoption of literature synthesis tools beyond controlled research prototypes. However, citation fidelity remains a bottleneck: ScholarCopilot [215] reports only 40.1% top-1 citation accuracy, suggesting that generating plausible related-work text is still easier than grounding each claim in the correct source.

3.2.3 Deep Research Agents

Deep research agents differ from single-pass retrieval or survey-generation systems by treating literature exploration as an *iterative, agentic* process. Given an open-ended query, they plan sub-queries, retrieve and read sources, update their internal state, and continue until a report can be synthesized with sufficient confidence. This loop makes deep research agents closer to a workflow for long-horizon information seeking than a single retrieval model.

Commercial systems have popularized this paradigm for broad information synthesis. OpenAI Deep Research, Google Deep Research, Perplexity, and Elicit all support multi-source retrieval and report generation, though they differ in latency, citation style, interactivity, and target use cases. *Open-source literature-specific systems* adapt this paradigm to scientific research. OpenScholar [9], published in Nature, is a retrieval-augmented LM that searches large-scale open-access scientific corpora and outperforms PaperQA2 and Perplexity Pro on scientific literature benchmarks. Tongyi DeepResearch [7] from Alibaba is an agentic LLM specialized for long-horizon deep information seeking, achieving strong results on deep research benchmarks.

Training-era approaches target the data and optimization bottlenecks that limit long-horizon research agents. O-Researcher [240] combines multi-agent distillation with agentic reinforcement learning to improve benchmark performance, while OpenResearcher [108] addresses the trajectory-data bottleneck by constructing an offline trajectory synthesis pipeline over large document collections. These synthesized trajectories provide long-horizon tool-use supervision for training research agents. *Domain-focused variants* remain important for specialized synthesis tasks: CHIME [68] provides LLM-assisted hierarchical organization of scientific studies, and ASReview [207], published in Nature Machine Intelligence, uses active-learning-based screening to reduce manual effort in systematic reviews while maintaining recall. Collectively, deep research agents span a spectrum from lightweight factual lookup to long-horizon autonomous synthesis, but increasingly converge on the same iterative architecture: plan, retrieve, read, update, and synthesize.

3.2.4 Assessment: Retrieval and Synthesis Quality


Evaluation has shifted from *retrieval accuracy* alone (“did the system find the right papers?”) toward broader *synthesis quality* (“did it produce a useful, accurate, and well-organized review?”). At the output level, DeepScholar-Bench [149] establishes a dedicated benchmark for research synthesis across coverage, coherence, and factual accuracy. ReportBench [103] scales this direction to deep research reports derived from survey-style prompts.

At the process level, ScholarGym [179] isolates the information-gathering stage of deep research by decomposing it into query planning, tool invocation, and relevance assessment. This is an early step toward evaluating *how* a system reaches its answer, rather than judging only the final output. Benchmarks have also begun probing structural and interactive dimensions of literature competence. SciNetBench [176] introduces a relation-aware benchmark for literature retrieval agents over large-scale AI literature, revealing that relation-aware retrieval accuracy often remains low. IDR Bench [40] addresses the human-in-the-loop dimension through interactive deep research tasks with on-demand user interaction.

Across these efforts, four evaluation dimensions have crystallized: *citation accuracy*, whether references are correctly attributed and faithfully support the associated claims; *coverage completeness*, whether the review captures the relevant landscape without major omissions; *narrative coherence*, whether the synthesis has logical flow, thematic organization, and readability; and *factual grounding*, whether claims are supported by cited evidence rather than hallucinated. SurveyX [110] exemplifies this multi-dimensional view by evaluating content quality, structure quality, and citation accuracy as separate axes. The main open challenge is to develop automated metrics that correlate with expert judgment on synthesis quality while remaining robust across domains, venues, and writing styles.


3.2.5 Findings and Observations

🔍 Stage 2: Literature Review



✅ State & Progress

- Fastest
- Convergence
- Open-Source
- Retrieval
- Synthesis
- DeepResearch



⚠️ Gaps & Limitations

- Relations
- Hallucination
- Citation
- Cross-Domain
- Coherence
- Scalability

- Fastest-maturing stage: four generations in two years (single-pass → structure-aware → multi-agent → editor-aware), with 35 systems spanning retrieval, synthesis, and deep research.
- Commercial and open-source systems increasingly converge on an iterative architecture: plan → retrieve → read → update → synthesize.
- Recent evidence suggests that trajectory data and long-horizon tool-use supervision can be as important as model scale for improving the performance of deep research systems [108].

- Multi-paper relational reasoning remains a core bottleneck: the citation top-1 accuracy metric remains largely limited [215], and relation-aware retrieval often remains weak [176].
- Hallucination has shifted from obvious fabrication to subtle misgrounding: generated claims may appear well-cited while not being faithfully supported.
- Nearly all benchmarks and systems target ML/NLP literature; cross-domain synthesis (chemistry, biology, physics) remains largely untested and likely requires domain-specific retrieval infrastructure.

3.3 Coding and Experiments

This stage translates research ideas into executable implementations, runs experiments, and analyzes the resulting evidence. Compared with literature review, coding and experimentation require AI systems to interact with external environments: repositories, dependencies, datasets, compute resources, test suites, and evaluation scripts. Existing work spans general code generation, paper-to-code translation, experiment orchestration, and result analysis. Across these directions, the central challenge is not whether LLMs can

write plausible code, but whether they can produce semantically correct research implementations, execute meaningful experiments, and interpret results reliably.

A comprehensive inventory of coding and experiment systems is provided in Table 5 (Appendix).

3.3.1 Code Generation

General-purpose code generation has become one of the most mature capabilities of current LLMs. On SWE-bench Verified [195], which evaluates real-world GitHub issue resolution, frontier systems now exceed 76%. Agent frameworks have played a central role in this progress. SWE-agent [231] established the *agent-computer interface* paradigm, giving LLMs structured access to files, tests, and tool calls rather than relying on unstructured shell interaction. OpenHands [212] extends this direction into a general open platform for software engineering agents and has become a common backbone for coding-oriented workflows.

However, high performance on standard software benchmarks does not directly imply readiness for research coding. SWE-bench Verified has been questioned for potential contamination, and more challenging variants expose a sharper limitation: performance drops to 23% on SWE-bench Pro [37] and 25% on SWE-EVO [201]. These results suggest that standard benchmarks may overestimate robustness when tasks are familiar, well-scaffolded, or pattern-matchable. This distinction becomes more pronounced in research settings, where the target is not only to fix existing software but to implement underspecified algorithms, reproduce implicit design choices, and validate scientific claims.

3.3.2 Paper-to-Code

Paper-to-code translation is a research-specific form of code generation. It is harder than conventional software engineering because research papers often mix natural-language descriptions, equations, pseudocode, ablation details, and domain conventions, while leaving key implementation choices implicit. PaperCoder [174] addresses this setting with a three-stage multi-agent framework for planning, analysis, and code generation, transforming ML papers into executable repositories.

Dedicated benchmarks quantify how difficult this setting remains. ResearchCodeBench [71] evaluates LLMs on 212 novel ML implementation tasks, where the best reported model achieves only 37.3% accuracy; notably, 58.6% of errors are semantic, meaning that the generated code runs but implements the wrong algorithm or behavior. SciReplicate-Bench [224] reports a similar ceiling of 39% across 100 tasks from 36 NLP papers. SciCode [203] extends research-level coding evaluation to mathematics, physics, and chemistry, while PaperBench [192] decomposes 20 ICML 2024 papers into individually gradable subtasks covering environment setup, experiment execution, and result reproduction. Together, these benchmarks reveal a substantial gap between general software issue resolution and faithful research implementation.

At the high end, FunSearch [161] demonstrates that LLM-generated programs can contribute to genuine mathematical discovery when embedded inside an evolutionary search loop. This result is important, but it also clarifies the boundary of current capability: success comes not from raw one-shot code generation alone, but from coupling generation with aggressive search, evaluation, and selection. The resulting contrast between strong performance on familiar software benchmarks and much lower performance on novel research code defines the capability cliff of this stage.

3.3.3 Experiment Execution and Orchestration

Once the code is available, the next challenge is to run experiments systematically and efficiently. Experiment orchestration systems provide infrastructure for planning runs, modifying code, launching jobs, monitoring results, and iterating over failures. MLAgentBench [73] evaluates language agents on ML experimentation; MLR-Copilot [104] separates autonomous research into idea and experiment agents; DS-Agent [58] targets end-to-end data-science workflows; and AIDE [81] frames ML engineering as tree search in code space. Broader evaluation environments, including MLR-Bench [24], MLE-Bench [20], MLGym [135], and CURIE [91], provide increasingly standardized testbeds for measuring autonomous experimentation.

Recent systems push this infrastructure toward higher-throughput and closed-loop research workflows. R&D-Agent [234] uses a Researcher-Developer dual-agent design for ML experimentation, while Karpathy’s autore-

search [87] demonstrates high-throughput experiment iteration. Closed-loop systems such as CodeScientist [77], Dolphin [245], and NovelSeek [249] attempt to connect hypothesis generation, implementation, execution, and verification. EvoScientist [127] further illustrates the ambition of this direction by reporting accepted papers generated through a self-evolving research pipeline. These systems show that experimental throughput and workflow automation are improving rapidly, but their reliability still depends heavily on task scaffolding, benchmark design, and verification quality.

A complementary line of work couples execution with search and learning signals. AlphaEvolve [140] improves algorithms through LLM-generated mutations and automated evaluation. Si *et al.* [185] use execution-grounded search with large-scale parallel GPU experiments, outperforming GRPO baselines. SciNav [254] uses pairwise tree-search judgments to select promising branches, while Yuksekgonul *et al.* [246] combine test-time training and reinforcement learning for continuous improvement across mathematics, GPU kernel optimization, and computational biology. AutoReproduce [260] addresses a different but related problem: reproducing cited experiments by extracting implicit knowledge from paper lineages.

Domain-specific systems illustrate how orchestration changes when the environment is scientific rather than purely computational. In chemistry, Coscientist [15] and ChemCrow [17] use LLM-driven tools to support autonomous research workflows. In biology, AlphaFold 3 [1] extends protein structure prediction to biomolecular complexes, while CRISPR-GPT [156], BioPlanner [141], and LAB-Bench [96] target gene-editing design, protocol planning, and biology research evaluation. For systems-level optimization, KernelBench [143] and TritonBench [101] evaluate whether LLMs can generate efficient GPU kernels and Triton operators. Cross-domain suites such as AstaBench [16] and EXP-Bench [92] broaden evaluation to multi-domain scientific tasks and autonomous experiment execution.

Overall, the execution layer has advanced quickly, especially when tasks are well specified and feedback is automated. The harder problem is experiment *planning*: deciding which experiments are worth running, in what order, and how to interpret failures. Many current systems perform well on prescribed task pools, but remain less reliable when asked to choose genuinely novel research directions. In this sense, coding and experimentation expose the same broader pattern as idea generation: execution capability is improving faster than the scientific judgment needed to decide what should be executed.

3.3.4 Assessment: Code Correctness and Reproducibility


Assessing coding and experiment systems requires more than checking whether the generated code runs. Research code must implement the intended algorithm, reproduce reported results, support meaningful ablations, and generate evidence that can be interpreted correctly. This makes semantic correctness and reproducibility central evaluation criteria.

Several benchmarks expose the difficulty of this interpretive layer. DiscoveryBench [130] and ScienceAgent-Bench [32] evaluate scientific reasoning over experimental data, showing that LLMs still struggle with multi-step analysis over complex result sets. DiscoveryWorld [76] provides a virtual environment with 120 challenge tasks for automated scientific discovery agents. InfiAgent-DABench [70] benchmarks end-to-end data-analysis workflows, including data cleaning, statistical testing, and visualization generation across diverse domains.

The core bottleneck is moving from raw outputs to trustworthy claims. Current systems can often produce plots, summary statistics, and local interpretations, but they are less reliable at identifying statistically meaningful trends, diagnosing failure modes, designing decisive ablations, and synthesizing results into a coherent empirical argument. This limitation is particularly consequential because coding errors and experimental misinterpretations can propagate into later writing and review stages, where polished narratives may obscure weak or incorrect evidence.

3.3.5 Findings and Observations


Stage 3: Coding & Experiments



State & Progress

37% Ceiling Closed-Loop Search

Throughput Orchestration Cross-Domain



Gaps & Limitations

Semantic Planning Fabrication

Insight Gap Benchmark Leak Verification

- Sharpest capability boundary across all stages: 76% on pattern-matching *vs.* 37–39% on novel research code, consistently reproduced across 4+ independent benchmarks [71, 224].
- Execution infrastructure is no longer the bottleneck: systems sustain ~12 experiments/hour in closed-loop, with generated papers accepted at academic venues [87, 127].
- Coupling generation with search (evolutionary, tree, RL) consistently outperforms raw code generation [81, 140, 161], suggesting that search strategy matters more than model capability alone.

- Semantic failures are especially problematic: generated code may execute successfully while implementing the wrong algorithm or producing misleading results [71].
- Current systems execute prescribed tasks more reliably than they choose meaningful experiments; experiment planning remains strongly dependent on human scientific judgment.
- 80% of fully autonomous results are fabricated [24], and downstream review catches only half of methodological issues [125], creating a compounding verification deficit.

3.4 Tables and Figures

Tables and figures transform experimental outputs, statistical summaries, algorithms, and conceptual designs into publication-ready research artifacts. Existing systems cover scientific figure generation, data visualization, table construction, formula generation, and algorithmic illustration. Compared with coding and experimentation, this stage is less about producing new evidence than about representing evidence faithfully. Across these artifact types, the central challenge is the gap between *visual plausibility* and *scientific correctness*: AI-generated outputs may look professional while containing incorrect labels, misleading layouts, invalid numerical relationships, or domain-specific notation errors.

A comprehensive inventory of figure and table generation systems is provided in Table 6 (Appendix).

3.4.1 Scientific Figure Generation

Scientific figure generation spans method diagrams, architecture illustrations, result plots, data visualizations, and pipeline figures. Standard result plots are comparatively tractable because they can often be grounded in structured data and executable plotting code. In contrast, method diagrams and framework figures are harder because they require faithful spatial organization, correct information flow, domain-specific symbols, and paper-specific visual conventions.

For *method and architecture diagrams*, AutoFigure-Edit [114] generates editable text-to-SVG scientific illustrations from long-form text, enabling users to revise generated figures rather than treating them as fixed images. Its companion system AutoFigure [270] introduces FigureBench for generating and refining publication-oriented scientific illustrations. PaperBanana [268] employs multiple specialized agents for retrieval, planning, styling, visualization, and critique, while StarVector [160] focuses on scalable vector graphics from textual descriptions. Together, these systems show a shift from static image generation toward editable, structured, and critique-aware figure construction.

For *result plots and data visualization*, MatPlotAgent [236] uses VLM-based visual feedback to improve data visualization quality, while PlotGen [53] and PlotCraft [253] study chart generation across diverse plot types and task difficulties. CoDA [31] explores multi-agent collaboration for visualization, and ChartGPT [204] decomposes chart generation into sequential reasoning steps for handling abstract natural-language inputs.

More recent systems broaden the scope of generation and evaluation: SciFig [74] introduces rubric-based evaluation for pipeline figures, VisCoder [139] studies code-based visualization generation at scale, Diagram-Agent [218] targets multiple diagram categories with specialized agents, and SciFlow-Bench [256] evaluates scientific framework figures through structure-first analysis. These efforts indicate that standard data plots are increasingly tractable, while complex framework figures remain difficult because they require structural consistency rather than only visual appeal.

For *figure editing and optimization*, VIS-Shepherd [145] provides constructive feedback for LLM-based data visualization, emphasizing critique and revision rather than direct generation alone. [22] surveys publisher policies on AI-generated figures and proposes best-practice guidelines for responsible use. The SAIL framework [172] separates domain logic from code syntax, allowing researchers to retain scientific oversight while delegating implementation details to AI. Across these systems, the emerging design principle is human-guided refinement: AI can accelerate layout, rendering, styling, and accessibility improvements, but researchers must verify whether the figure faithfully represents the underlying method or data.

3.4.2 Table Understanding and Generation

Table generation spans two complementary tasks: understanding existing tables and creating new ones. On the understanding side, Chain-of-Table [216] performs table reasoning through multi-step table transformations, reflecting the fact that many table tasks require sequential operations rather than single-pass extraction. On the generation side, ArxivDIGESTables [136] synthesizes scientific literature into structured comparison tables, ShowTable [121] introduces collaborative reflection and refinement for creative table visualization, and Table2LaTeX-RL [115] converts table images into LaTeX code using reinforced multimodal language models.

Compared with standard figure generation, table generation remains less mature because scientific tables must satisfy stricter semantic constraints. Comparison tables require consistent axes, fair grouping of methods, complete citation coverage, and correct numerical transcription. Ablation tables are even more demanding because they encode experimental design choices, not only final results. AbGen [261] evaluates LLMs on ablation study design using expert-annotated examples from NLP papers, revealing a significant gap between LLM-generated table plans and human expert judgments. This suggests that table generation is not merely a formatting problem; it requires understanding which comparisons are scientifically meaningful and how evidence should be organized.

3.4.3 Mathematical Formulas and Algorithm Pseudocode

Mathematical formulas, TikZ diagrams, and algorithm pseudocode are compact representations of scientific reasoning, making them particularly sensitive to small errors. Unlike ordinary prose or standard charts, these artifacts require exact syntax and exact semantics simultaneously: a misplaced symbol, index, operator, arrow, or dependency can change the meaning of the method. As a result, formula and pseudocode generation remain less robust than natural-language polishing or standard visualization.

Recent systems address this challenge through specialized datasets, multimodal inputs, and iterative refinement. AutomataTikZ [12] introduces DaTikZ, a large-scale TikZ dataset, and shows that fine-tuned models can outperform general-purpose LLMs on scientific vector graphics. DeTikZify [13] extends this line with multimodal input and MCTS-based iterative refinement over a larger collection of TikZ graphics. TikZilla [56] further suggests that domain-specific training with supervised fine-tuning and reinforcement learning can make smaller open-source models competitive with larger general-purpose models on TikZ generation. TeXpert [85] highlights the remaining difficulty: accuracy drops sharply as LaTeX tasks become more complex, especially for tables with merged cells, nested environments, and nontrivial formatting constraints. These results reinforce the broader pattern of table and figure generation: specialized training and iterative refinement help, but human verification remains necessary when visual or symbolic artifacts carry scientific meaning.

3.4.4 Assessment: Visual Fidelity and Scientific Accuracy

Evaluation for table and figure generation must assess both *visual fidelity* and *scientific accuracy*. Visual fidelity asks whether an artifact is readable, aesthetically coherent, and consistent with publication conventions. Scientific accuracy asks whether the artifact faithfully represents the underlying data, method, comparison,

or mathematical relation. The distinction is crucial: an AI-generated figure may look professional while containing misaligned arrows, incorrect labels, invalid quantitative relationships, or domain-specific notation errors.

Recent benchmarks increasingly target this gap. SciFlow-Bench [256] uses inverse-parsing evaluation to detect structurally incorrect but visually plausible framework figures. FigureBench [270] evaluates scientific illustration generation and refinement. PlotCraft [253] studies chart generation across diverse chart types, while SciFig [74] provides a rubric-based evaluation for pipeline figures. TeXpert [85] evaluates LaTeX generation across difficulty levels, exposing steep performance degradation on hard cases. AbGen [261] extends evaluation to ablation study design, where the challenge is not only formatting a table but selecting scientifically meaningful comparisons.

Across artifact types, maturity remains uneven. Standard result plots are the most tractable because they can be generated from structured data and validated through executable plotting code. Method diagrams and framework figures remain harder because they require spatial organization and semantic consistency. Tables are difficult when they encode comparison logic or ablation design rather than simple formatting. Mathematical formulas, TikZ diagrams, and pseudocode exhibit steep accuracy cliffs because small syntactic errors can alter scientific meaning. Together, these benchmarks show that **S4** evaluation is moving from appearance-based judgment toward structure-, semantics-, and task-aware assessment.

3.4.5 Findings and Observations

💡 Stage 4: Tables & Figures

<div style="display: flex; align-items: center; margin-bottom: 10px;"> <div> <p style="margin: 0;">✔ State & Progress</p> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Emerging Visualization Small Models </div> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Multi-Agent Benchmarks Editing </div> </div> </div> <ul style="list-style-type: none"> Fastest-growing stage from zero: first dedicated tools appeared in late 2025, yet 20+ systems already span figures, tables, formulas, and editing. Standard data visualization becomes increasingly tractable: 90%+ execution pass rate on Matplotlib/Seaborn [139], with multi-agent approaches boosting quality > 40% over baselines [31]. Domain-specific training and iterative refinement can make smaller specialized models competitive for structured visual languages such as TikZ [13, 56]. 	<div style="display: flex; align-items: center; margin-bottom: 10px;"> <div> <p style="margin: 0;">⚠ Gaps & Limitations</p> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Correctness Uneven Tables </div> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Formulas Spatial Symbols </div> </div> </div> <ul style="list-style-type: none"> Visual plausibility \neq scientific correctness: generated artifacts may look polished while misrepresenting data, structure, notation, or information flow. Maturity is sharply uneven: figures are most advanced, table generation lags with no high-traction LaTeX tool, and formula accuracy drops from 78.8% to 15% with complexity [85]. Tools remain assistants, not producers: AI-generated figures frequently require human modification for domain-specific symbols, spatial relationships, and paper-specific visual languages.
---	--

3.5 Summary and Transition: Creation

The four Creation stages are tightly coupled in practice. **S1** (*Idea Generation*) generates candidate hypotheses, **S2** (*Literature Review*) situates them within prior work, **S3** (*Coding and Experiments*) turns them into executable implementations and empirical evidence, and **S4** (*Tables and Figures*) converts the resulting outputs into visual and structured artifacts for communication. Progress across these stages shows a consistent pattern: AI systems are increasingly effective at producing the artifacts of research, including ideas, literature summaries, code, experiments, figures, and tables, but they remain less reliable at verifying whether these artifacts are novel, faithful, executable, and scientifically meaningful.

This gap appears differently at each stage. In **S1**, plausible novelty often weakens after implementation. In **S2**, fluent synthesis can conceal citation errors or incomplete coverage. In **S3**, executable code can still be

semantically wrong, and automated runs do not guarantee meaningful experimental design. In [S4](#), polished visual artifacts may misrepresent data, notation, or methodological structure. These failure modes suggest that Creation-stage automation is most credible when coupled with grounding, execution feedback, explicit verification, and human scientific judgment.

The outputs of Phase 1 (*Creation*) constitute the raw material for Phase 2 (*Writing*). Ideas, retrieved literature, validated experiments, statistical summaries, comparison tables, and scientific figures must be organized into a coherent manuscript that explains the contribution, justifies its significance, and prepares it for external scrutiny. We therefore next turn to *Writing*, where AI assistance shifts from producing research artifacts to structuring evidence into a scholarly argument.

4 Phase 2: Writing

This phase consists of a single stage: *Paper Writing* ([S5](#)). Writing merits its own phase because it transforms the artifacts produced in Phase 1 into a scholarly argument. It is not merely a formatting step: a manuscript must select evidence, structure claims, situate contributions in the literature, explain methods with sufficient detail for reproducibility, and anticipate objections before external scrutiny in Phase 3 (*Validation*). Compared with Phase 1 (*Creation*), which emphasizes artifact production, Phase 2 (*Writing*) emphasizes rhetorical organization and evidential justification.

This distinction matters for AI-assisted research. Writing tools are among the most widely adopted systems in the AI-for-research ecosystem, spanning grammar correction, sentence polishing, section drafting, citation support, and full-paper generation. At the same time, Writing is one of the most ethically sensitive phases because questions of authorship, attribution, disclosure, and the boundary between assistance and generation remain unresolved. The central challenge is therefore not whether AI can produce fluent academic prose, but whether it can preserve factual grounding, argumentative depth, citation fidelity, and human accountability.

4.1 Paper Writing

AI-assisted writing has moved from occasional support to mainstream research practice. Large-scale corpus analyses estimate detectable AI modification in up to 17.5% of computer science abstracts [\[109\]](#) and 13.5% of biomedical abstracts [\[90\]](#), while self-reported adoption is higher: a 2025 Nature survey found that more than half of researchers report seeking AI writing help [\[134\]](#). These measurements are imperfect, but together they indicate a clear shift: AI writing assistance is now embedded in everyday scientific workflows. This makes the quality, transparency, and governance of AI-assisted writing increasingly important.

A comprehensive inventory of AI-assisted writing systems is provided in [Table 7](#) (Appendix).

4.1.1 Semi-Automated Writing Assistance

Semi-automated writing assistance supports different parts of the manuscript workflow, from planning and drafting to polishing and revision. At the planning stage, systems help generate titles, outlines, section structures, and citation suggestions. ScholarCopilot [\[215\]](#), for example, trains LLMs for academic writing with integrated citation recommendation, reflecting a broader trend toward tools that combine text generation with literature grounding.

During drafting, commercial tools such as Grammarly, Writefull, Paperpal, and GPT-based editors support paragraph generation, sentence polishing, citation insertion, and style refinement. Open-source prompt templates [\[106\]](#) provide lightweight alternatives, while CoAuthor [\[97\]](#) studies human–AI collaborative writing workflows. The dominant paradigm is increasingly shifting from “*AI writes for you*” to “*AI writes with you*”: AI handles mechanical or local operations, such as polishing, citation formatting, and initial drafting, while researchers retain responsibility for novelty, argumentation, experimental interpretation, and scientific judgment.

Editor-integrated systems make this collaboration more explicit. PaperDebugger [\[67\]](#) embeds a multi-agent system into Overleaf, running Reviewer, Enhancer, Scoring, and Researcher agents within the writing environment. A complementary line of work emphasizes cognitive engagement and transparency. Script&Shift [\[186\]](#)

structures AI-assisted writing around *source transformation* rather than direct text generation, aiming to preserve the writer’s active reasoning. DraftMarks [188] provides visual traces of revision intensity and AI-generated content, making the human–AI writing process more transparent to readers and reviewers. Empirical evidence [187] further suggests that purposeful AI support can assist student writing without fully displacing cognitive effort.

Post-writing systems focus on revision, consistency, and style. XtraGPT [25] provides an open-source LLM suite for instruction-guided scientific paper revision, SciIG [46] benchmarks introduction writing using recent NAACL and ICLR papers, OpenDraft [36] uses specialized agents to generate long research drafts with citation support, and LimAgents [6] integrates OpenReview comments and citation networks to generate research-limitation statements. Together, these systems show that semi-automated writing assistance is most credible when it augments researcher control rather than replaces the intellectual work of framing, interpreting, and defending a contribution.

4.1.2 Fully Automated Paper Generation

Fully automated paper generation attempts to move beyond local assistance toward end-to-end manuscript production. Existing systems can be grouped into three directions. First, end-to-end research systems such as The AI Scientist [122, 123] and Agent Laboratory [171] generate full papers as part of broader automated research pipelines. These systems demonstrate the feasibility of producing complete paper-like artifacts, but their outputs often remain limited by shallow argumentation, weak experimental validation, or insufficient novelty.

Second, benchmarked paper-generation systems aim to approach human review standards. CycleReviewer [220] reports generated papers scoring 5.36 on the ICLR scale, approaching but still below the reported accepted-paper average of 5.69. This gap is important because it suggests that the main bottleneck is no longer surface fluency alone. Rather, near-threshold papers often lack the argumentative depth, experimental rigor, and reviewer anticipation that distinguish publishable work from plausible drafts.

Third, rubric-guided and section-specific systems improve parts of the manuscript rather than generating the entire paper from scratch. APRES [258] discovers rubrics predictive of citation counts and revises papers accordingly, with human experts preferring revised papers 79% of the time. FutureGen [5] targets “Future Work” section generation. PaperWritingBench [191], introduced as part of the PaperOrchestra framework, provides a dedicated benchmark for automated paper writing by evaluating multi-agent systems against reverse-engineered top-tier conference papers. These systems indicate that automated writing is increasingly measurable, but also reinforce that high-quality papers require more than fluent text: they require evidence-grounded reasoning and coherent scientific contribution.

4.1.3 Assessment: Writing Quality and AI Detection

Assessment of AI-assisted writing involves two related but distinct questions: whether AI use can be detected, and whether the resulting manuscript is scientifically strong. Detection remains unreliable as a governance mechanism. Current detectors can produce unacceptable false positives, especially for formal, non-native, or highly edited academic prose, motivating a shift at major venues from attempting to *detect* AI use toward requiring authors to *declare* AI use. Watermarking offers a more principled route under controlled settings [159], but it requires model-provider cooperation and remains vulnerable to paraphrasing, translation, and post-editing.

Quality evaluation is more important, but also harder. Good academic writing must be assessed along multiple dimensions: factual correctness, citation accuracy, argumentative coherence, methodological completeness, novelty of framing, and stylistic appropriateness. LLM-as-Judge frameworks are increasingly used to approximate parts of this evaluation. CycleReviewer [220] reports a 26.89% reduction in Proxy MAE relative to individual human reviewers for score prediction, while the Stanford Agentic Reviewer [80] achieves review-score correlations comparable to human inter-rater agreement ($\rho = 0.42$ vs. human $\rho = 0.41$). These results suggest that automated evaluators can provide useful review-style signals, but they should not be treated as substitutes for expert assessment: score prediction and agreement metrics only partially capture factual grounding, evidential rigor, novelty, and scientific contribution.

The central failure mode of AI writing is therefore not ungrammatical prose, but unsupported persuasion: text that is fluent, well-structured, and citation-like, yet insufficiently grounded in evidence or scientific judgment. This issue is amplified by the productivity–quality divergence observed in recent studies: AI use can increase publication output, but AI-assisted papers with complex language may be less likely to be accepted [93]. As in Phase 1 (*Creation*), greater artifact production does not necessarily imply stronger research.

4.1.4 Findings and Observations

💡 Stage 5: Paper Writing

<div style="display: flex; align-items: center; margin-bottom: 10px;"> <div> <p style="margin: 0;">✔ State & Progress</p> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Commercial Near Threshold Cognitive </div> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Collaboration Rubric-Guided Transparency </div> </div> </div> <hr style="border-top: 1px dotted #ccc;"/> <ul style="list-style-type: none"> Paper writing is among the most widely adopted AI-assisted research stages, with tools supporting planning, drafting, polishing, citation assistance, and manuscript revision [109, 134]. Strong automated systems score 5.36 on the ICLR scale (<i>vs.</i> 5.69 accepted [220]); rubric-guided revision achieves 79% expert preference [258]. The gap to acceptance is argumentative depth, not fluency. Cognitive engagement and transparency are emerging design principles, aiming to preserve human understanding rather than merely producing polished text [186, 188]. 	<div style="display: flex; align-items: center; margin-bottom: 10px;"> <div> <p style="margin: 0;">⚠ Gaps & Limitations</p> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Paradox Mediocrity Detect Failing </div> <div style="display: flex; justify-content: space-between; font-size: 0.8em; margin: 5px 0;"> Shallow Args Attribution Deskilling </div> </div> </div> <hr style="border-top: 1px dotted #ccc;"/> <ul style="list-style-type: none"> Productivity and quality can diverge: AI-assisted workflows may increase output, but more fluent or complex language does not necessarily translate into stronger acceptance outcomes [93]. Valley of mediocrity: papers are fluent enough to look real but lack argumentative depth, experimental rigor, and reviewer-anticipation, the skills that separate publishable from near-publishable work. Detection tools have unacceptable false-positive rates, forcing a shift from “detect” to “declare” policies, while 17.5% of CS papers already carry detectable AI modification [109].
---	--

4.2 Summary and Transition: Writing

This phase shifts the focus from producing research artifacts to organizing them into a scholarly argument. **S5** (*Paper Writing*) takes the outputs of Phase 1, including ideas, retrieved literature, experiments, figures, and tables, and converts them into a manuscript that explains what was done, why it matters, and how the evidence supports the claims. Progress in this phase shows that AI systems are increasingly effective at assisting the writing workflow, from planning and drafting to polishing, citation support, revision, and even full-paper generation.

The central limitation is that fluent writing can conceal weak reasoning. AI-generated or AI-assisted text may improve readability and productivity while leaving deeper scientific requirements unresolved: whether claims are grounded, whether citations faithfully support them, whether experiments are sufficient, and whether the contribution is argued with appropriate nuance. This limitation appears both in semi-automated writing assistance and in fully automated paper generation. The former is most credible when it preserves researcher control over framing, interpretation, and final responsibility; the latter increasingly approaches reviewable quality in selected settings, but still struggles with argumentative depth, evidential rigor, and reviewer anticipation.

The output of this phase is a manuscript ready for external scrutiny. We therefore next turn to Phase 3 (*Validation*), where the manuscript is evaluated through peer review and revised through author rebuttal. This transition shifts the role of AI from structuring evidence into a coherent argument to assessing whether that argument is sound, fair, and sufficiently supported.

5 Phase 3: Validation

This phase encompasses the stages through which a manuscript produced in Phase 2 (*Writing*) is externally scrutinized and iteratively refined: peer review and rebuttal with revision. Together, these stages address a different question from *Creation* or *Writing*: *does this contribution meet the epistemic standards of the field?*

Validation is distinct because it introduces adversarial evaluation by reviewers who are expected to identify unsupported claims, methodological flaws, missing comparisons, unclear writing, and insufficient novelty. This makes Phase 3 a high-stakes setting for AI assistance. Automated systems can help summarize manuscripts, draft reviews, synthesize reviewer opinions, identify weaknesses, and support rebuttal preparation, but they also risk amplifying leniency, bias, adversarial manipulation, and weakly grounded criticism. The central challenge is therefore not whether AI can produce review-like text, but whether it can support fair, critical, and evidence-grounded evaluation without replacing independent expert judgment. The two stages also form a feedback loop: reviewer critiques in [s6](#) (*Peer Review*) may require additional experiments in [s3](#) (Coding and Experiments), revised figures in [s4](#) (*Tables and Figures*), or manuscript rewrites in [s5](#) (*Paper Writing*), while rebuttal and revision in [s7](#) (*Rebuttal and Revision*) determine how those critiques are addressed.

5.1 Peer Review

Peer review is the gateway to validation. It evaluates whether a manuscript is technically sound, sufficiently novel, clearly presented, and supported by appropriate evidence. Existing systems span automated review generation, meta-review drafting, reviewer–paper matching, and review quality assessment. Across these directions, the central limitation is that LLMs can often produce structured and plausible critiques, but may under-detect methodological flaws, over-score weak submissions, and remain vulnerable to adversarial manipulation.

A comprehensive inventory of automated review systems is provided in Table 8 (Appendix).

5.1.1 Automated Review Generation

Automated review generation aims to produce structured critiques of manuscripts, including summaries, strengths, weaknesses, questions, and rating recommendations. Existing approaches can be grouped into four broad families. *Fine-tuned reviewer models* specialize LLMs on expert review data to improve domain alignment and review format. DeepReviewer-14B [269] reports strong performance against GPT-o1 and on ICLR 2024 accept/reject prediction, while OpenReviewer [75] fine-tunes Llama-8B on 79K expert reviews. These systems show that supervised review data can improve review-style generation, but acceptance prediction remains a narrow proxy for review quality.

Multi-agent review systems decompose reviewing into specialized roles. MARG [35] uses multi-agent collaboration to generate multiple substantive comments per manuscript, the open-source ai-peer-review tool [150] uses multiple LLMs to produce independent reviews followed by meta-review synthesis, and ScholarPeer [55] extends this paradigm with literature search and claim verification. Such decomposition is useful because high-quality reviewing requires several distinct operations: understanding the manuscript, checking related work, evaluating claims, identifying weaknesses, and producing actionable feedback.

RL-optimized review systems attempt to improve review quality through more explicit training signals. REMOR [196] optimizes review generation with multi-objective reinforcement learning, while ReviewRL [247] combines retrieval-augmented context with RL to produce more comprehensive and grounded reviews. *Prompt-based systems* provide a lighter-weight alternative. Reviewer2 [45] introduces a two-stage framework that models the distribution of review aspects, while ChatReviewer [138] provides a deployed ChatGPT-based tool for analyzing strengths, weaknesses, and possible improvements. Overall, automated review generation has become increasingly structured, but review-like prose should not be mistaken for reliable validation: the core difficulty is whether critiques are accurate, calibrated, and grounded in the manuscript and relevant literature.

5.1.2 Meta-Review Generation

Meta-review generation synthesizes multiple reviewer opinions into a coherent area-chair-style assessment. This task differs from single-review generation because it must compare reviewer concerns, resolve disagreements, identify consensus, and justify a final recommendation. Bhatia *et al.* [66] evaluate GPT-3.5, LLaMA-2, and PaLM-2 on composing meta-review drafts from 40 ICLR papers, finding that LLMs are useful for multi-perspective summarization but struggle with nuanced judgment calls. AgentReview [83] simulates the full review lifecycle, including meta-review and final decisions, and shows that social influence and authority bias can affect outcomes.

The main challenge is not summarization alone, but decision-making under disagreement. When reviewers fundamentally disagree about a manuscript’s contribution, LLMs often produce diluted compromises rather than taking a defensible substantive position. This limitation reflects a broader issue in AI-assisted validation: current systems are better at consolidating stated opinions than at independently adjudicating technical disputes.

5.1.3 Reviewer Matching

Reviewer–paper matching supports the editorial process by assigning manuscripts to reviewers with relevant expertise while accounting for conflicts of interest. This task is less visible than review generation, but it is crucial for review quality at scale: even a well-written review is less useful if the reviewer lacks appropriate domain expertise. RelevAI-Reviewer [34] has been deployed at major venues, while RATE [118] improves expertise-based matching through profile distillation, aiming to capture a reviewer’s competence signature beyond keyword overlap.

Compared with automated reviewing, reviewer matching is a more appropriate setting for operational AI support because it assists the allocation process rather than replacing expert judgment. However, matching systems still require transparent conflict handling, robust expertise modeling, and human oversight, especially in interdisciplinary areas where surface-level keyword similarity can be misleading.

5.1.4 Assessment: Review Consistency, Bias, and Robustness

Assessing AI-assisted peer review requires more than measuring whether generated reviews resemble human reviews. A useful review must be consistent, critical, fair, grounded, and robust to manipulation. On consistency, recent systems show measurable progress. The Stanford Agentic Reviewer [80] achieves Spearman correlation of 0.42, comparable to human–human correlation of 0.41. ReviewAgents [43] uses a multi-agent framework trained on a Review-CoT dataset of 37,403 papers and 142,324 reviews. The reviewer component reported in the Nature version of The AI Scientist [123] reaches 69% balanced accuracy on ICLR acceptance prediction, while ClaimCheck [142] evaluates whether LLM critiques are grounded in the reviewed manuscript. ReViewGraph [105] further models multi-round reviewer–author debates as heterogeneous graphs, improving debate outcome prediction without LLM fine-tuning.


However, consistency is not sufficient. A reviewer can be consistent while being systematically lenient, biased, or shallow. LLM-based reviewers have been shown to assign inflated scores relative to humans, and in some settings to misclassify rejected papers as acceptable [267]. This makes standalone AI review risky: it may produce coherent critiques while failing to identify decisive methodological weaknesses. A more credible deployment mode is to use LLMs to improve human reviews rather than to replace reviewers. In a randomized ICLR 2025 study across 22,467 reviews, LLM feedback on reviews improved review quality in 89% of cases, with reviewers updating their reviews 26.6% of the time, without affecting acceptance rates [202]. Chen *et al.* [29] further study how reviewers engage with AI-generated feedback during a live ICLR 2025 process, and Zhuang *et al.* [272] provide a broader taxonomy of automated scholarly paper review methods.

Robustness and governance remain major concerns. The “AI Review Lottery” [164] estimates that at least 15.8% of ICLR 2024 reviews were AI-assisted, with 49.4% of submissions receiving at least one AI-assisted review. A 2025 Nature survey similarly reports that many academics use AI in peer review despite restrictive venue policies [134], and a major 2026 conference rejected 497 papers for AI-use policy violations [50]. These trends indicate that AI involvement in peer review is already widespread, while enforceable governance remains immature.

Adversarial manipulation further complicates deployment. Breaking the Reviewer [113] studies adversarial robustness of LLM-based review assessments, while Keuper [88] shows that simple prompt injections, such as white text on a white background, can manipulate LLM reviews. Ye *et al.* [242] show that covert content injection can substantially raise review scores and that manipulating a small fraction of reviews can alter rankings. Zhou *et al.* [266] further demonstrate that in-paper prompt injection can raise LLM scores under static and iterative attacks. At the lexical level, Raina *et al.* [157] show that benign adjectives can function as adversarial triggers, while Sahoo *et al.* [166] evaluate indirect manipulation across multiple LLMs and attack strategies, with frontier models beginning to show stronger resistance in some settings. Finally, detection-based policy enforcement remains fragile: Saha *et al.* [165] show that state-of-the-art AI text detectors misclassify LLM-polished reviews, and Yu *et al.* [244] evaluate 18 detection algorithms on 788,984 AI-written peer reviews, highlighting the difficulty of identifying AI-generated review text at the individual-review level.

5.1.5 Findings and Observations

💡
Stage 6: Peer Review



✔ **State & Progress**


Deployment
Consistency
Mapped Risks

Multi-Agent
RL-Trained
Matching

- The strongest validated deployment mode is LLM feedback on *reviews*, not standalone AI review: in ICLR 2025, review feedback improved quality in 89% of cases without affecting acceptance rates [202].

- Automated reviewers can approach human-level consistency on selected metrics, with the Stanford Agentic Reviewer matching human inter-rater agreement ($\rho = 0.42$ vs. 0.41) [80].

- Reviewer matching and meta-review generation are promising support tasks because they assist editorial coordination and opinion synthesis rather than directly replacing expert judgment.



⚠ **Gaps & Limitations**

Leniency
Fragility
Policy

Inflation
Injection
Undetectable

- Standalone AI review remains unsafe: LLMs assign inflated scores (6.86 vs. human 5.70 [267]), misclassifying 95.8% of rejected papers as acceptable. Standalone AI review is unsafe.

- Adversarial fragility persists: prompt injection reaches 10 scores [266]; benign adjectives function as universal triggers [157]; 5% manipulation flips 12% of rankings [242].

- Governance is difficult because AI-assisted reviewing is already prevalent [164], yet all five SOTA detectors fail on polished reviews [165]. Prevalence has outpaced governance.

5.2 Rebuttal and Revision

Rebuttal and revision form the second stage of Phase 3 (*Validation*), where authors respond to external critique and revise the manuscript before a final decision or camera-ready submission. This stage is epistemologically important because it is the only point in the publication process where authors engage directly with reviewers’ objections. Existing work spans reviewer-comment analysis, automated rebuttal generation, and revision tracking. Across these directions, the central challenge is not merely generating persuasive responses, but ensuring that rebuttals are evidence-grounded, faithful to the manuscript, and followed by actual revisions.

A comprehensive inventory of rebuttal systems is provided in Table 9 (Appendix).

5.2.1 Reviewer Comment Analysis

Reviewer-comment analysis decomposes critiques into actionable concerns, such as missing experiments, unclear motivation, insufficient baselines, unsupported claims, or presentation issues. This analysis is a prerequisite for effective rebuttal generation because reviewer comments are often long, mixed in priority, and partially overlapping across reviews. ReviewMT [197] models peer review as multi-turn, long-context dialogue with role-based interactions, covering 26,841 papers and 92,017 reviews from ICLR and Nature

Communications. Re² [250] further provides a consistency-ensured dataset for full-stage peer review and multi-turn rebuttal, covering 19,926 submissions, 70,668 reviews, and 53,818 rebuttals from 24 conferences.

Empirical studies show that rebuttal can materially affect outcomes, especially for borderline submissions. Analysis of ICLR 2024–2025 [86] reports that 75–81% of scores remain unchanged after rebuttal, 17–23% improve, and only approximately 1% decrease, with the most common transition being 5 → 6 from borderline to acceptable. These numbers suggest that rebuttal is not a universal remedy, but it is consequential for the subset of papers where reviewer concerns can be clarified, corrected, or supported with additional evidence.

5.2.2 Automated Rebuttal Generation

Automated rebuttal generation attempts to produce author responses that address reviewer concerns clearly and strategically. Early systems treated rebuttal as direct text generation, but this formulation is prone to hallucination, missed reviewer points, and unverifiable claims. More recent systems therefore decompose rebuttal into intermediate steps such as concern extraction, evidence retrieval, response planning, and final generation.

RebuttalAgent [63] uses Theory-of-Mind modeling to craft strategically persuasive responses, reporting an average 18.3% improvement over the base model. Paper2Rebuttal [129] introduces evidence-centric planning by decomposing reviewer comments into atomic concerns and retrieving supporting literature, improving Coverage and Specificity by up to +0.78 and +1.33. ReviewerToo [167] includes a rebuttal module within a broader modular framework and reports 81.8% accept/reject accuracy.

A newer direction emphasizes planning and author control. DRPG [60] proposes a four-step Decompose–Retrieve–Plan–Generate pipeline, reporting 98%+ planning accuracy and stronger-than-average human rebuttal quality with an 8B model. Author-in-the-Loop [163] integrates author expertise and intent into response generation, aiming to ensure that rebuttals reflect the paper’s actual contributions rather than generic LLM output. These systems indicate that rebuttal automation is moving from fluent response generation toward evidence-grounded and author-aware revision support.

5.2.3 Assessment: Rebuttal Effectiveness


Assessing rebuttal systems requires measuring both immediate effectiveness and downstream accountability. Immediate effectiveness concerns whether a rebuttal addresses reviewer concerns, clarifies misunderstandings, provides evidence, and improves reviewer confidence. ICLR 2024–2025 analysis shows that papers whose scores improve after rebuttal achieve acceptance rates of 55.7–57.6%, compared to 7.8–12.4% for papers with unchanged scores [86]. This makes rebuttal especially important for borderline papers, where small changes in reviewer confidence can affect final outcomes.

However, effective rebuttal is not only persuasion. Many reviewer requests require new experiments, additional ablations, corrected figures, or manuscript restructuring, creating a feedback loop from (S7) back to (S3), (S4), and (S5). Ruan and Gurevych [163] provide large-scale aligned review–response–revision triplets, enabling the study of how rebuttals translate into actual manuscript changes. Current rebuttal systems can retrieve evidence and draft responses, but they generally cannot generate new experimental evidence in response to reviewer requests. This makes the rebuttal–experiment loop one of the most practically important gaps in current auto-research pipelines.


The second evaluation dimension is accountability. A recent audit [21] finds that ICLR 2025 authors make an average of 11.8 commitments per paper during rebuttal, but approximately 25% of these commitments are not fulfilled in the camera-ready version, with missing experiments among the most common unfulfilled promises. This gap exposes the same capability-versus-integrity tension observed throughout the survey: AI systems may generate plausible and persuasive responses, but the scientific validity of a rebuttal depends on whether its claims are supported and its promises are later implemented. Rebuttal systems should therefore be evaluated not only by response quality, but also by concern coverage, evidence grounding, revision traceability, and fulfillment of author commitments.

5.2.4 Findings and Observations

Stage 7: Rebuttal & Revision

**State & Progress**

- Newest
- Human-Level
- Decomposition
- Decisive
- Evidence
- Planning

**Gaps & Limitations**

- No New Expts
- Commitment
- Overlooked
- Accountability
- Only 10 Tools
- Loop Gap

- Rebuttal automation is emerging as a distinct stage, with recent systems moving from direct response generation toward decomposition, evidence retrieval, response planning, and author-aware generation.
- Rebuttal is consequential for borderline submissions: 17–23% of ICLR 2024–2025 submissions improve scores after rebuttal, and improved-score papers achieve much higher acceptance rates [86].
- Evidence-centric planning helps address common failures of direct generation, including missed reviewer concerns, unsupported responses, and generic rebuttal text [63, 129].
- Current systems cannot reliably generate new experimental evidence in response to reviewer requests; the **S7** (*Rebuttal and Revision*) → **S3** (*Coding*) feedback loop remains a major unautomated gap.
- The quality of the rebuttal process must be tied to revision fulfillment: approximately 25% of ICLR 2025 rebuttal commitments are not fulfilled in camera-ready versions [21].
- Rebuttal remains under-served relative to its practical importance, despite being the stage where authors directly negotiate reviewer concerns before final decisions.

5.3 Summary and Transition: Validation

This phase shifts the focus from constructing a manuscript to testing whether its claims withstand external scrutiny. **S6** (*Peer Review*) evaluates the manuscript through independent critique, while **S7** (*Rebuttal and Revision*) gives authors an opportunity to clarify, defend, and revise the work in response. Together, these stages form a feedback loop rather than a one-way checkpoint: reviewer comments can trigger new experiments, revised analyses, updated figures, and substantial manuscript rewriting.

Progress in Validation shows a consistent pattern. AI systems are increasingly capable of producing review-like critiques, summarizing reviewer opinions, supporting reviewer matching, and drafting rebuttal responses. However, the hard part of validation is not producing plausible evaluative text; it is making fair, critical, evidence-grounded judgments and ensuring that critique leads to accountable revision. In peer review, standalone AI reviewers may be consistent but lenient, biased, or vulnerable to manipulation, while the strongest validated deployment mode is to use AI to improve human reviews. In rebuttal, AI can help decompose concerns and draft evidence-aware responses, but cannot yet generate missing experimental evidence or guarantee that author commitments are fulfilled.

The output of this phase is a manuscript that has been externally challenged, defended, and revised. Once validated, the research must be communicated beyond the review process through posters, slides, videos, project pages, and social media. We therefore next turn to Phase 4 (*Dissemination*), where AI assistance shifts from evaluating scholarly claims to adapting validated research artifacts for different audiences, formats, and communication goals.

6 Phase 4: Dissemination

This phase converts the validated manuscript into formats accessible to audiences beyond specialist venue readers. The discussion covers the transformation of papers into posters, slides, videos, social media posts, and interactive agents. Compared with earlier phases, Dissemination is less about producing or validating scientific claims than about adapting those claims to different audiences, media, and interaction modes.

Dissemination merits a separate phase because its outputs are independent knowledge artifacts rather than

simple derivatives of the paper. A poster must compress the contribution into a single visual narrative; a slide deck must support oral explanation; a video must synchronize visual, textual, and spoken channels; a social media post must balance accessibility with precision; and an interactive agent must expose the paper’s methods for downstream use. The central challenge is therefore not whether AI can reformat a paper, but whether it can preserve scientific fidelity while adapting the work to new modalities, audiences, and levels of interactivity.

6.1 Research Dissemination (Paper2X)

Research dissemination converts a completed paper into audience-adaptive artifacts. Unlike Phases 1–3, which primarily target specialist authors, reviewers, and readers, Paper2X outputs must serve diverse audiences: conference attendees, oral-session audiences, online readers, prospective users, journalists, practitioners, and future researchers who may interact with the work through tools rather than text. Existing systems cover poster generation, slide generation, narrated video and talk generation, social media and web-page generation, and emerging paper-to-agent conversion. Across these formats, the core bottleneck is trust: researchers may use AI to draft public-facing materials, but they remain reluctant to delegate final communication to systems that may distort results, overstate claims, or omit important limitations.

A comprehensive inventory of Paper2X systems across poster, slides, video, web, social media, and agentic formats is provided in Table 10 (Appendix).

6.1.1 Paper to Posters

Paper-to-poster generation transforms a full manuscript into a compact visual narrative. This requires more than summarization: the system must select the central message, allocate space across motivation, method, results, and conclusion, preserve key figures and tables, and arrange them into a readable layout. Compared with slide generation, poster generation has stronger spatial constraints because all content must be legible and coherent on a single canvas.

Early systems established agentic poster generation as a feasible task. Paper2Poster [146] introduces PosterAgent with binary-tree layout planning and a Painter–Commenter feedback loop, showing that poster generation can be decomposed into layout construction, rendering, and critique. Subsequent systems add stronger design and hierarchy awareness. PosterGen [257] incorporates aesthetic-aware multi-agent generation, PosterForest [33] uses hierarchical multi-agent collaboration, and P2P [194] introduces P2PInstruct with specialized agents and instruction data for poster design.

Recent systems move from one-shot poster creation toward editing and unified poster manipulation. APEX [180] supports interactive poster editing with fine-grained control, addressing the practical need for human post-editing in conference preparation. PosterOmni [27] unifies multiple poster tasks, including rescaling, filling, extension, layout-driven generation, style-driven generation, and identity-driven generation, while PosterCraft [28] further explores quality-aware poster generation in a unified framework. Together, these systems suggest that poster automation is shifting from direct paper summarization toward editable, design-aware poster production.

6.1.2 Paper to Slides

Paper-to-slides systems convert manuscripts into sequential visual narratives for oral presentation. Unlike posters, slides unfold over time and must support speaker delivery. This requires content selection, section-to-slide mapping, visual layout, speaker-note synthesis, and often iterative refinement based on rendered slide quality. The key challenge is preserving the paper’s argument while changing its rhetorical structure from written exposition to spoken explanation.

Early datasets and pipelines established the task. DOC2PPT [41] provides paired document–slide data, while PPTAgent [262] generates and evaluates presentations with PPTEval across content, design, and coherence. Environment-grounded refinement then closes the gap between symbolic planning and rendered slides. DeepPresenter [263] conditions revision on rendered slide images rather than only internal reasoning traces, showing that visual feedback is important for presentation quality.

Multi-agent and interactive systems further decompose slide generation into specialized subtasks. SlideGen [111] uses agents for outlining, content mapping, arrangement, note synthesis, and iterative refinement to produce editable PPTX slides. Auto-Slides [235] targets Beamer generation with multi-agent collaboration and interactive editing. SlideTailor [248] conditions generation on user preference from a single example pair using a chain-of-speech mechanism. Other systems focus on task-specific capabilities: PASS [3] combines slide generation with AI audio delivery, AutoPresent [48] fine-tunes a slide-generation model on SlidesBench, Paper2Slides [64] provides one-click conversion through a multi-stage RAG pipeline, Talk to Your Slides [84] supports natural-language slide editing, and Office Raccoon [173] targets page-level editing with template and brand-guideline learning. Across these systems, the main trend is from static slide generation toward editable, feedback-aware, and user-preference-conditioned presentation design.

6.1.3 Paper to Videos and Talks

Paper-to-video and paper-to-talk systems extend dissemination from visual artifacts to multimodal explanation. These systems must coordinate slides, subtitles, narration, cursor motion, pacing, and sometimes avatar or talking-head video. This makes the task substantially harder than poster or slide generation: errors can arise not only from content selection, but also from temporal alignment, speech clarity, visual synchronization, and duration constraints.

PresentAgent [181] provides an end-to-end document-to-narrated-video pipeline with synchronized slides, text-to-speech narration, and the PresentEval benchmark. Paper2Video [271] introduces a benchmark of paper–video pairs and the PaperTalker framework, decomposing video generation into slide, subtitle, cursor, and talker builders. Preacher [116] uses top-down decomposition followed by bottom-up generation with Progressive Chain of Thought across multiple research fields.

Although these systems show promising progress, video remains one of the hardest Paper2X formats. Unlike posters and slides, video generation requires coordination across at least four modalities: visual slides, subtitles, speech audio, and temporal or avatar-based presentation. The resulting artifact must also remain faithful to the paper while being concise enough for viewers to follow. Current systems, therefore, work best as first-draft generators that produce synchronized presentation assets for human review, rather than final public-facing videos requiring no editing.

6.1.4 Paper to Social Media

Paper-to-social-media and paper-to-web generation aims to make research discoverable outside the publication venue. Outputs include project pages, blog posts, press-release-style summaries, short-form research posts, and X/Twitter threads. These formats require stronger audience modeling than posters or slides: a thread for ML practitioners, a lay summary for journalists, and a project-page introduction for potential users should emphasize different details, use different vocabulary, and make different assumptions about background knowledge.

Paper2Web [30] converts papers into interactive multimedia-rich academic homepages and provides a benchmark for this task. More generally, researchers increasingly use general-purpose LLMs to draft online summaries, figure captions, project-page text, and social media announcements. However, dedicated research-to-social-media tools remain comparatively underdeveloped. The bottleneck is not text generation alone, but audience-adaptive fidelity: systems must simplify without distorting, emphasize contributions without exaggeration, and preserve limitations while remaining engaging.

This makes social dissemination a distinct trust problem. Public-facing outputs are often read without the paper, so any overclaim, missing caveat, or misleading comparison can shape how the work is perceived. AI assistance is therefore most credible when it provides audience-specific drafts, style variants, and claim-checking support, while leaving final messaging and factual responsibility to the authors.

6.1.5 Paper to Agents and Tools

A newer dissemination direction converts papers from static documents into interactive agents or tools. This changes the function of dissemination: instead of only explaining a contribution, the system exposes the

paper’s methods, code, data, or workflows through natural-language interaction. In this setting, the reader becomes a user who can query, reproduce, adapt, or extend the work.

Paper2Agent [132] exemplifies this shift by converting research papers and associated codebases into interactive paper agents. The system analyzes the paper and code, constructs a Model Context Protocol (MCP) server with tools, resources, and prompts, and iteratively tests the resulting agent so that users can interact with the paper’s methods through natural language. This reframes dissemination as operational access: a paper is no longer only read, but queried and executed.

Related systems broaden this idea from paper-specific agents to tool-using scientific agents. Gao *et al.* [42] study how scientific tool ecosystems can democratize AI scientists by exposing computational capabilities through agent-accessible interfaces. ProteinMCP [228] applies an MCP-based agentic framework to protein engineering, illustrating how domain-specific workflows can be wrapped into interactive, tool-using systems. These systems suggest that future dissemination may increasingly involve executable interfaces, reproducible workflows, and domain agents that make research easier to reuse.

This direction also introduces new risks. An interactive paper agent must not only summarize the paper faithfully, but also execute tools correctly, respect the limitations of the original method, and avoid presenting unsupported extrapolations as valid conclusions. Evaluation therefore requires both communication metrics and reproducibility metrics: whether the agent explains the paper clearly, whether it invokes the correct tools, whether it reproduces expected results, and whether it handles out-of-scope user queries responsibly.

6.1.6 Assessment: Fidelity, Usability, and Adoption


Assessment for Paper2X must evaluate three dimensions: *fidelity*, whether the generated artifact accurately represents the paper; *usability*, whether the artifact supports its intended communication or interaction goal; and *adoption*, whether researchers trust the artifact enough to use it publicly. Fidelity is the most important dimension because dissemination artifacts often circulate independently of the paper. A polished poster, slide, video, or thread can misrepresent a contribution if it omits caveats, changes baselines, simplifies methods incorrectly, or exaggerates results.

Poster and slide generation have the most mature evaluation infrastructure. Paper2Poster [146] evaluates poster quality under cost-efficient generation settings, while PPTAgent [262] introduces PPTEval to assess slide quality along content, design, and coherence dimensions. Video evaluation is newer: PresentEval [181] evaluates narrated video pipelines, and Paper2Video [271] introduces comprehension-oriented evaluation through paper–video pairs. These benchmarks reflect a broader shift from surface aesthetics toward whether generated materials preserve content and improve audience understanding.

Agentic dissemination requires additional evaluation criteria. A paper agent should be assessed not only by answer quality, but also by tool correctness, reproducibility, error handling, and boundary awareness. This makes Paper2X evaluation closer to the assessment problems in Coding and Experiments (S3): a system may appear helpful in natural language while invoking the wrong workflow or returning unsupported outputs. Across formats, no large-scale adoption study has yet established whether Paper2X tools reduce or increase misrepresentation compared with manual author-created materials. As a result, Paper2X systems are currently best understood as drafting and interaction aids rather than final producers.


6.1.7 Findings and Observations

Stage 8: Dissemination (Paper2X)



State & Progress

Low Cost Paper2Poster Paper2Slides
Paper2Video Multi-Agent Interactive



Gaps & Limitations

Trust 4-Modal Hard No Adaptation
Adoption Fidelity Social Media

- Cost barrier eliminated: \$0.005/poster with 87% fewer tokens [146]; 8B models match frontier on slides [263]. Most cost-efficient stage to automate.
- Poster and slide generation are the most developed directions, moving from one-shot conversion toward editable, feedback-aware, and user-preference-conditioned workflows [146, 262, 263].
- Paper-to-agent systems extend dissemination from static explanation to interactive reuse, exposing paper methods, code, and workflows through tool-using agents [132, 228].
- Trust, not generation cost, is the bottleneck: researchers need confidence that AI-generated public artifacts preserve claims, caveats, and limitations.
- Video remains difficult because it must coordinate slides, subtitles, narration, pacing, and sometimes avatar or cursor motion under strict time constraints [116, 271].
- Social media and web dissemination remain limited by audience modeling: simplifying a paper for broader reach without exaggerating or distorting the contribution remains unresolved.

6.2 Summary and Transition: Dissemination

This phase shifts the focus from validated scholarly argument to audience-adaptive communication and reuse. The goal is to convert the manuscript into posters, slides, videos, project pages, social media posts, and increasingly interactive agents or tools. Progress in this phase shows that AI can substantially lower the cost of producing dissemination artifacts, especially when the input paper is complete and the target format has a predictable structure.

The central limitation is fidelity under format change. Each dissemination artifact compresses, reorders, or re-expresses the paper, creating opportunities for omission, overstatement, or distortion. This risk appears differently across formats: posters and slides may oversimplify the contribution, videos may misalign narration and visual evidence, social media posts may trade nuance for engagement, and paper agents may expose tools or workflows beyond their validated scope. Dissemination-stage automation is therefore most credible when it supports draft generation, format adaptation, editing, and interaction while preserving author oversight over claims and limitations.

The completion of this phase closes the research lifecycle: a contribution has been created, written, validated, and communicated. The remaining question is what patterns cut across all phases. We therefore next synthesize the common architectures, capability boundaries, deployment principles, and open challenges that define AI-assisted research as a whole.

7 Cross-Cutting Analysis

The preceding sections analyzed AI-assisted research stage by stage. We now synthesize patterns that emerge across the complete lifecycle. This cross-cutting view is necessary because many of the most important limitations do not appear within a single stage, but at the boundaries between stages: ideas that weaken after implementation, retrieved evidence that is misrepresented in writing, experiments that produce unsupported claims, reviews that miss methodological flaws, rebuttals that promise revisions without fulfilling them, and dissemination artifacts that simplify results beyond the evidence.

We organize this analysis around **four** questions. First, how do end-to-end systems integrate multiple stages of the research lifecycle? Second, how should research automation be evaluated across heterogeneous artifacts

and long-horizon workflows? Third, what recurring capability boundaries and deployment principles appear across phases? Finally, what open challenges must be addressed before AI systems can be trusted as reliable research collaborators rather than artifact generators?

7.1 End-to-End Research Systems

Some systems discussed earlier, especially in Section 4, also qualify as end-to-end research systems because they generate complete manuscripts. Here, however, we analyze them from a different perspective: not as paper-writing tools, but as lifecycle-scale architectures. The key question is how these systems connect ideation, literature review, coding, experimentation, writing, validation, and dissemination, and where the handoffs between stages remain fragile.

Most current end-to-end systems emphasize Phase 1 (*Creation*) and Phase 2 (*Writing*), connecting idea generation, implementation, experiment execution, and manuscript drafting into a single workflow. Far fewer systems incorporate substantive Phase 3 (*Validation*), such as adversarial review, rebuttal planning, or revision tracking, and Phase 4 (*Dissemination*) remains mostly outside current end-to-end pipelines. This imbalance reflects a broader pattern observed throughout the survey: it is easier to generate research artifacts than to validate, revise, and communicate them with accountable fidelity.

Existing systems can be grouped into four architectural families: sequential pipelines, search-based and self-improving systems, skill-based and tool-integrated systems, and multi-agent or community-scale frameworks.

7.1.1 Sequential and Pipeline-Based Systems

Sequential systems connect research stages in a mostly linear order, typically moving from idea generation to experiment execution and manuscript drafting. The AI Scientist [122] established this paradigm by demonstrating that hypothesis generation, code execution, experimental analysis, and paper writing can be assembled into a single automated workflow. Agent Laboratory [171], AI-Researcher [199], CycleResearcher [220], Kosmos [133], Dolphin [245], CodeScientist [77], and InternAgent [249] instantiate related pipeline designs with different choices of base models, task scopes, and evaluation targets.

The advantage of sequential architectures is operational simplicity. Each stage produces an artifact that becomes the input to the next stage, making the workflow interpretable and relatively easy to implement. The limitation is error propagation. A weak idea can lead to irrelevant experiments; incorrect code can produce misleading results; and unsupported experimental claims can be polished into a plausible manuscript. Sequential pipelines therefore expose the same phase-boundary risk observed throughout the lifecycle: producing an artifact at one stage does not guarantee that the next stage represents it faithfully or verifies it adequately.

7.1.2 Search-Based and Self-Improving Systems

A second family introduces search, evolution, or self-improvement to avoid the brittleness of one-pass generation. AI Scientist v2 [229] uses agentic tree search to explore research trajectories more systematically than its predecessor. ASI-Evolve [227], AutoSOTA [107], CORAL [155], and related evolutionary systems search over architectures, algorithms, data curation strategies, or multi-agent behaviors to discover stronger solutions.

Search-based designs are important because research rarely proceeds as a single pass. Strong work typically emerges from branching alternatives, failed experiments, ablation-driven refinement, and selective continuation of promising directions. These systems therefore better match the iterative structure of scientific practice than direct pipelines. However, search alone does not solve validation. Without reliable evaluators, search can optimize toward benchmark-specific artifacts, superficial novelty, or brittle improvements. The core design question is thus not only how broadly a system searches, but what signals guide selection and whether those signals reflect scientific value rather than local metric gains.

7.1.3 Skill-Based and Tool-Integrated Systems

Skill-based systems package research workflows as composable capabilities, often built around coding agents, retrieval tools, experiment runners, document editors, and evaluation modules. ARIS [233] represents this direction by organizing research automation into reusable workflows for idea discovery, auto-review, and

paper writing. AutoResearchClaw [117] similarly implements a multi-stage pipeline with internal agents for coding, benchmarking, and figure generation. Broader tool-integrated systems such as AutoAgent [198], Biomni [72], SciSciGPT [177], and ResearchClaw [232] emphasize retrieval, analytics, code execution, document understanding, and domain-specific tool use.

The strength of skill-based architectures is modularity. Rather than relying on one monolithic agent to perform all research activities, these systems expose explicit tools and reusable skills for individual operations. This makes it easier to inspect intermediate artifacts, swap components, and insert human checkpoints. The limitation is coordination: modular systems still need reliable state management across stages. If the idea, literature trace, code state, experimental logs, manuscript claims, review feedback, and revision plan are not represented in a shared and updateable workspace, then phase handoffs remain fragile despite the presence of many tools.

7.1.4 Multi-Agent and Community-Scale Systems

Multi-agent systems distribute research tasks across specialized agents, such as researchers, engineers, reviewers, analyzers, writers, or simulated community members. freephdlabor [100], SciMaster [19], EvoScientist [127], UniScientist [99], Medical AI Scientist [222], AiScientist-LH [23], FARS [8], and AutoResearchClaw [117] illustrate different forms of multi-agent orchestration. Related community-scale systems such as VirSci [193], AgentRxiv [170], and ResearchTown [243] further simulate aspects of scientific collaboration, including idea exchange, manuscript writing, review, and revision.

The motivation for multi-agent architectures is that research requires heterogeneous expertise and adversarial feedback. A single model asked to generate, execute, write, and critique its own work is prone to self-confirmation. Separating roles can reduce this risk by introducing specialization and cross-agent critique. However, multi-agent systems also introduce coordination problems: agents may duplicate work, reinforce shared misconceptions, defer to weak signals, or produce verbose deliberation without improving scientific quality. The strongest multi-agent systems therefore require more than agent count; they require clear role separation, shared memory, grounded tools, and explicit verification mechanisms.

7.1.5 Assessment: Lifecycle Coverage and Phase Boundaries


End-to-end systems should be evaluated not only by the quality of their final manuscript, but also by lifecycle coverage and phase-boundary reliability. Current systems are strongest at generating ideas, code, experiments, and paper drafts, but weaker at external validation, author-style revision, and audience-adaptive dissemination. This uneven coverage is not accidental: Stage 1 (*Creation*) and Stage 2 (*Writing*) produce artifacts, whereas Stage 3 (*Validation*) and Stage 4 (*Dissemination*) require judgment, accountability, and audience-aware fidelity.

The most important failure mode is not isolated stage failure, but unverified handoff. An idea may appear novel but fail during execution; code may run but implement the wrong algorithm; experimental logs may be summarized into unsupported claims; an automated review may be coherent but lenient; and a rebuttal may promise changes that are not fulfilled. End-to-end systems amplify this risk because errors can propagate silently across stages. A mature lifecycle-scale research system must therefore preserve traceable links between hypotheses, retrieved evidence, code, experiments, figures, manuscript claims, reviews, rebuttals, and revisions.

Reported costs and quality metrics further suggest that the token budget alone is not the decisive factor. Systems vary widely in cost, but stronger results often come from search strategy, tool integration, structured decomposition, and verification design rather than brute-force generation. The central evaluation question is therefore shifting from *Can the system produce a paper?* to *Can the system maintain scientific fidelity across the complete lifecycle?*

7.1.6 Findings and Observations

End-to-End Research Systems




State & Progress

Pipelines Search Skills

Multi-Agent Validation Handoffs

- End-to-end systems increasingly move beyond linear pipelines toward search-based, skill-based, and multi-agent architectures that better reflect the iterative structure of research.
- Phase coverage remains uneven: most systems cover *Creation* and *Writing*, while substantially fewer incorporate *Validation*, and none yet provide mature *Dissemination* coverage.
- Systems that include review, critique, or revision mechanisms point toward more credible lifecycle automation, but their success depends on verification quality rather than review-like text alone.



Gaps & Limitations

Propagation Self-Critique Fragmented

Validation Gap State Loss Overclaiming

- Error propagation is the main lifecycle risk: weak ideas, semantic code errors, unsupported claims, and lenient reviews can compound when phase handoffs are not explicitly verified.
- Single-model self-critique remains structurally limited; credible validation usually requires role separation, external evidence, adversarial review, or human oversight.
- Most existing E2E systems do not maintain a fully traceable, updateable state across hypotheses, literature review, coding, experiments, manuscript claims, reviews, and revisions.

7.2 Evaluation Across the Research Lifecycle

Evaluation is the central bottleneck for AI-assisted research. Each stage produces different artifacts—ideas, literature summaries, code, experiments, figures, manuscripts, reviews, rebuttals, and dissemination materials—so no single metric can capture research quality across the full lifecycle. Existing benchmarks have therefore evolved from narrow task-specific evaluations toward broader, process-aware, and increasingly execution-grounded protocols. Table 2, introduced in Section 2.3, summarizes major benchmarks across the eight stages.

Across the benchmark landscape, three trends are clear. First, evaluation is moving from isolated outputs to a multi-dimensional assessment. Early benchmarks often measured a single capability, such as citation prediction, code execution, or writing fluency. Recent benchmarks instead evaluate multiple axes, such as novelty and feasibility in ideation, coverage and citation accuracy in literature synthesis, semantic correctness in research code, review consistency in peer review, and fidelity in Paper2X artifacts. Second, benchmarks are becoming more domain- and workflow-aware. Specialized evaluations now target GPU kernel optimization [101, 143], biology research [96], scientific experimentation [92], and broader scientist-aligned workflows [16, 226]. Third, a persistent gap remains between benchmark performance and real-world research value: systems can perform well on measurable proxies while still producing outputs that experts judge as shallow, incremental, or insufficiently grounded [220, 267].

7.2.1 Stage-Specific Benchmarks

Stage-specific benchmarks remain necessary because each part of the research lifecycle requires different evaluation criteria.

- For **S1** (*Idea Generation*), benchmarks assess novelty, feasibility, diversity, and downstream potential. These evaluations are difficult because apparent novelty may not survive implementation, and expert judgments of research promise are inherently noisy.
- For **S2** (*Literature Review*), benchmarks emphasize retrieval precision, citation fidelity, coverage completeness, and synthesis quality. The central challenge is not only whether the system finds relevant papers, but whether it uses them faithfully when constructing a narrative.

- For [§3](#) (*Coding and Experiments*), benchmarks increasingly move beyond code execution toward semantic correctness and reproducibility. Research-code benchmarks ask whether generated implementations match the intended algorithm, while broader workflow benchmarks such as EXP-Bench [92] evaluate experiment design, execution, and analysis.
- For [§4](#) (*Tables and Figures*), evaluation must distinguish visual plausibility from scientific correctness: a figure or table may look publication-ready while misrepresenting data, notation, or comparison structure.
- For [§5](#) (*Paper Writing*), evaluation combines writing quality, citation accuracy, factual grounding, and review-style judgment.
- For [§6](#) (*Peer Review*), benchmarks assess consistency, grounding, bias, and robustness to manipulation rather than review fluency alone.
- For [§7](#) (*Rebuttal and Revision*), emerging datasets align reviews, author responses, and manuscript changes, enabling evaluation of whether rebuttals actually address concerns and lead to fulfilled revisions.
- For [§8](#) (*Dissemination*), evaluation centers on fidelity, usability, and audience adaptation across posters, slides, videos, project pages, social media posts, and interactive agents. This stage is especially difficult because dissemination artifacts often circulate independently of the paper and can shape public understanding of the work.

7.2.2 Evaluation Methodologies

Current evaluation methodologies can be grouped into five families. *Expert evaluation* remains the most credible approach for assessing novelty, significance, correctness, and scientific contribution. Si *et al.* [183], for example, recruited over 100 NLP researchers to evaluate generated ideas. However, expert evaluation is expensive, slow, and noisy: even peer-review-style judgments show limited inter-rater agreement, with the Stanford Agentic Reviewer study reporting human-human correlation around $\rho = 0.41$ [80]. This makes expert evaluation indispensable but difficult to scale.

LLM-as-Judge and Agent-as-Judge methods provide scalable approximations of human assessment. CycleReviewer [220] reports a 26.89% reduction in Proxy MAE relative to individual human reviewers for score prediction, while the Stanford Agentic Reviewer [80] achieves review-score correlations comparable to human inter-rater agreement ($\rho = 0.42$ vs. human $\rho = 0.41$). These results show that automated evaluators can provide useful review-style signals, but they remain imperfect proxies. They can exhibit positivity bias, length bias, authority bias, self-preference, and vulnerability to adversarial prompts [241, 267]. As a result, LLM-based evaluation is most reliable when calibrated against expert judgments and combined with task-specific verification.

Automated metrics offer objective but narrow signals. Code execution success, unit-test pass rates, citation accuracy, acceptance prediction, and traditional text-generation metrics such as BLEU [148], ROUGE [112], and BERTScore [255] are easy to compute, but they capture only fragments of research quality. For example, executable code can still be semantically wrong, accurate citations can still be used to support misleading claims, and fluent text can still lack contribution. Over-optimization on any one metric risks Goodhart’s law.

Execution-grounded evaluation verifies research outputs by running code, reproducing experiments, or checking claims against generated evidence. This paradigm is especially important for [§3](#), but it also affects Writing and Validation because manuscript claims should be traceable to executed experiments. PaperBench [192], for example, decomposes papers into individually gradable subtasks that can be checked through implementation and execution. Si *et al.* [185] further show that execution-guided search can improve discovery workflows by using empirical feedback rather than textual judgment alone.

Process- and trace-based evaluation assesses how a system reaches an output, not only the final artifact. This includes tool-use trajectories in deep research, reviewer-comment decomposition in rebuttal, revision fulfillment after author responses, and fidelity between paper content and dissemination materials. This paradigm is increasingly important because many lifecycle failures occur at handoffs: a system may retrieve the right paper but cite it incorrectly, run an experiment but summarize it inaccurately, or promise a rebuttal revision without implementing it.

7.2.3 Emerging Evaluation Paradigms

Several emerging paradigms are reshaping evaluation for AI-assisted research. The first is *execution-grounded evaluation*, where claims are checked against executable artifacts rather than judged only from text. This is essential for research coding, paper replication, and experimental analysis, where surface plausibility is insufficient. It also provides a path toward evaluating Writing: a manuscript claim should be traceable to a figure, table, log, or executed experiment.

The second is *adversarial evaluation*. As discussed in §6 (Peer Review), LLM-based reviewers and judges are vulnerable to prompt injection, lexical triggers, and covert content manipulation. Breaking the Reviewer [113] and related studies show that robustness to manipulation must be treated as an evaluation dimension rather than a peripheral security issue. This is particularly important for Validation, where a manipulated review or judge can affect acceptance decisions.

The third is *long-horizon evaluation*. Many benchmarks remain short-horizon, evaluating tasks that take minutes or hours rather than weeks or months. However, real research involves delayed feedback, failed attempts, changing hypotheses, and evolving evidence. METR’s analysis suggests that AI task horizons are rapidly increasing [131], while RE-Bench [221] provides open-ended ML R&D environments that begin to approximate longer research workflows. Still, current long-horizon benchmarks remain far shorter and cleaner than authentic research projects.

The fourth is *lifecycle-level evaluation*. Existing benchmarks usually evaluate one stage at a time, but many important failures occur between stages. A future lifecycle benchmark should test whether an idea remains valid after implementation, whether retrieved literature is faithfully represented in writing, whether experimental evidence supports manuscript claims, whether rebuttal commitments are fulfilled, and whether dissemination artifacts preserve claims and limitations. Such evaluation would better match the real risk profile of end-to-end research automation.

7.2.4 Evaluation Gaps

Despite rapid progress, several gaps remain unresolved. First, *novelty and significance are still difficult to define*. Expert judgments vary across reviewers, venues, and fields, and automated novelty scores can reward ideas that sound original but fail after execution. Second, *benchmark contamination and temporal validity are persistent concerns*. Many tasks are derived from public papers, code, or reviews that may appear in model training data. Temporal splits help, but they introduce changes in topic, difficulty, and community standards.

Third, *cross-system comparison remains difficult*. Systems are often evaluated with different base models, prompts, tools, datasets, compute budgets, and human-in-the-loop assumptions. This makes reported results hard to compare even when they target the same stage. Fourth, *cross-domain generalization remains under-tested*. Most benchmarks focus on machine learning and NLP, while chemistry, biology, materials science, physics, medicine, and social science require different evidence standards, experimental workflows, and domain-specific tools.

Fifth, *computational cost is itself an evaluation dimension*. Some research tasks, such as paper replication, long-horizon experimentation, and multimodal dissemination, require substantial token, compute, or tool-use budgets. A system that performs well under unlimited sampling may not be practically useful if its cost is prohibitive or its results are not reproducible under realistic constraints.

Finally, *no existing benchmark evaluates the complete research lifecycle with human-equivalent rigor*. PaperBench [192] makes important progress for replication, and process-aware benchmarks are emerging across literature review, coding, peer review, rebuttal, and Paper2X. However, no benchmark yet evaluates the full chain from ideation to dissemination while preserving traceability across artifacts. This lifecycle evaluation gap is central: without it, systems may appear strong within individual stages while failing to maintain scientific fidelity across the research process.

7.3 Cross-Cutting Insights

The preceding sections reveal a consistent pattern across the research lifecycle: AI systems are increasingly capable of producing research-like artifacts, but remain less reliable at verifying whether those artifacts

are novel, faithful, executable, and scientifically meaningful. We distill five cross-cutting insights from the stage-level analysis. These insights are not tied to a single tool or benchmark; rather, they describe recurring capability boundaries and deployment principles that appear across Creation, Writing, Validation, and Dissemination.

7.3.1 Artifact Generation Outpaces Scientific Verification

Across the lifecycle, AI systems are better at producing artifacts than at verifying their scientific validity. In [s1](#) (*Idea Generation*), generated ideas can appear novel and well-motivated, yet weaken after implementation. Si *et al.* [184] show that AI-generated ideas degrade more sharply after execution than human ideas, exposing a gap between apparent novelty and executable substance. In [s3](#) (*Coding and Experiments*), generated code may run successfully while implementing the wrong algorithm, with semantic failures forming a major source of error [71]. In [s4](#) (*Tables and Figures*), generated visual artifacts may look polished while misrepresenting data, notation, or information flow. In [s5](#) (*Paper Writing*), fluent prose can conceal weak reasoning or unsupported claims.

This gap also appears in Validation and Dissemination. In [s6](#) (*Peer Review*), automated reviews can be coherent and consistent while under-detecting decisive methodological flaws or assigning inflated scores [142, 267]. In [s7](#) (*Rebuttal and Revision*), generated responses may sound persuasive, but their value depends on whether promised evidence or manuscript changes are actually fulfilled [21]. In [s8](#) (*Dissemination*), posters, slides, videos, and social media summaries can simplify a paper in ways that overstate claims or omit limitations. The central lifecycle problem is therefore not artifact production alone, but artifact verification: each output must remain traceable to evidence, assumptions, and limitations.

7.3.2 Human-Governed Collaboration Remains the Most Reliable Deployment Mode

The strongest deployment pattern across stages is not full autonomy, but human-governed collaboration. In Writing, semi-automated systems are most credible when they assist planning, drafting, polishing, and citation support while researchers retain control over argumentation, interpretation, and final responsibility. In Peer Review, the strongest validated setting is not standalone AI review, but AI feedback on human reviews: the ICLR 2025 randomized study shows that LLM feedback improved review quality in 89% of cases without affecting acceptance rates [202]. In Rebuttal, author-aware systems are more appropriate than generic response generation because rebuttals must reflect the paper’s actual contributions and the authors’ intended revisions [163]. In Dissemination, AI-generated posters, slides, videos, and public summaries are best treated as editable drafts whose claims and emphasis remain under author control.

This pattern explains why direct automation is risky in high-stakes research settings. Research requires judgment under uncertainty: deciding whether an idea is worth pursuing, whether an experiment is sufficient, whether a critique is valid, whether a rebuttal promise is feasible, and whether a public-facing summary is faithful. These decisions are precisely where current systems remain fragile. AI assistance is therefore most useful when it expands researcher capacity while preserving human oversight over scientific claims, evidence interpretation, and accountability.

7.3.3 Capability Boundaries Emerge in Open-Ended Research Tasks

The sharpest capability boundaries appear when tasks become novel, underspecified, or long-horizon. Current systems perform strongly on structured tasks with clear feedback, such as standard software issue resolution, grammar correction, simple plotting, and format conversion. Performance drops when the task requires interpreting implicit assumptions, designing meaningful experiments, reproducing underspecified methods, or judging scientific contribution. This is most visible in research coding: while frontier systems perform well on familiar software benchmarks, performance falls sharply on novel research-code tasks, with reported ceilings around 37–39% on dedicated benchmarks [71, 224].

Similar boundaries recur elsewhere. Literature review systems retrieve and summarize individual papers increasingly well, but struggle with multi-paper relational reasoning and citation fidelity. Idea-generation systems produce plausible hypotheses but face persistent novelty–feasibility tradeoffs. Paper-writing systems generate fluent manuscripts but remain weaker at argumentative depth and reviewer anticipation. Peer-review

systems can approximate review style but remain vulnerable to leniency, bias, and manipulation. These failures share a common structure: the task cannot be solved by pattern matching alone, because success depends on implicit domain knowledge, causal reasoning, long-horizon feedback, and expert judgment.

7.3.4 Effective Systems Converge on Layered Architectures

Across phases, the most capable systems increasingly combine three layers: *exploration*, *execution*, and *verification*. The exploration layer searches over hypotheses, paper collections, code variants, response plans, or design alternatives. The execution layer interacts with tools: retrieval engines, code interpreters, experiment runners, plotting libraries, document editors, or presentation generators. The verification layer checks whether intermediate outputs are grounded, correct, and useful, through execution feedback, citation validation, critique, reviewer simulation, or human review.

This layered view explains why simple prompting is insufficient for research automation. Research tasks rarely require only one generation step; they require proposing alternatives, testing them, revising based on feedback, and preserving state across iterations. Search-based systems improve exploration, tool-integrated systems strengthen execution, and multi-agent systems can support specialization and critique. However, more agents do not automatically improve performance. Multi-agent systems appear most useful when the task can be decomposed into parallel or role-specialized subtasks; they can degrade on sequential reasoning tasks when coordination overhead and error propagation dominate [89]. Thus, the important design principle is not agent count, but whether the architecture matches the task structure and includes reliable verification.

7.3.5 AI Use Has Become a Governance Problem, Not a Detection Problem

AI assistance is already embedded in the research ecosystem. Corpus-level studies estimate detectable AI modification in a nontrivial fraction of scientific writing, including up to 17.5% of computer science abstracts [109] and 13.5% of biomedical abstracts [90]. Self-reported adoption is higher, with many researchers using AI for writing or review-related tasks [134]. At the same time, linguistic marker studies show that AI-associated words can surge after the introduction of LLMs, but such signals are unreliable for adjudicating individual cases and can change as users and models adapt.

The policy implication is that the community should move from detection-centered enforcement toward disclosure, attribution, and accountability. Detection tools can produce false positives, especially for formal or non-native academic prose, while watermarking remains dependent on provider cooperation and robustness to paraphrasing [159]. The more durable governance questions are therefore: What forms of AI assistance must be disclosed? Which uses are allowed during review? Who is responsible for AI-generated claims, citations, rebuttal commitments, or public summaries? How should venues audit high-risk uses without penalizing legitimate writing support? As AI becomes a routine part of research practice, governance must focus less on whether AI was used at all and more on whether its use preserved scientific integrity.

7.4 Open Challenges and Future Directions

Despite rapid progress across the research lifecycle, the preceding analysis shows that the main barriers to reliable AI-assisted research are not merely missing tools. The harder problems concern whether AI systems can preserve faithfulness across phase boundaries, evaluate scientific value, verify evidence, support responsible governance, generalize across domains, and preserve human expertise. We organize the remaining challenges around these six themes.

7.4.1 Faithfulness Across Phase Boundaries

Many of the most consequential failures occur not within a single stage, but when artifacts move from one phase to the next. An idea that appears promising in [S1](#) (*Idea Generation*) may weaken after implementation in [S3](#) (*Coding and Experiments*); retrieved evidence in [S2](#) (*Literature Review*) may be misrepresented in [S5](#) (*Paper Writing*); experimental results from [S3](#) (*Coding and Experiments*) may be summarized into claims that are stronger than the data support; reviewer concerns in [S6](#) (*Peer Review*) may lead to rebuttal promises in [S7](#) (*Rebuttal and Revision*) that are not fulfilled; and dissemination outputs in [S8](#) (*Dissemination*) may simplify the contribution beyond its evidence.

This phase-boundary problem is especially important for end-to-end systems. A lifecycle-scale system must not only generate artifacts, but preserve traceable links between them: hypotheses should connect to retrieved literature, code should connect to experiments, figures should connect to logs, manuscript claims should connect to evidence, rebuttal commitments should connect to revisions, and public-facing summaries should connect to the validated paper. Current systems rarely maintain this level of provenance across the full lifecycle. Future systems should therefore treat phase handoffs as explicit verification checkpoints rather than implicit transitions between modules.

7.4.2 Scientific Judgment and Novelty Assessment

Scientific judgment remains difficult to automate because research quality is not reducible to surface novelty, fluency, or benchmark score. In ideation, generated proposals can appear novel before execution but fail to remain feasible or impactful after implementation. Diversity is also a persistent concern: LLM-generated ideas may cluster in narrow regions of the idea space, limiting their ability to explore genuinely distinct research directions [79]. In literature review, systems increasingly retrieve and summarize individual papers well, but still struggle with multi-paper relational reasoning, methodological lineage, and cross-paper contradictions.

The deeper challenge is that novelty, significance, and contribution are socially and temporally situated. A good research idea depends on field-specific context, feasibility, timing, community standards, and the availability of evidence. Automated novelty scoring can therefore reward ideas that sound original while missing whether they are executable, important, or meaningfully different from prior work. Future progress will likely require evaluation methods that combine retrieval, temporal splits, expert judgment, execution feedback, and downstream impact analysis, rather than relying on LLM-as-Judge scores alone.

7.4.3 Verification, Reproducibility, and Accountability

Verification is the central unresolved problem for autonomous research systems. In coding and experiments, generated code may execute successfully while implementing the wrong algorithm, and automated experiment runners can produce outputs that appear quantitative without being scientifically meaningful. Paper replication remains particularly difficult: PaperBench [192] shows that current agents still fall far short of human performance on reproducing research results. This indicates that even verifying existing work is not yet solved, let alone generating new work that is independently reproducible.

Rebuttal and revision expose a parallel accountability problem. A rebuttal is scientifically meaningful only if its claims are supported and its commitments are fulfilled. The commitment–fulfillment gap observed in ICLR 2025 [21] shows that persuasive response text is insufficient: systems must track whether promised experiments, clarifications, and revisions are actually incorporated. Future AI research systems should therefore include explicit evidence ledgers, experiment provenance, versioned manuscript diffs, and revision-tracking mechanisms. The goal is not only to produce stronger artifacts, but to make every claim auditable.

7.4.4 Citation, Versioning, and Source Provenance

Citation verification is not solved by adding retrieval or web search. Scientific records are versioned: the same contribution may appear as an arXiv preprint, workshop paper, conference version, journal extension, or revised technical report, with changes to title, authors, venue, year, DOI, and sometimes content. Bibliographic databases may merge or separate these records differently, and prior work on arXiv–publisher citation consolidation shows that version merging is itself a nontrivial bibliometric problem [44].

This creates a challenge for AI-assisted literature review, writing, and dissemination. A generated manuscript may cite the correct idea but assemble metadata from inconsistent versions, or quote a claim from one version while citing another. Existing citation-audit tools and benchmarks target fabricated or unsupported references [151, 215], but end-to-end research agents also need *version-consistent citation assembly*: title, authors, venue, year, URL/DOI, and quoted claims should come from the same selected record, with provenance preserved. Future systems should therefore treat citation not as a formatting task, but as a versioned source-grounding problem.

7.4.5 Governance, Disclosure, and Research Integrity

AI use in research is no longer hypothetical. Writing assistance, review support, literature search, code generation, and dissemination drafting are already part of many researchers' workflows. This makes governance a central challenge. Detection-based enforcement is unreliable because AI text detectors can produce false positives, especially for formal academic writing, non-native prose, or heavily edited text. As discussed in §5 (*Paper Writing*) and §6 (*Peer Review*), the community is therefore shifting from trying to detect every instance of AI use toward requiring disclosure, attribution, and accountability.

The open question is how to define responsible AI use across stages. Assistance with grammar correction is different from generating experimental claims; drafting a rebuttal is different from promising new experiments; using AI to improve a review is different from delegating the review itself. Venues, publishers, and institutions need policies that distinguish low-risk assistance from high-risk substitution, specify what must be disclosed, and clarify who is accountable for AI-generated content. The central governance principle should be that authors remain responsible for claims, citations, experiments, rebuttal commitments, and public-facing summaries, regardless of which AI tools contributed to their production.

7.4.6 Cross-Domain Generalization and Infrastructure Access

Most current systems and benchmarks are concentrated in computer science, machine learning, and NLP. Extending AI-assisted research to chemistry, biology, medicine, materials science, physics, and social science requires more than retraining on domain papers. These fields differ in evidence standards, experimental infrastructure, safety constraints, data availability, and community norms. Systems such as Google AI Co-scientist [54], Biomni [72], Medical AI Scientist [222], and domain-specific laboratory agents point toward this direction, but broad cross-domain generalization remains unresolved.

Infrastructure access is part of the same challenge. Some domains require specialized instruments, wet-lab protocols, proprietary datasets, or expensive compute. If advanced AI research tools are available only to well-resourced laboratories or companies, they may amplify existing inequalities in scientific production. Future systems should therefore be evaluated not only by performance, but also by accessibility, reproducibility, and deployability under realistic resource constraints. Open-source tools, standardized interfaces, shared benchmarks, and transparent provenance mechanisms will be important for preventing research automation from becoming an infrastructure privilege.

7.4.7 Human Expertise and Cognitive Ownership

A final challenge concerns the long-term development of researchers themselves. Many AI tools automate the external products of research: summaries, code, plots, manuscripts, reviews, rebuttals, and slides. However, the cognitive value of research lies in forming hypotheses, understanding prior work, diagnosing failures, interpreting results, constructing arguments, and responding to critique. If AI tools bypass these processes too aggressively, they may increase short-term productivity while weakening the skills that define scientific expertise.

This concern is most visible in Writing, where AI assistance is already widely adopted, but it applies across the lifecycle. A junior researcher who delegates literature synthesis may not develop field judgment; one who delegates experiment planning may not learn what makes evidence decisive; one who delegates rebuttal may not learn how to reason from criticism. Tools such as Script&Shift [186] and DraftMarks [188] suggest a better design direction: AI should support source transformation, process transparency, and reflective revision rather than replacing the user's cognitive engagement. The practical principle is that AI should handle mechanical, repetitive, or scaffolded tasks, while humans retain ownership of judgment, interpretation, argumentation, and accountability.

7.4.8 Toward Reliable AI-Assisted Research

Taken together, these challenges suggest a shift in the goal of AI-assisted research. The near-term objective should not be fully autonomous science in which AI systems independently generate, validate, publish, and promote research without oversight. A more credible objective is reliable human-governed research automation:

systems that expand the scale and speed of research while preserving traceability, verification, expert judgment, and accountability.

Future progress will likely come from systems that integrate four design principles. First, they should maintain provenance across the full lifecycle, linking ideas, evidence, code, figures, claims, reviews, rebuttals, and dissemination artifacts. Second, they should use execution and retrieval grounding wherever possible, replacing purely textual self-judgment with verifiable signals. Third, they should include human checkpoints at phase boundaries, where errors are most likely to propagate. Fourth, they should make AI involvement transparent, so that readers, reviewers, and institutions can assess how a research artifact was produced. These principles define the path from artifact-generating systems toward trustworthy research collaborators.

8 Conclusion

This work presented an end-to-end analysis of AI-assisted academic research across the complete lifecycle. We organized the field into four epistemological phases: *Creation*, *Writing*, *Validation*, and *Dissemination*, and eight stages spanning ideation, literature review, coding and experiments, tables and figures, paper writing, peer review, rebuttal and revision, and Paper2X dissemination. This lifecycle framing connects tools that are often studied in isolation and exposes where current systems succeed, where they fail, and how errors propagate across stage boundaries.

The central finding is that AI systems are increasingly capable of producing research artifacts, but remain less reliable at verifying their scientific meaning. Across the lifecycle, plausible outputs can conceal deeper failures: ideas may weaken after execution, retrieved evidence may be misrepresented, executable code may implement the wrong algorithm, fluent manuscripts may lack argumentative depth, reviews may miss methodological flaws, rebuttals may promise unfulfilled revisions, and dissemination artifacts may overstate claims. The core bottleneck is therefore not generation alone, but maintaining novelty, faithfulness, reproducibility, and accountability across the research process.

The most credible path forward is human-governed AI-assisted research. AI should reduce mechanical friction in retrieval, drafting, coding, visualization, review support, and dissemination, while researchers retain ownership over judgment, interpretation, experimental design, argumentation, and final responsibility. Future systems should maintain provenance across artifacts, use retrieval and execution grounding wherever possible, support human checkpoints at phase boundaries, and make AI involvement transparent. If developed with these principles, AI can amplify human creativity and rigor; without them, it risks scaling the production of plausible but unreliable research artifacts.

Acknowledgments

We thank Josh Susskind for insightful discussions and careful proofreading of this manuscript.

We also thank the researchers and open-source contributors whose systems, benchmarks, datasets, and technical reports made this survey possible, as well as the broader community for ongoing discussions on responsible AI use, research integrity, peer review, and the future of scientific work.

Responsible Use and Limitations

This work is intended to inform responsible use of AI-assisted research tools, not to endorse replacing human scientific judgment with full automation. Current systems are most reliable when used to assist retrieval, drafting, coding, visualization, review support, and dissemination, while humans retain responsibility for novelty, interpretation, verification, authorship, and accountability. Because the field evolves rapidly, this paper should be read as a structured snapshot through our search cutoff, and AI-generated research outputs should be independently verified before scholarly use.

Appendix

A Auto-Research Tool Inventory

This appendix provides a comprehensive inventory of all surveyed works, organized by stage.

A.1 Phase 1: Creation

Table 3 Comprehensive inventory: S1 Idea Generation. †Evaluation information uncertain.


#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
<div style="border: 1px solid green; padding: 10px;">  <p>S1 Idea Generation Generating, refining, and evaluating novel research hypotheses. Techniques range from knowledge graph reasoning and retrieval-augmented generation to Multi-Agent collaboration for structured hypothesis formation.</p> </div>							
LLM Internal Knowledge-Based Generation							
1	Chain of Ideas	[102]	arXiv'24	↗	LLM Internal		Comparable to human quality; \$0.50/idea min cost
2	ResearchAgent	[10]	NAACL'25	↗	LLM Internal		Human + model eval; academic graph feedback
3	SciMON	[209]	ACL'24	↗	LLM Internal		Mitigates shallow novelty; iterative refinement
4	Idea Gen Agent	[183]	arXiv'24	↗	LLM Internal	-	100+ NLP researchers; LLM ideas higher novelty ($p < 0.05$)
5	IRIS	[47]	ACL'25	↗	LLM Internal		MCTS adaptive reasoning; human-in-the-loop platform
6	Spark	[168]	ICCC'25	↗	LLM Internal	-	Judge model trained on 600K OpenReview reviews
External Signal-Driven Generation							
7	MOOSE-Chem	[238]	ICLR'25	↗	External Signal	-	Rediscovered hypotheses from 51 high-impact papers
8	Nova	[69]	arXiv'24	↗	External Signal	-	3.4× more novel ideas; 2.5× more top-rated
9	SciAgents	[49]	arXiv'24	↗	External Signal		Multi-agent reasoning over knowledge graphs
10	SciPIP	[211]	arXiv'24	↗	External Signal		Multi-domain; paper-anchored idea generation
11	IdeaSynth	[152]	CHI'25	↗	External Signal	-	20-user study; more alternatives explored vs baseline
12	MOOSE-Chem2	[237]	NeurIPS'25	↗	External Signal	-	Fine-grained, experimentally actionable hypotheses
Multi-Agent Collaborative Generation							
13	Combi. Creativity	[57]	arXiv'24	↗	Multi-Agent	-	+7–10% similarity scores; cross-domain composition
14	Deep Ideation	[259]	arXiv'25	↗	Multi-Agent		+10.67% quality; surpasses conference acceptance levels
15	VirSci	[193]	ACL'25	↗	Multi-Agent		Outperforms single-agent on novelty [†]
16	Multi-Agent Dial.	[206]	SIGDIAL'25	↗	Multi-Agent	-	Optimal at 3 critique-revision rounds [†]
17	Artificial Hivemind	[79]	NeurIPS'25	↗	Multi-Agent	-	26K queries; diversity collapse across models
Novelty and Feasibility Assessment							
18	IdeaBench	[59]	KDD'25	↗	Evaluation	-	2,374 papers; 8 domains; novelty >0.6, feasibility <0.5
19	LiveIdeaBench	[162]	arXiv'24	↗	Evaluation	-	40+ models; 1,180 keywords; 22 scientific domains
20	AI Idea Bench 2025	[154]	arXiv'25	↗	Evaluation		3,495 papers; alignment + general reference eval
21	HeurekaBench	[147]	ICLR'26	↗	Evaluation		+22% with critic module; open-ended science tasks
22	ResearchBench	[120]	ACL'26	↗	Evaluation	-	12 disciplines; inspiration retrieval + ranking
23	HindSight	[78]	arXiv'26	↗	Evaluation	-	LLM novelty negatively correlated with impact ($\rho = -0.29$)
24	Rubric Rewards	[52]	arXiv'25	↗	LLM Internal	-	70% expert preference; RL with rubric self-grading
25	DeepInnovator	[39]	arXiv'26	↗	LLM Internal		80–94% win rates vs frontier; 14B model
26	FlowPIE	[210]	arXiv'26	↗	External Signal	-	Higher novelty, feasibility, diversity vs baselines

Table 4 Comprehensive inventory: S2 Literature Review. †Evaluation information uncertain.



S2 Literature Review





Retrieving, synthesizing, and organizing prior work into coherent narratives. Modern approaches span semantic retrieval, citation-graph traversal, and Deep Research agents that iteratively explore the literature.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
Literature Retrieval							
1	CiteME	[151]	arXiv'24	↗	Retrieval	-	Citation fidelity benchmark
2	LitLLM	[2]	arXiv'24	↗	Retrieval	-	LLM + academic database integration
3	LitSearch	[4]	arXiv'24	↗	Retrieval	🔄	Retrieval precision benchmark
4	PaperQA2	[189]	arXiv'24	↗	Retrieval	🔄	Matches/exceeds expert on 3 tasks; 70% contradiction validation
5	OpenResearcher	[265]	EMNLP'24	↗	Retrieval	-	RAG + graph traversal for literature exploration
6	PaSa	[62]	arXiv'25	↗	Retrieval	🔄	Agentic multi-step iterative retrieval
Survey & Related Work Generation							
7	ChatPaper	[124]	GitHub'23	↗	Generation	🔄	19K+ GitHub stars; arXiv summarization tool
8	PaperQA	[128]	arXiv'23	↗	Generation	🔄	8K+ GitHub stars; RAG for scientific Q&A
9	AutoSurvey	[214]	arXiv'24	↗	Generation	🔄	First end-to-end LLM survey drafting system
10	GPT Researcher	[38]	GitHub'24	↗	Generation	🔄	26K+ GitHub stars; comprehensive report generation
11	LLMs for Lit. Review	[200]	arXiv'24	↗	Generation	-	Hallucination analysis; models still generate errors [†]
12	STORM	[178]	arXiv'24	↗	Generation	🔄	Multi-perspective question-asking for outlines
13	Agentic AutoSurvey	[119]	arXiv'25	↗	Generation	-	Multi-agent role decomposition [†]
14	Citegeist	[11]	arXiv'25	↗	Generation	-	Dynamic RAG pipeline on arXiv corpus
15	IterSurvey	[251]	arXiv'25	↗	Generation	🔄	Iterative outline planning with stability checks
16	LiRA	[51]	arXiv'25	↗	Generation	-	Multi-agent retrieval + verification + narrative
17	SurveyForge	[230]	arXiv'25	↗	Generation	🔄	Outperforms AutoSurvey on outline quality [†]
18	SurveyG	[137]	arXiv'25	↗	Generation	-	Three-layer citation graph (Foundation/Dev/Frontier)
19	SurveyX	[110]	arXiv'25	↗	Generation	-	+0.259 content quality improvement; near expert level
20	InteractiveSurvey	[219]	arXiv'25	↗	Generation	🔄	User-customizable reference categorization + outlines
21	CiteLLM	[65]	arXiv'26	↗	Generation	-	Hallucination-free via trusted repository routing
Deep Research Agents							
22	ASReview	[207]	Nature MI'21	↗	Deep Research	🔄	Active learning; up to 95% effort reduction
23	CHIME	[68]	arXiv'24	↗	Deep Research	-	Hierarchical organization of scientific studies
24	DeepResearch-Agent	[190]	GitHub'25	↗	Deep Research	🔄	Hierarchical multi-agent; planner + sub-agents
25	DeerFlow	[18]	GitHub'25	↗	Deep Research	🔄	Sub-agents with shared memory; sandboxed execution
26	OpenScholar	[9]	Nature'26	↗	Deep Research	-	45M papers; +6.1% over GPT-4o, +5.5% over PaperQA2
27	AutoAgent	[198]	arXiv'25	↗	Deep Research	-	Universal LLM compatibility; GAIA benchmark
28	Tongyi DeepResearch	[7]	GitHub'25	↗	Deep Research	🔄	30.5B params (3.3B activated); SOTA on Deep Research
29	O-Researcher	[240]	arXiv'26	↗	Deep Research	-	Multi-agent distillation + agentic RL
30	OpenResearcher	[108]	arXiv'26	↗	Deep Research	🔄	54.8% BrowseComp-Plus; 97K+ trajectories
Retrieval and Synthesis Quality Assessment							
31	DeepScholar-Bench	[149]	arXiv'25	↗	Evaluation	🔄	Coverage, coherence, factual accuracy benchmark
32	ReportBench	[103]	arXiv'25	↗	Evaluation	🔄	100-prompt benchmark from 678 filtered survey papers
33	IDRBench	[40]	arXiv'26	↗	Evaluation	-	100 tasks; interactive Deep Research evaluation
34	ScholarGym	[179]	arXiv'26	↗	Evaluation	-	2,536 queries; query planning + tool invocation
35	SciNetBench	[176]	arXiv'26	↗	Evaluation	-	18M papers; relation-aware retrieval <20%

Table 5 Comprehensive inventory: S3 Coding & Experiments. †Evaluation information uncertain.


#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
<div style="display: flex; align-items: center;">  <div> <p>S3 Coding & Experiments</p> <p>Translating ideas into executable code, running experiments at scale, and analyzing results. This stage spans code generation, Paper-to-Code translation, autonomous experiment orchestration, and result interpretation.</p> </div> </div>							
Code Generation							
1	SWE-bench	[82]	ICLR'24	↗	Code Gen.		2,294 real GitHub issues; Verified split (500 problems)
2	SWE-agent	[231]	arXiv'24	↗	Code Gen.		Agent-computer interface paradigm for coding
3	OpenHands	[212]	ICLR'25	↗	Code Gen.		Open platform for generalist coding agents
4	SWE-bench Pro	[37]	arXiv'25	↗	Code Gen.	-	1,865 enterprise problems; best score 23%
5	SWE-EVO	[201]	arXiv'25	↗	Code Gen.	-	Software evolution benchmark; best score 25%
Paper-to-Code							
6	FunSearch	[161]	Nature'24	↗	Paper-to-Code		New cap-set solutions; evolutionary program search
7	SciCode	[203]	arXiv'24	↗	Paper-to-Code		Research-level coding across math, physics, chemistry
8	PaperBench	[192]	arXiv'25	↗	Paper-to-Code		20 ICML'24 papers; 8,316 gradable subtasks
9	PaperCoder	[174]	arXiv'25	↗	Paper-to-Code		3-stage multi-agent; ML papers to code repos
10	ResearchCodeBench	[71]	arXiv'25	↗	Paper-to-Code	-	212 novel ML tasks; best 37.3% (Gemini-2.5-Pro)
11	SciReplicate-Bench	[224]	arXiv'25	↗	Paper-to-Code		100 tasks from 36 NLP papers; 39% ceiling
Experiment Execution & Orchestration							
12	BioPlanner	[141]	arXiv'23	↗	Execution		Biological protocol planning evaluation
13	CRISPR-GPT	[156]	arXiv'24	↗	Execution	-	Gene-editing experiment design assistance
14	DS-Agent	[58]	arXiv'24	↗	Execution		End-to-end data science workflow automation
15	MLE-Bench	[20]	arXiv'24	↗	Execution	-	75 Kaggle competitions benchmark
16	MLAgentBench	[73]	arXiv'24	↗	Execution		13 ML experimentation tasks benchmark
17	MLR-Copilot	[104]	arXiv'24	↗	Execution	-	IdeaAgent + ExperimentAgent dual-agent pipeline
18	AIDE	[81]	arXiv'25	↗	Execution	-	SOTA on MLE-Bench + RE-Bench; tree search in code space
19	AlphaEvolve	[140]	arXiv'25	↗	Execution	-	LLM-generated mutations + automated evaluators
20	AutoReproduce	[260]	arXiv'25	↗	Execution		Paper lineage algorithm for experiment reproduction
21	CURIE	[91]	arXiv'25	↗	Execution		Rigorous automated experimentation framework
22	MLGym	[135]	arXiv'25	↗	Execution	-	AI research agent gym benchmark
23	MLR-Bench	[24]	arXiv'25	↗	Execution	-	201 tasks (NeurIPS/ICLR/ICML); 80% fabrication rate
24	Execution-Grounded	[185]	arXiv'26	↗	Execution	-	69.4% vs 48.0% GRPO; parallel GPU search
25	Learn to Discover	[246]	arXiv'26	↗	Execution	-	Test-time training + RL; math, GPU kernel, biology
26	SciNav	[254]	arXiv'26	↗	Execution	-	Pairwise tree-search branch selection
27	FrontierScience	[208]	arXiv'26	↗	Execution	-	Expert-level tasks; Olympiad + PhD difficulty
Code Correctness and Reproducibility Assessment							
28	DiscoveryBench	[130]	arXiv'24	↗	Analysis		Data-driven insight extraction benchmark
29	DiscoveryWorld	[76]	arXiv'24	↗	Analysis		120 tasks; 8 topics; 3 difficulty levels
30	InfiAgent-DABench	[70]	arXiv'24	↗	Analysis	-	End-to-end data analysis workflow benchmark
31	ScienceAgentBench	[32]	arXiv'24	↗	Analysis	-	Rigorous data-driven scientific discovery assessment
32	LAB-Bench	[96]	arXiv'24	↗	Execution		Multi-domain biology research task benchmark
33	KernelBench	[143]	arXiv'25	↗	Execution		GPU kernel generation benchmark
34	TritonBench	[101]	arXiv'25	↗	Execution		Triton operator generation benchmark
35	AstaBench	[16]	arXiv'25	↗	Execution		2,400+ problems; multi-domain scientific research
36	ResearchClawBench	[225]	arXiv'25	↗	Execution		Scientist-aligned workflow benchmark
37	EXP-Bench	[92]	ICLR'26	↗	Execution		461 tasks from 51 AI papers
38	PostTrainBench	[158]	arXiv'26	↗	Execution		LLM post-training automation benchmark

Table 6 Comprehensive inventory: S4 Tables & Figures. †Evaluation information uncertain.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
<div style="display: flex; align-items: center;">  <div> <p>S4 Tables & Figures</p> <p>Creating Method Diagrams, result plots, comparison tables, and LaTeX formulas. Scientific visualization transforms raw experimental outputs into publication-quality charts, illustrations, and structured tables.</p> </div> </div>							
Scientific Figure Generation							
1	ChartGPT	[204]	arXiv'23	↗	Data Viz	-	6-step reasoning for chart generation
2	MatPlotAgent	[236]	arXiv'24	↗	Data Viz	-	+12.3 over GPT-4 base; VLM visual feedback†
3	CoDA	[31]	arXiv'25	↗	Data Viz	-	+41.5% over baselines; multi-agent collaboration
4	PlotGen	[53]	arXiv'25	↗	Data Viz	-	4–6% improvement over baselines†
5	VIS-Shepherd	[145]	arXiv'25	↗	Figure Editing	-	Constructive critique feedback framework
6	DiagramAgent	[218]	CVPR'25	↗	Data Viz	-	4 specialized agents; 8 diagram categories
7	StarVector	[160]	CVPR'25	↗	Method Diagrams	-	Scalable SVG generation from descriptions†
8	VisCoder	[139]	EMNLP'25	↗	Data Viz	-	VisCode-200K dataset; 90%+ execution pass rate†
9	AI-Generated Figures	[22]	arXiv'26	↗	Policy	-	Publisher policy survey (Nature, Science, etc.)
10	AutoFigure-Edit	[114]	arXiv'26	↗	Method Diagrams		Editable text-to-SVG scientific illustrations†
11	AutoFigure	[270]	ICLR'26	↗	Method Diagrams		FigureBench (3,300 pairs); publication-ready illust.
12	PaperBanana	[268]	arXiv'26	↗	Method Diagrams	-	292 test cases; outperforms baselines†
13	SAIL	[172]	arXiv'26	↗	Figure Editing	-	Domain logic / code syntax separation
Table Understanding & Generation							
14	ArxivDIGESTables	[136]	EMNLP'24	↗	Table Gen.	-	Literature comparison table synthesis
15	Chain-of-Table	[216]	ICLR'24	↗	Table Reasoning	-	Multi-step table reasoning chains
16	ShowTable	[121]	CVPR'26	↗	Table Viz	-	Collaborative reflection and refinement†
17	Table2LaTeX-RL	[115]	arXiv'25	↗	Table Conversion	-	Image-to-LaTeX via reinforced multimodal LM
Mathematical Formulas & TikZ							
18	AutomaTikZ	[12]	ICLR'24	↗	TikZ	-	DaTikZ: first large-scale dataset (120K drawings)
19	DeTikZify	[13]	NeurIPS'24	↗	TikZ	-	360K TikZ graphics; MCTS iterative refinement
20	TikZilla	[56]	arXiv'26	↗	TikZ	-	3B/8B matches GPT-5; SFT+RL on expanded DaTikZ
Visual Fidelity and Scientific Accuracy Assessment							
21	PlotCraft	[253]	arXiv'25	↗	Benchmark	-	1K-task benchmark; 48 chart types
22	TeXpert	[85]	SDP'25	↗	Benchmark	-	3-level difficulty; 78.8%/58.7%/17.5%†
23	AbGen	[261]	ACL'25	↗	Benchmark	-	1,500 ablation studies; 807 NLP papers
24	SciFig	[74]	arXiv'26	↗	Benchmark	-	Rubric-based evaluation; 2K+ pipeline figures
25	SciFlow-Bench	[256]	arXiv'26	↗	Benchmark	-	500 framework figures; inverse-parsing evaluation
26	FigureBench	[270]	ICLR'26	↗	Benchmark		3,300 text-figure pairs; publication-ready eval

A.2 Phase 2: Writing

Table 7 Comprehensive inventory: S5 Paper Writing. †Evaluation information uncertain.


#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
<div style="display: flex; align-items: center;">  <div> <p>S5 Paper Writing</p> <p>Drafting, editing, and polishing academic manuscripts. AI assistance ranges from semi-automated grammar and citation tools to fully automated paper generation: the most commercially mature yet ethically contested stage.</p> </div> </div>							
Semi-Automated Writing Assistance							
1	CoAuthor	[97]	arXiv'22	↗	Collaborative	-	Human-AI collaborative writing workflows
2	Script&Shift	[186]	CHI'25	↗	Source Transform	-	CHI Honorable Mention; preserves cognitive engagement
3	AI Writing Study	[187]	AIED'25	↗	Empirical Study	-	90-student RCT; purposeful AI fosters writing
4	OpenDraft	[36]	-	↗	Full Draft Gen.		19 agents; 20K+ words in 10 min; verified citations
5	DraftMarks	[188]	arXiv'25	↗	Transparency	-	Skeuomorphic visual traces for AI process transparency
6	PaperDebugger	[67]	arXiv'25	↗	In-Editor Assist		Multi-agent Overleaf plugin (Reviewer+Enhancer+Scorer)
7	ScholarCopilot	[215]	arXiv'25	↗	Citation Assist	-	40.1% top-1 citation accuracy (vs 15.0% E5-Mistral)
8	XtraGPT	[25]	arXiv'25	↗	Post-Writing	-	1.5B-14B models; 7K papers; 140K revision pairs
9	LimAgents	[6]	arXiv'26	↗	Limitations Gen.	-	OpenReview comments + citation network integration
Fully Automated Paper Generation							
10	CycleResearcher	[220]	ICLR'25	↗	E2E Gen.	-	5.36 ICLR scale (vs 5.24 preprint, 5.69 accepted)
11	Agent Laboratory	[171]	EMNLP'25	↗	E2E Gen.	-	\$2-13/paper; 84% cost reduction; 3.5-4.0 score
12	FutureGen	[5]	arXiv'25	↗	Section Gen.	-	RAG-based Future Work section generation
13	AI Scientist	[122]	Nature'26	↗	E2E Gen.		\$15/paper; end-to-end across 3 ML subfields
14	APRES	[258]	arXiv'26	↗	Rubric Revision	-	79% expert preference; citation-predictive rubrics
Societal Analysis							
15	AI Writing Adoption	[61]	Nature'26	↗	Measurement	-	41.3M papers; AI expands impact but contracts focus
16	Nature AI Survey	[134]	Nature'26	↗	Survey	-	57% of researchers use AI in peer review
Writing Quality and AI Detection Assessment							
17	Mapping LLM Use	[109]	arXiv'24	↗	Detection	-	Up to 17.5% of CS papers AI-modified
18	CycleReviewer	[220]	ICLR'25	↗	AI Judge	-	26.89% MAE reduction vs individual human reviewers
19	Stanford Agentic	[80]	Web'25	↗	AI Judge	-	$\rho = 0.42$ vs human $\rho = 0.41$; matches consistency
20	SciIG	[46]	arXiv'25	↗	Writing Bench	-	NAACL/ICLR 2025 introduction writing benchmark [†]
21	Watermarking	[159]	arXiv'25	↗	Detection	-	Near-zero false-positive rate under controlled conditions
22	PaperWritingBench	[191]	arXiv'26	↗	Benchmark	-	200 reverse-engineered top-tier conference papers

A.3 Phase 3: Validation

Table 8 Comprehensive inventory: S6 Peer Review. †Evaluation information uncertain.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
Automated Review Generation							
1	ChatReviewer	[138]	GitHub'23	↗	Review Gen.		ChatGPT-based strengths/weaknesses analysis
2	AI-Peer-Review	[150]	GitHub'24	↗	Review Gen.		Multi-LLM independent reviews + meta-review synthesis
3	MARG	[35]	arXiv'24	↗	Review Gen.	-	3.7 good comments/paper (2.2× over baseline)
4	Reviewer2	[45]	arXiv'24	↗	Review Gen.	-	Two-stage prompt-based review aspect modeling
5	ReviewRL	[247]	EMNLP'25	↗	Review Gen.	-	RL + RAG; factually grounded reviews
6	DeepReviewer	[269]	arXiv'25	↗	Review Gen.	-	88.21% win rate vs GPT-o1; 64% accept/reject acc.
7	OpenReviewer	[75]	NAACL'25	↗	Review Gen.	-	Llama-8B fine-tuned on 79K expert reviews
8	REMOR	[196]	arXiv'25	↗	Review Gen.	-	Multi-objective RL review optimization
9	ScholarPeer	[55]	arXiv'26	↗	Review Gen.	-	Context-aware multi-agent; literature verification
Meta-Review & Reviewer Matching							
10	AgentReview	[83]	EMNLP'24	↗	Meta-Review	-	Full review lifecycle simulation; social/authority bias
11	Meta-Review LLMs	[66]	NAACL'25	↗	Meta-Review	-	40 ICLR papers; GPT-3.5/LLaMA-2/PaLM-2 compared
12	RATE	[118]	arXiv'26	↗	Matching	-	Expertise-based matching via profile distillation
Adversarial Attacks & Bias Analysis							
13	Raina <i>et al.</i>	[157]	EMNLP'24	↗	Adversarial	-	Benign adjectives as universal adversarial triggers
14	AI Review Lottery	[164]	arXiv'24	↗	Bias Analysis	-	15.8% ICLR reviews AI-assisted; +4.9pp borderline
15	Ye <i>et al.</i>	[242]	arXiv'24	↗	Adversarial	-	Scores inflated to ~8; 5% manipulation flips 12%
16	Breaking the Reviewer	[113]	arXiv'25	↗	Adversarial	-	Systematic adversarial robustness evaluation
17	LLM Reviewer Bias	[267]	arXiv'25	↗	Bias Analysis	-	1,441 papers; 95.8% rejected misclassified as acceptable
18	Prompt Injection	[88]	arXiv'25	↗	Adversarial	-	White-text injection; up to 100% acceptance scores
19	Sahoo <i>et al.</i>	[166]	arXiv'25	↗	Adversarial	-	+13.95 on Mistral; 13 LLMs; 15 attack strategies
20	Zhou <i>et al.</i>	[266]	arXiv'25	↗	Adversarial	-	+1.24 to +2.80 from hype prose; 10.00 under iteration
Detection & Policy							
21	AI Detection	[244]	arXiv'25	↗	Detection	-	788,984 AI-written reviews; 18 detection algorithms
22	AI Use Rejects	[50]	Nature'26	↗	Policy	-	497 papers rejected (~2% of submissions)
23	Nature AI Survey	[134]	Nature'26	↗	Survey	-	1,600 academics; 57% use AI in peer review
24	Policy Enforcement	[165]	arXiv'26	↗	Policy	-	All 5 SOTA detectors misclassify LLM-polished reviews
25	Reviewer Feedback	[29]	CHI'26	↗	Empirical	-	ICLR 2025 live process; reviewer engagement study†
Review Consistency and Bias Assessment							
26	Review Survey	[272]	IF'25	↗	Survey	-	Comprehensive taxonomy of review methods
27	Stanford Agentic	[80]	Web'25	↗	Quality	-	$\rho = 0.42$ vs human $\rho = 0.41$; matches consistency
28	ClaimCheck	[142]	EMNLP'25	↗	Quality	-	LLM critique grounding; gaps in factual basis
29	ReViewGraph	[105]	AAAI'26	↗	Quality	-	+15.73% avg improvement via heterogeneous graph
30	ReviewAgents	[43]	arXiv'25	↗	Quality	-	37,403 papers; 142,324 reviews; Review-CoT dataset
31	ICLR 2025 Study	[202]	NMI'26	↗	Quality	-	22,467 reviews; 89% quality improved; no acceptance effect

Table 9 Comprehensive inventory: S7 Rebuttal & Revision. †Evaluation information uncertain.




S7 Rebuttal & Revision
 Analyzing reviewer comments and generating evidence-grounded responses. The newest frontier in AI auto-research, rebuttal automation is decisive for roughly one in five borderline submissions at major venues.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
Reviewer Comment Analysis							
1	ReviewMT	[197]	arXiv'24	↗	Analysis	-	26,841 papers; 92,017 reviews; multi-turn dialogue
2	ICLR Rebuttal Study	[86]	arXiv'25	↗	Analysis	-	ICLR 2024–2025; score transition analysis
Automated Rebuttal Generation							
3	ReviewerToo	[167]	arXiv'25	↗	Modular Pipeline	-	81.8% accept/reject accuracy (vs 83.9% human)
4	RebuttalAgent	[63]	ICLR'26	↗	Rebuttal Gen.		18.3% avg improvement; ToM-grounded
5	Author-in-the-Loop	[163]	ACL'26	↗	Author-Aware	-	Integrates author expertise and intent
6	DRPG	[60]	arXiv'26	↗	Rebuttal Gen.		98%+ planning accuracy; surpasses avg human quality
7	Paper2Rebuttal	[129]	arXiv'26	↗	Rebuttal Gen.	-	Evidence-centric rebuttal planning
Rebuttal Effectiveness Assessment							
8	Re ²	[250]	arXiv'25	↗	Dataset	-	19,926 submissions; 70,668 reviews; 53,818 rebuttals
9	Commitment Checklist	[21]	arXiv'26	↗	Benchmark	-	11.8 commitments/paper; ~25% unfulfilled
10	Re ³ Align	[163]	ACL'26	↗	Dataset	-	First large-scale aligned review–response–revision triplets

A.4 Phase 4: Dissemination

Table 10 Comprehensive inventory: S8 Dissemination (Paper2X). †Evaluation information uncertain.



S8 Dissemination / Paper2X

Converting papers into posters, slides, videos, websites, and social media content. Each output format targets a different audience and demands its own design logic, AI tool chain, and trust considerations.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
Paper2Poster							
1	P2P	[194]	ICLR'26	🔗	Paper2Poster	-	P2PInstruct 30K+ examples; 3 specialized agents
2	Paper2Poster	[146]	NeurIPS'25	🔗	Paper2Poster	🔄	\$0.005/poster; 87% fewer tokens vs GPT-4o
3	PosterForest	[33]	arXiv'25	🔗	Paper2Poster	-	Hierarchical multi-agent collaboration†
4	PosterGen	[257]	arXiv'25	🔗	Paper2Poster	-	Aesthetic-aware multi-agent generation†
5	APEX	[180]	arXiv'26	🔗	Paper2Poster	🔄	First agentic interactive poster editing
6	PosterOmni	[27]	arXiv'26	🔗	Paper2Poster	-	6 unified poster tasks; outperforms open-source
Paper2Slides							
7	DOC2PPT	[41]	AAAI'22	🔗	Paper2Slides	-	5,873 paired document–slide decks
8	PPTAgent	[262]	EMNLP'25	🔗	Paper2Slides	🔄	PPTEval benchmark; 10,448 curated presentations
9	AutoPresent	[48]	CVPR'25	🔗	Paper2Slides	-	8B Llama model; SlidesBench (7K train, 585 test)
10	Paper2Slides	[64]	GitHub'25	🔗	Paper2Slides	🔄	4-stage RAG pipeline; one-click conversion
11	Auto-Slides	[235]	arXiv'25	🔗	Paper2Slides	-	Multi-agent Beamer generation; interactive editing†
12	PASS	[3]	arXiv'25	🔗	Paper2Slides	-	First combined slides + AI audio delivery†
13	SlideGen	[111]	arXiv'25	🔗	Paper2Slides	-	Multi-agent VLM coordination; editable PPTX output
14	Talk to Your Slides	[84]	arXiv'25	🔗	Paper2Slides	-	+34% instruction fidelity; 87% lower cost vs GUI
15	SlideTailor	[248]	AAAI'26	🔗	Paper2Slides	🔄	User-preference conditioned; chain-of-speech
16	DeepPresenter	[263]	arXiv'26	🔗	Paper2Slides	🔄	9B model competitive with frontier at lower cost
17	Office Raccoon	[173]	Web'26	🔗	Paper2Slides	-	Page-level editing; template/brand-guideline learning†
Paper2Video							
18	Preacher	[116]	ICCV'25	🔗	Paper2Video	🔄	Top-down decomposition; 5 research fields
19	Paper2Video	[271]	arXiv'25	🔗	Paper2Video	🔄	101 paper–video pairs; +10% PresentQuiz accuracy
20	PresentAgent	[181]	EMNLP'25	🔗	Paper2Video	🔄	PresentEval benchmark; approaches human-level
Paper2Web & Social Media							
21	Paper2Web	[30]	arXiv'25	🔗	Paper2Web	🔄	10,716 papers; multimedia-rich academic home-pages
Fidelity and Adoption Assessment							
22	PPTEval	[262]	EMNLP'25	🔗	Benchmark	🔄	Content, design, coherence; 10,448 presentations
23	PresentQuiz	[271]	arXiv'25	🔗	Benchmark	🔄	101 paper–video pairs; +10% over human on comprehension
24	PresentEval	[181]	EMNLP'25	🔗	Benchmark	🔄	End-to-end narrated video quality; near human-level

A.5 Cross-Phase: End-to-End Systems

Table 11 Comprehensive inventory: End-to-End and Cross-Phase Systems. Systems that span multiple stages of the research lifecycle. †Evaluation information uncertain.

#	Method	Ref	Venue	Link	Category	GitHub	Evaluation
Fully Automated Research Systems							
1	AI Scientist	[122]	arXiv’24		E2E Pipeline		Pioneered E2E at \$15/paper; ICLR-scale review
2	Agent Laboratory	[171]	EMNLP’25		E2E Pipeline	-	\$2–13/paper; 3.5–4.0 NeurIPS scale
3	AI-Researcher	[199]	arXiv’25		E2E Pipeline	-	Scientist-Bench; approaches human-level quality
4	CycleResearcher	[220]	ICLR’25		E2E Pipeline	-	5.36 ICLR scale; cyclic write–review
5	AI Scientist v2	[229]	arXiv’25		E2E + Tree Search		ICLR 2025 ICBINB workshop; score 6.33
6	Kosmos	[133]	arXiv’25		E2E Pipeline	-	79.4% claim accuracy; 7 discoveries; 4 domains
7	Dolphin	[245]	ACL’25		E2E Pipeline	-	Closed-loop auto-research pipeline
8	CodeScientist	[77]	ACL’25		E2E Pipeline		Hypothesis to verification; closed-loop
9	InternAgent	[249]	arXiv’25		E2E Pipeline		Closed-loop hypothesis to verification
10	freephdlabor	[100]	arXiv’25		Multi-Agent	-	Personalized research group; continual automation
11	SciMaster	[19]	arXiv’25		Multi-Agent		General-purpose scientific AI agents
12	ARIS	[233]	GitHub’26		Skill Library		31 skills; score 5.0→7.5; 20+ GPU experiments
13	EvoScientist	[127]	arXiv’26		Multi-Agent		6 papers accepted at ICAIS’25
14	UniScientist	[99]	Web’26		Multi-Agent	-	30B open-source; beats Claude Opus 4.5
15	ASI-Evolve	[227]	GitHub’26		Multi-Agent		+0.97 DeltaNet; +18 MMLU; +12.5 GRPO
16	AiScientist-LH	[23]	arXiv’26		E2E + Hierarchical	-	Long-horizon ML engineering; File-as-Bus
17	AutoSOTA	[107]	arXiv’26		E2E Pipeline		Paper-to-code-to-SOTA optimization
18	CORAL	[155]	arXiv’26		Multi-Agent		Multi-agent evolution; SOTA on 10 tasks
19	FARS	[8]	Web’26		Multi-Agent	-	100 papers in 228h; avg 5.05 ICLR scale
20	AutoResearchClaw	[117]	GitHub’26		Multi-Agent		fully autonomous 23-stage pipeline
Domain-Specific Systems							
21	Coscientist	[15]	Nature’23		Chemistry	-	Autonomous chemistry; LLM-driven tool use
22	AlphaFold 3	[1]	Nature’24		Biology	-	Biomolecular structure prediction
23	ChemCrow	[17]	NMI’24		Chemistry	-	Chemistry tool orchestration
24	Medical AI Scientist	[222]	arXiv’26		Medicine	-	Clinical research automation
Evolutionary & Self-Improving Systems							
25	ShinkaEvolve	[95]	arXiv’25		Evolutionary		Open-ended sample-efficient program evolution
26	Darwin Godel Machine	[252]	arXiv’25		Self-Improving		Open-ended evolution of self-improving agents
Research Platforms & Infrastructure							
27	R&D-Agent	[234]	arXiv’25		Infrastructure		Researcher+Developer dual-agent; MLE-Bench top
28	autoresearch	[87]	GitHub’25		Infrastructure		~12 exp/hour overnight
29	Google AI Co-scientist	[54]	arXiv’25		Platform	-	Multi-agent hypothesis gen + validation
30	ResearchTown	[243]	ICML’25		Multi-Agent		Simulates research community with LLM agents
31	AgentRxiv	[170]	arXiv’25		Multi-Agent	-	11.4% improvement on MATH-500
32	LabClaw	[223]	Web’26		Skill Library		206 biomedical skills; always-on autonomous lab agent
33	PiFlow	[153]	arXiv’25		Multi-Agent	-	Principle-aware scientific discovery

B Survey Coverage Comparison & Taxonomy Analysis

Table 12 compares our coverage with five closely related concurrent efforts. The goal is not to rank surveys, but to clarify how our lifecycle framework differs in scope and organization.

Our eight-stage framework subsumes several prior taxonomies while making two distinctions explicit: AI auto-research should be analyzed across the complete lifecycle, and stages should be grouped by epistemological function rather than only by task name or autonomy level.

- **AI4Research** [26] defines five task categories: Comprehension, Survey, Discovery, Writing, and Review. These overlap with our **S1–S3**, **S5**, and **S6**. Our framework newly elevates **S4** (Tables & Figures), **S7** (Rebuttal & Revision), and **S8** (Dissemination) as independent lifecycle stages.
- **From Automation to Autonomy** [264] organizes systems by autonomy level, from tool-like assistance to scientist-level automation. This axis is complementary: each of our stages can be instantiated at different autonomy levels, while our framework specifies *where* in the research lifecycle the system operates.
- **LLM4SR** [126] proposes a four-part view centered on hypothesis, experiment, writing, and review. This structure is close to ours, but does not separately model rebuttal and revision as a feedback stage. Our Validation phase separates **S6** from **S7**, making the review–response loop explicit.
- **Automated Scholarly Paper Review** [272] provide in-depth coverage of review generation, quality assessment, and reviewer–paper matching. They are complementary to our work: they focus on **S6**, while our framework places peer review within the broader lifecycle.
- **AI Scientist Survey** [205] focus on autonomous or semi-autonomous scientific discovery, overlapping mainly with **S1–S3** and **S5**. Our framework extends this view by also covering scientific visualization, peer validation, rebuttal, and dissemination.

These comparisons show that prior taxonomies often list research tasks sequentially, while leaving functional distinctions and feedback loops implicit. Our four-phase framework makes these dependencies explicit. For example, **S6** (*Peer Review*) and **S7** (*Rebuttal & Revision*) do not simply follow paper writing as isolated downstream steps; they can redirect the workflow back to **S3** for additional experiments, **S4** for revised figures or tables, and **S5** for manuscript restructuring. Similarly, dissemination artifacts in **S8** may expose ambiguities in the original framing, requiring revisions to claims, explanations, or visual evidence.

These cross-stage dependencies are central to real research practice and are especially important for AI-assisted workflows, where errors can propagate from generated ideas to experiments, from experiments to claims, and from claims to public-facing summaries. By organizing the field into *Creation*, *Writing*, *Validation*, and *Dissemination*, our framework highlights not only which stages are covered by existing systems, but also where evidence, claims, critique, and communication must remain aligned.

Table 12 Comparison of survey coverage across our four-phase research lifecycle framework. ✓ = in-depth coverage, ◦ = partial coverage, — = not covered, ★ new = newly introduced stage.

Stage	Ours	LLM4SR ₂₀₀₁	PR Survey ₂₀₀₁	Auto→Auton ₂₀₀₅	AI4Research ₂₀₀₇	AI Scientist ₂₀₁₀
Phase 1: Creation						
S1 : Idea Generation novelty, feasibility, multi-agent	✓	✓	—	✓	✓	✓
S2 : Literature Review retrieval, survey gen, deep research	✓	—	—	✓	✓	✓
S3 : Coding & Experiments paper-to-code, execution, analysis	✓	✓	—	✓	✓	✓
S4 : Figures & Tables diagrams, plots, formulas	✓	—	—	◦	◦	—
Phase 2: Writing						
S5 : Paper Writing semi-auto, full-auto, detection	✓	✓	—	◦	✓	✓
Phase 3: Validation						
S6 : Peer Review auto-review, matching, quality	✓	✓	✓	◦	✓	—
S7 : Rebuttal & Revision comment analysis, rebuttal gen	★ new	—	—	—	—	—
Phase 4: Dissemination						
S8 : Dissemination poster, slides, video, social	★ new	—	—	—	—	—
Stages covered	8/8	4	1	5	5	4

References

- [1] J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A. J. Ballard, J. Bambrick, S. W. Bodenstern, D. A. Evans, C.-C. Hung, M. O’Neill, D. Reiman, K. Tunyasuvunakool, Z. Wu, A. Zémgulytė, E. Arvaniti, C. Beattie, O. Bertolusso, A. Sherwood, J. M. Jumper, and D. Hassabis. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630:493–500, 2024.
- [2] S. Agarwal, G. Sahu, A. Puri, I. H. Laradji, K. D. Dvijotham, J. Stanley, L. Charlin, and C. Pal. LitLLM: A toolkit for literature review with large language models. *arXiv preprint arXiv:2402.01788*, 2024.
- [3] T. Aggarwal and A. Bhand. PASS: Presentation automation for slide generation and speech. *arXiv preprint arXiv:2501.06497*, 2025.
- [4] A. Ajith, M. Xia, A. Chevalier, T. Goyal, D. Chen, and T. Gao. LitSearch: A retrieval benchmark for scientific literature search. In *Conference on Empirical Methods in Natural Language Processing*, 2024.
- [5] I. Al Azher, M. J. Mokarrama, Z. Guo, S. R. Choudhury, and H. Alhoori. FutureGen: A RAG-based approach to generate the future work of scientific article. *arXiv preprint arXiv:2503.16561*, 2025.
- [6] I. Al Azher, Z. Guo, and H. Alhoori. Multi-agent LLMs for generating research limitations. *arXiv preprint arXiv:2601.11578*, 2026.
- [7] Alibaba NLP. Tongyi DeepResearch: An agentic LLM for long-horizon deep information seeking. <https://github.com/Alibaba-NLP/DeepResearch>, 2025.
- [8] Analemma.ai. FARS: Fully automated research system. <https://analemma.ai/blog/introducing-fars>, 2026.
- [9] A. Asai, J. He, R. Shao, W. Shi, A. Singh, J. C. Chang, K. Lo, L. Soldaini, S. Feldman, M. D’Arcy, D. Wadden, M. Latzke, J. Sparks, J. D. Hwang, V. Kishore, M. Tian, P. Ji, S. Liu, H. Tong, B. Wu, Y. Xiong, L. Zettlemoyer, G. Neubig, D. S. Weld, D. Downey, W. tau Yih, P. W. Koh, and H. Hajishirzi. Synthesizing scientific literature with retrieval-augmented language models. *Nature*, 650:857–863, 2026. doi: 10.1038/s41586-025-10072-4. URL <https://doi.org/10.1038/s41586-025-10072-4>.
- [10] J. Baek, S. K. Jauhar, S. Cucerzan, and S. J. Hwang. ResearchAgent: Iterative research idea generation over scientific literature with large language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6709–6738, 2025. doi: 10.18653/v1/2025.naacl-long.342. URL <https://aclanthology.org/2025.naacl-long.342/>.
- [11] C. Beger and C.-L. Henneking. Citegeist: Automated generation of related work analysis on the arXiv corpus. *arXiv preprint arXiv:2503.23229*, 2025.
- [12] J. Belouadi, A. Lauscher, and S. Eger. AutomaTikZ: Text-guided synthesis of scientific vector graphics with TikZ. In *International Conference on Learning Representations*, 2024.
- [13] J. Belouadi, S. P. Ponzetto, and S. Eger. DeTikZify: Synthesizing graphics programs for scientific figures and sketches with TikZ. In *Advances in Neural Information Processing Systems*, volume 37, pages 85074–85108, 2024.
- [14] Blog. 228 hours of non-stop work to produce 100 papers, burning through 11.4 billion tokens: Fars has gone crazy. <https://eu.36kr.com/en/p/3696795271966336>, 2026.
- [15] D. A. Boiko, R. MacKnight, B. Kline, and G. Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- [16] J. Bragg, M. D’Arcy, N. Balepur, D. Bareket, B. Dalvi, S. Feldman, D. Haddad, J. D. Hwang, P. Jansen, V. Kishore, B. P. Majumder, A. Naik, S. Rahamimov, K. Richardson, A. Singh, H. Surana, A. Tiktinsky, R. Vasu, G. Wiener, C. Anastasiades, S. Candra, J. Dunkelberger, D. Emery, R. Evans, M. Hamada, R. Huff, R. Kinney, M. Latzke, J. Lochner, R. Lozano-Aguilera, C. Nguyen, S. Rao, A. Tanaka, B. Vlahos, P. Clark, D. Downey, Y. Goldberg, A. Sabharwal, and D. S. Weld. AstaBench: Rigorous benchmarking of AI agents with a scientific research suite. In *International Conference on Learning Representations*, 2026.
- [17] A. M. Bran, S. Cox, O. Schilter, C. Baldassari, A. D. White, and P. Schwaller. ChemCrow: Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5):525–535, 2024.
- [18] ByteDance. DeerFlow: A deep research framework orchestrating sub-agents, memory, and sandboxes. <https://github.com/bytedance/deer-flow>, 2025.

- [19] J. Chai, S. Tang, R. Ye, Y. Du, X. Zhu, M. Zhou, Y. Wang, W. E, Y. Zhang, L. Zhang, and S. Chen. SciMaster: Towards general-purpose scientific AI agents, part I. X-Master as foundation: Can we lead on humanity’s last exam? *arXiv preprint arXiv:2507.05241*, 2025.
- [20] J. S. Chan, N. Chowdhury, O. Jaffe, J. Aung, D. Sherburn, E. Mays, G. Starace, K. Liu, L. Maksin, T. Patwardhan, L. Weng, and A. Madry. MLE-Bench: Evaluating machine learning agents on machine learning engineering. In *International Conference on Learning Representations*, 2025.
- [21] C.-C. Chen and I. Gurevych. Commitment checklist: Auditing author commitments in peer review. *arXiv preprint arXiv:2603.00003*, 2026.
- [22] D. Chen. AI-generated figures in academic publishing: Policies, tools, and practical guidelines. *arXiv preprint arXiv:2603.16159*, 2026.
- [23] G. Chen, J. Chen, L. Chen, J. Zhao, F. Meng, W. X. Zhao, R. Song, C. Chen, J.-R. Wen, and K. Jia. Toward autonomous long-horizon engineering for ML research. *arXiv preprint arXiv:2604.13018*, 2026.
- [24] H. Chen, M. Xiong, Y. Lu, W. Han, A. Deng, Y. He, J. Wu, Y. Li, Y. Liu, and B. Hooi. MLR-Bench: Evaluating AI agents on open-ended machine learning research. *arXiv preprint arXiv:2505.19955*, 2025.
- [25] N. Chen, A. H. Lin, J. Wu, J. Hou, Z. Zhang, Q. Wang, X. Wang, and B. He. XtraGPT: Context-aware and controllable academic paper revision. *arXiv preprint arXiv:2505.11336*, 2025.
- [26] Q. Chen, M. Yang, L. Qin, J. Liu, Z. Yan, J. Guan, D. Peng, Y. Ji, H. Li, M. Hu, Y. Zhang, Y. Liang, Y. Zhou, J. Wang, Z. Chen, and W. Che. AI4Research: A survey of artificial intelligence for scientific research. *arXiv preprint arXiv:2507.01903*, 2025.
- [27] S. Chen, J. Lai, J. Gao, H. Shi, Z. Liu, T. Ye, J. Luo, X. Wei, and L. Zhu. PosterOmni: Generalized artistic poster creation via task distillation and unified reward feedback. *arXiv preprint arXiv:2602.12127*, 2026.
- [28] S. Chen, J. Lai, J. Gao, T. Ye, H. Chen, H. Shi, S. Shao, Y. Lin, S. Fei, Z. Xing, Y. Jin, J. Luo, X. Wei, and L. Zhu. PosterCraft: Rethinking high-quality aesthetic poster generation in a unified framework. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026. URL <https://arxiv.org/abs/2506.10741>.
- [29] S. Chen, S. Zhong, D. P. Brumby, and A. L. Cox. What happens when reviewers receive AI feedback in their reviews? In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, 2026. doi: 10.1145/3772318.3791431. URL <https://doi.org/10.1145/3772318.3791431>.
- [30] Y. Chen, T. Lv, S. Zhang, Y. Yin, Y. Wan, P. S. Yu, and D. Chen. Paper2Web: Let’s make your paper alive! *arXiv preprint arXiv:2510.15842*, 2025.
- [31] Z. Chen, J. Chen, S. O. Arik, M. Sra, T. Pfister, and J. Yoon. CoDA: Agentic systems for collaborative data visualization. *arXiv preprint arXiv:2510.03194*, 2025.
- [32] Z. Chen, S. Chen, Y. Ning, Q. Zhang, B. Wang, B. Yu, Y. Li, Z. Liao, C. Wei, Z. Lu, V. Dey, M. Xue, F. N. Baker, B. Burns, D. Adu-Ampratwum, X. Huang, X. Ning, S. Gao, Y. Su, and H. Sun. ScienceAgentBench: Toward rigorous assessment of language agents for data-driven scientific discovery. In *International Conference on Learning Representations*, 2025.
- [33] J. Choi, S. Park, S. Song, and H. Shim. PosterForest: Hierarchical multi-agent collaboration for scientific poster generation. *arXiv preprint arXiv:2508.21720*, 2025.
- [34] P. H. Couto, Q. P. Ho, N. Kumari, B. K. Rachmat, T. G. H. Khuong, I. Ullah, and L. Sun-Hosoya. RelevAI-Reviewer: A benchmark on AI reviewers for survey paper relevance. *arXiv preprint arXiv:2406.10294*, 2024. URL <https://arxiv.org/abs/2406.10294>.
- [35] M. D’Arcy, T. Hope, L. Birnbaum, and D. Downey. MARG: Multi-agent review generation for scientific papers. *arXiv preprint arXiv:2401.04259*, 2024.
- [36] F. De Ponte. OpenDraft: 19-agent research draft generation. <https://github.com/federicodeponte/opendraft>, 2025.
- [37] X. Deng, J. Da, E. Pan, Y. Y. He, C. Ide, K. Garg, N. Lauffer, A. Park, N. Pasari, C. Rane, K. Sampath, M. Krishnan, S. Kundurthy, S. Hendryx, Z. Wang, V. Bharadwaj, J. Holm, R. Aluri, C. B. C. Zhang, N. Jacobson, B. Liu, and B. Kenstler. SWE-Bench Pro: Can AI agents solve long-horizon software engineering tasks? *arXiv preprint arXiv:2509.16941*, 2025.

- [38] A. Elovic. GPT Researcher: Autonomous agent for comprehensive online research. <https://github.com/assafelovic/gpt-researcher>, 2024.
- [39] T. Fan, F. Zhang, Y. Zheng, B. Chen, X. Niu, C. Huang, J. Lin, and C. Huang. DeepInnovator: Triggering the innovative capabilities of LLMs. *arXiv preprint arXiv:2602.18920*, 2026.
- [40] Y. Feng, Q. Huang, X. Xie, Z. Yang, J. Yu, W. Chen, and A. K. H. Tung. IDR Bench: Interactive deep research benchmark. *arXiv preprint arXiv:2601.06676*, 2026.
- [41] T.-J. Fu, W. Y. Wang, D. McDuff, and Y. Song. DOC2PPT: Automatic presentation slides generation from scientific documents. In *AAAI Conference on Artificial Intelligence*, volume 36, pages 634–642, 2022.
- [42] S. Gao, R. Zhu, P. Sui, Z. Kong, S. Aldogom, Y. Huang, A. Noori, R. Shamji, K. Parvataneni, T. Tsiligkaridis, and M. Zitnik. Democratizing AI scientists using ToolUniverse. *arXiv preprint arXiv:2509.23426*, 2025.
- [43] X. Gao, J. Ruan, Z. Zhang, J. Gao, T. Liu, and Y. Fu. ReviewAgents: Bridging the gap between human and AI-generated paper reviews. *arXiv preprint arXiv:2503.08506*, 2025.
- [44] Y. Gao, Q. Wu, and L. Zhu. Merging the citations received by arXiv-deposited e-prints and their corresponding published journal articles: Problems and perspectives. *Information Processing & Management*, 57(5):102267, 2020. doi: 10.1016/j.ipm.2020.102267. URL <https://doi.org/10.1016/j.ipm.2020.102267>.
- [45] Z. Gao, K. Brantley, and T. Joachims. Reviewer2: Optimizing review generation through prompt generation. *arXiv preprint arXiv:2402.10886*, 2024.
- [46] K. Garg, F. Shaik, S. Bandyopadhyay, and C. Caragea. Let’s use ChatGPT to write our paper! benchmarking LLMs to write the introduction of a research paper. *arXiv preprint arXiv:2508.14273*, 2025.
- [47] A. Garikaparathi, M. Patwardhan, L. Vig, and A. Cohan. IRIS: Interactive research ideation system for accelerating scientific discovery. In *Annual Meeting of the Association for Computational Linguistics*, pages 592–603, 2025. doi: 10.18653/v1/2025.acl-demo.57. URL <https://aclanthology.org/2025.acl-demo.57/>.
- [48] J. Ge, Z. Z. Wang, X. Zhou, Y.-H. Peng, S. Subramanian, Q. Tan, M. Sap, A. Suhr, D. Fried, G. Neubig, and T. Darrell. AutoPresent: Designing structured visuals from scratch. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2902–2911, 2025.
- [49] A. Ghafarollahi and M. J. Buehler. SciAgents: Automating scientific discovery through multi-agent intelligent graph reasoning. *arXiv preprint arXiv:2409.05556*, 2024.
- [50] E. Gibney. Major conference catches illicit AI use — and rejects hundreds of papers. *Nature News*, 652:281–282, 2026. doi: 10.1038/d41586-026-00893-2. URL <https://www.nature.com/articles/d41586-026-00893-2>.
- [51] G. H. T. Go, K. Ly, A. Sogaard, A. Tabatabaei, M. de Rijke, and X. Chen. LiRA: A multi-agent framework for reliable and readable literature review generation. *arXiv preprint arXiv:2510.05138*, 2025.
- [52] S. Goel, R. Hazra, D. Jayalath, T. Willi, P. Jain, W. F. Shen, I. Leontiadis, F. Barbieri, Y. Bachrach, J. Geiping, and C. Whitehouse. Training AI co-scientists using rubric rewards. *arXiv preprint arXiv:2512.23707*, 2025.
- [53] K. Goswami, P. Mathur, R. Rossi, and F. DERNONCOURT. PlotGen: Multi-agent LLM-based scientific data visualization via multimodal feedback. *arXiv preprint arXiv:2502.00988*, 2025.
- [54] J. Gottweis, W.-H. Weng, A. Daryin, T. Tu, A. Palepu, P. Sirkovic, A. Myaskovsky, F. Weissenberger, K. Rong, R. Tanno, K. Saab, D. Popovici, J. Blum, F. Zhang, K. Chou, A. Hassidim, B. Gokturk, A. Vahdat, P. Kohli, Y. Matias, A. Carroll, K. Kulkarni, N. Tomasev, Y. Guan, V. Dhillon, E. D. Vaishnav, B. Lee, T. R. D. Costa, J. R. Penadés, G. Peltz, Y. Xu, A. Pawlosky, A. Karthikesalingam, and V. Natarajan. Towards an AI co-scientist. *arXiv preprint arXiv:2502.18864*, 2025.
- [55] P. Goyal, M. Parmar, Y. Song, H. Palangi, T. Pfister, and J. Yoon. ScholarPeer: A context-aware multi-agent framework for automated peer review. *arXiv preprint arXiv:2601.22638*, 2026.
- [56] C. Greisinger and S. Eger. TikZilla: Scaling text-to-TikZ with high-quality data and reinforcement learning. *arXiv preprint arXiv:2603.03072*, 2026.
- [57] T. Gu, J. Wang, Z. Zhang, and H. Li. LLMs can realize combinatorial creativity: Generating creative ideas via LLMs for scientific research. *arXiv preprint arXiv:2412.14141*, 2024.
- [58] S. Guo, C. Deng, Y. Wen, H. Chen, Y. Chang, and J. Wang. DS-Agent: Automated data science by empowering large language models with case-based reasoning. *arXiv preprint arXiv:2402.17453*, 2024.

- [59] S. Guo, A. H. Shariatmadari, G. Xiong, A. Huang, M. Kim, C. M. Williams, S. Bekiranov, and A. Zhang. IdeaBench: Benchmarking large language models for research idea generation. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5888–5899, 2025. doi: 10.1145/3711896.3737419. URL <https://doi.org/10.1145/3711896.3737419>.
- [60] P. Han, Y. Yu, J. Xu, and J. You. DRPG (decompose, retrieve, plan, generate): An agentic framework for academic rebuttal. *arXiv preprint arXiv:2601.18081*, 2026.
- [61] Q. Hao, F. Xu, Y. Li, and J. Evans. Artificial intelligence tools expand scientists’ impact but contract science’s focus. *Nature*, 649:1237–1243, 2026.
- [62] Y. He, G. Huang, P. Feng, Y. Lin, Y. Zhang, H. Li, and W. E. PaSa: An LLM agent for comprehensive academic paper search. *arXiv preprint arXiv:2501.10120*, 2025.
- [63] Z. He, Z. Lyu, and Y. R. Fung. RebuttalAgent: Strategic persuasion in academic rebuttal via theory of mind. *arXiv preprint arXiv:2601.15715*, 2026.
- [64] HKU Data Intelligence Lab. Paper2Slides: From paper to presentation in one click. <https://github.com/HKUDS/Paper2Slides>, 2025.
- [65] M. Hong, D. Jiang, C. J. Zhang, Z. Guo, Y. Li, J. Chen, S. Cui, and Z. Su. CiteLLM: An agentic platform for trustworthy scientific reference discovery. *arXiv preprint arXiv:2602.23075*, 2026.
- [66] E. Hossain, S. K. Sinha, N. Bansal, R. A. Knipper, S. Sarkar, J. Salvador, Y. Mahajan, S. R. P. K. Guttikonda, M. Akter, M. M. Hassan, M. Freestone, M. C. W. Jr., D. Feng, and S. Karmaker. LLMs as meta-reviewers’ assistants: A case study. In *Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*, pages 7763–7803, 2025. doi: 10.18653/v1/2025.naacl-long.395. URL <https://aclanthology.org/2025.naacl-long.395/>.
- [67] J. Hou, A. H. Lin, N. Chen, Y. Gong, and B. He. PaperDebugger: A plugin-based multi-agent system for in-editor academic writing, review, and editing. *arXiv preprint arXiv:2512.02589*, 2025.
- [68] C.-C. Hsu, E. Bransom, J. Sparks, B. Kuehl, C. Tan, D. Wadden, L. L. Wang, and A. Naik. CHIME: LLM-assisted hierarchical organization of scientific studies for literature review support. *arXiv preprint arXiv:2407.16148*, 2024.
- [69] X. Hu, H. Fu, J. Wang, Y. Wang, Z. Li, R. Xu, Y. Lu, Y. Jin, L. Pan, and Z. Lan. Nova: An iterative planning and search approach to enhance novelty and diversity of LLM generated ideas. *arXiv preprint arXiv:2410.14255*, 2024.
- [70] X. Hu, Z. Zhao, S. Wei, Z. Chai, Q. Ma, G. Wang, X. Wang, J. Su, J. Xu, M. Zhu, Y. Cheng, J. Yuan, J. Li, K. Kuang, Y. Yang, H. Yang, and F. Wu. InfiAgent-DABench: Evaluating agents on data analysis tasks. *arXiv preprint arXiv:2401.05507*, 2024.
- [71] T. Hua, H. Hua, V. Xiang, B. Klieger, S. T. Truong, W. Liang, F.-Y. Sun, and N. Haber. ResearchCodeBench: Benchmarking LLMs on implementing novel machine learning research code. *arXiv preprint arXiv:2506.02314*, 2025.
- [72] K. Huang, S. Zhang, H. Wang, Y. Qu, Y. Lu, Y. Roohani, R. Li, L. Qiu, G. Li, J. Zhang, D. Yin, S. Marwaha, J. N. Carter, X. Zhou, M. Wheeler, J. A. Bernstein, M. Wang, P. He, J. Zhou, M. Snyder, L. Cong, A. Regev, and J. Leskovec. Biomni: A general-purpose biomedical AI agent. <https://github.com/snap-stanford/Biomni>, 2025.
- [73] Q. Huang, J. Vora, P. Liang, and J. Leskovec. MLAgentBench: Evaluating language agents on machine learning experimentation. In *International Conference on Machine Learning*, 2024.
- [74] S. Huang, Y. Gao, J. Bai, Y. Zhou, Z. Yin, X. Liu, R. Chellappa, C. P. Lau, S. Nag, C. Peng, and S. Pramanick. SciFig: Towards automating scientific figure generation. *arXiv preprint arXiv:2601.04390*, 2026.
- [75] M. Idahl and Z. Ahmadi. OpenReviewer: A specialized large language model for generating critical scientific paper reviews. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: System Demonstrations)*, pages 550–562, 2025. doi: 10.18653/v1/2025.naacl-demo.44. URL <https://aclanthology.org/2025.naacl-demo.44/>.
- [76] P. Jansen, M.-A. Cote, T. Khot, E. Bransom, B. Dalvi Mishra, B. P. Majumder, O. Tafjord, and P. Clark. DiscoveryWorld: A virtual environment for developing and evaluating automated scientific discovery agents. In *Advances in Neural Information Processing Systems*, 2024.

- [77] P. Jansen, O. Tafjord, M. Radensky, P. Siangliulue, T. Hope, B. Dalvi Mishra, B. P. Majumder, D. S. Weld, and P. Clark. CodeScientist: End-to-end semi-automated scientific discovery with code-based experimentation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 13370–13467, 2025. doi: 10.18653/v1/2025.findings-acl.692. URL <https://aclanthology.org/2025.findings-acl.692/>.
- [78] B. Jiang. HindSight: Evaluating LLM-generated research ideas via future impact. *arXiv preprint arXiv:2603.15164*, 2026.
- [79] L. Jiang, Y. Chai, M. Li, M. Liu, R. Fok, N. Dziri, Y. Tsvetkov, M. Sap, A. Albalak, and Y. Choi. Artificial hivemind: The open-ended homogeneity of language models (and beyond). In *Advances in Neural Information Processing Systems*, 2025.
- [80] Y. Jiang and A. Y. Ng. Automated scientific reviewing with agentic AI. <https://paperreview.ai/tech-overview>, 2025.
- [81] Z. Jiang, D. Schmidt, D. Srikanth, D. Xu, I. Kaplan, D. Jacenko, and Y. Wu. AIDE: AI-driven exploration in the space of code. *arXiv preprint arXiv:2502.13138*, 2025.
- [82] C. E. Jimenez, J. Yang, A. Wettig, S. Yao, K. Pei, O. Press, and K. Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *International Conference on Learning Representations*, 2024.
- [83] Y. Jin, Q. Zhao, Y. Wang, H. Chen, K. Zhu, Y. Xiao, and J. Wang. AgentReview: Exploring peer review dynamics with LLM agents. In *Conference on Empirical Methods in Natural Language Processing*, pages 1208–1226, 2024.
- [84] K. Jung, H. Cho, J. Yun, S. Yang, J. Jang, and J. Choo. Talk to your slides: High-efficiency slide editing via language-driven structured data manipulation. *arXiv preprint arXiv:2505.11604*, 2025. doi: 10.48550/arXiv.2505.11604. URL <https://arxiv.org/abs/2505.11604>.
- [85] S. Kale and V. Nadadur. TeXpert: A multi-level benchmark for evaluating LaTeX code generation by LLMs. In *Proceedings of the Fifth Workshop on Scholarly Document Processing (SDP 2025)*, pages 7–16, 2025. doi: 10.18653/v1/2025.sdp-1.2. URL <https://aclanthology.org/2025.sdp-1.2/>.
- [86] A. H. Kargaran, N. Nikeghbal, J. Yang, and N. Ousidhoum. Insights from ICLR peer review and rebuttal process. *arXiv preprint arXiv:2511.15462*, 2025.
- [87] A. Karpathy. autoresearch. <https://github.com/karpathy/autoresearch>, 2025.
- [88] J. Keuper. Prompt injection attacks on LLM generated reviews of scientific publications. *arXiv preprint arXiv:2509.10248*, 2025.
- [89] Y. Kim, K. Gu, C. Park, C. Park, S. Schmidgall, A. A. Heydari, Y. Yan, Z. Zhang, Y. Zhuang, M. Malhotra, P. P. Liang, H. W. Park, Y. Yang, X. Xu, Y. Du, S. Patel, T. Althoff, D. McDuff, and X. Liu. Towards a science of scaling agent systems. *arXiv preprint arXiv:2512.08296*, 2025.
- [90] D. Kobak, R. González-Márquez, E. Ákos Bernárdez-Vilaboa, and J. Lause. Delving into LLM-assisted writing in biomedical publications through excess vocabulary. *Science Advances*, 11(27):eadt3813, 2025.
- [91] P. T. J. Kon, J. Liu, Q. Ding, Y. Qiu, Z. Yang, Y. Huang, J. Srinivasa, M. Lee, M. Chowdhury, and A. Chen. Curie: Toward rigorous and automated scientific experimentation with AI agents. *arXiv preprint arXiv:2502.16069*, 2025.
- [92] P. T. J. Kon, Q. Ding, J. Liu, X. Zhu, J. Peng, J. Xing, Y. Huang, Y. Qiu, J. Srinivasa, M. Lee, M. Chowdhury, M. Zaharia, and A. Chen. EXP-Bench: Can AI conduct AI research experiments? In *International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=KjgyAm383Z>.
- [93] K. Kusumegi, X. Yang, P. Ginsparg, M. de Vaan, T. Stuart, and Y. Yin. Scientific production in the era of large language models. *Science*, 390:1240–1243, 2025.
- [94] J. Lála, O. O’Donoghue, A. Shtedritski, S. Cox, S. G. Rodrigues, and A. D. White. PaperQA: Retrieval-augmented generative agent for scientific research. *arXiv preprint arXiv:2312.07559*, 2023.
- [95] R. T. Lange, Y. Imajuku, and E. Cetin. ShinkaEvolve: Towards open-ended and sample-efficient program evolution. *arXiv preprint arXiv:2509.19349*, 2025.
- [96] J. M. Laurent, J. D. Janizek, M. Ruzo, M. M. Hinks, M. J. Hammerling, S. Narayanan, M. Ponnappati, A. D. White, and S. G. Rodrigues. Lab-Bench: Measuring capabilities of language models for biology research. *arXiv preprint arXiv:2407.10362*, 2024.

- [97] M. Lee, P. Liang, and Q. Yang. CoAuthor: Human-AI collaborative writing with language models. *arXiv preprint arXiv:2201.06796*, 2022.
- [98] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474, 2020.
- [99] B. Li, J. Wu, Y. Zhao, W. Xu, X. Chen, H. Yin, L. Chen, W. Zhang, and K. Li. UniScientist: Advancing universal scientific research intelligence. <https://unipat.ai/blog/UniScientist>, 2026.
- [100] E. Li, J. Ren, X. Pan, C. Yan, C. Li, D. Bergemann, and Z. Yang. Build your personalized research group: A multiagent framework for continual and interactive science automation. *arXiv preprint arXiv:2510.15624*, 2025.
- [101] J. Li, S. Li, Z. Gao, Q. Shi, Y. Li, Z. Wang, J. Huang, H. Wang, J. Wang, X. Han, Z. Liu, and M. Sun. TritonBench: Benchmarking large language model capabilities for generating triton operators. In *Findings of the Association for Computational Linguistics: ACL 2025*, 2025.
- [102] L. Li, W. Xu, J. Guo, R. Zhao, X. Li, Y. Yuan, B. Zhang, Y. Jiang, Y. Xin, R. Dang, D. Zhao, Y. Rong, T. Feng, and L. Bing. Chain of ideas: Revolutionizing research via novel idea development with LLM agents. *arXiv preprint arXiv:2410.13185*, 2024.
- [103] M. Li, Y. Zeng, Z. Cheng, C. Ma, and K. Jia. ReportBench: Evaluating deep research agents via academic survey tasks. *arXiv preprint arXiv:2508.15804*, 2025.
- [104] R. Li, T. Patel, Q. Wang, and X. Du. MLR-Copilot: Autonomous machine learning research based on large language models agents. *arXiv preprint arXiv:2408.14033*, 2024.
- [105] S. Li, L. Fan, Y. Lin, Z. Li, X. Wei, S. Ni, H. Alinejad-Rokny, and M. Yang. Automatic paper reviewing with heterogeneous graph reasoning over LLM-simulated reviewer-author debates. *arXiv preprint arXiv:2511.08317*, 2026. AAAI-2026.
- [106] Y. Li. awesome-ai-research-writing. <https://github.com/Leey21/awesome-ai-research-writing>, 2024.
- [107] Y. Li, C. Shao, X. Liu, R. Zhao, P. Liu, H. Su, Z. Chen, Q. Yang, A. Xu, Y. Fang, Q. Zeng, T. Li, J. Xu, F. Xu, Y. Li, and T.-Y. Liu. AutoSOTA: An end-to-end automated research system for state-of-the-art AI model discovery. *arXiv preprint arXiv:2604.05550*, 2026.
- [108] Z. Li, D. Jiang, X. Ma, H. Zhang, P. Nie, Y. Zhang, K. Zou, J. Xie, Y. Zhang, and W. Chen. OpenResearcher: A fully open pipeline for long-horizon deep research trajectory synthesis. *arXiv preprint arXiv:2603.20278*, 2026.
- [109] W. Liang, Y. Zhang, Z. Wu, H. Lepp, W. Ji, X. Zhao, H. Cao, S. Liu, S. He, Z. Huang, D. Yang, C. Potts, C. D. Manning, and J. Y. Zou. Mapping the increasing use of LLMs in scientific papers. *arXiv preprint arXiv:2404.01268*, 2024.
- [110] X. Liang, J. Yang, Y. Wang, C. Tang, Z. Zheng, S. Song, Z. Lin, Y. Yang, S. Niu, H. Wang, B. Tang, F. Xiong, K. Mao, and Z. Li. SurveyX: Academic survey automation via large language models. *arXiv preprint arXiv:2502.14776*, 2025.
- [111] X. Liang, X. Zhang, Y. Xu, S. Sun, and C. You. SlideGen: Collaborative multimodal agents for scientific slide generation. *arXiv preprint arXiv:2512.04529*, 2025.
- [112] C.-Y. Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, 2004. URL <https://aclanthology.org/W04-1013/>.
- [113] T.-L. Lin, W.-C. Chen, T.-F. Hsiao, H.-I. Liu, Y.-H. Yeh, Y.-K. Chan, W.-S. Lien, P.-Y. Kuo, P. S. Yu, and H.-H. Shuai. Breaking the reviewer: Assessing the vulnerability of large language models in automated peer review under textual adversarial attacks. *arXiv preprint arXiv:2506.11113*, 2025.
- [114] Z. Lin, Q. Xie, M. Zhu, S. Li, Q. Sun, E. Gu, Y. Ding, K. Sun, F. Guo, P. Lu, Z. Ning, Y. Weng, and Y. Zhang. AutoFigure-Edit: Generating editable scientific illustration. *arXiv preprint arXiv:2603.06674*, 2026.
- [115] J. Ling, Y. Qi, T. Huang, S. Zhou, Y. Huang, J. Yang, Z. Song, Y. Zhou, Y. Yang, H. T. Shen, and P. Wang. Table2LaTeX-RL: High-fidelity LaTeX code generation from table images via reinforced multimodal language models. *arXiv preprint arXiv:2509.17589*, 2025.
- [116] J. Liu, L. Yang, H. Luo, F. Wang, H. Li, and M. Wang. Preacher: Paper-to-video agentic system. In *IEEE/CVF International Conference on Computer Vision*, pages 17129–17139, 2025.

- [117] J. Liu, P. Xia, S. Han, S. Qiu, L. Zhang, G. Chen, H. Tu, X. Yang, J. Zhou, H. Zhu, Y. Li, J. Zhang, Y. Zhou, Z. Zheng, C. Xie, M. Ding, and H. Yao. AutoResearchClaw: Fully autonomous research from idea to paper. <https://github.com/aiming-lab/AutoResearchClaw>, 2026.
- [118] W. Liu, Z. Yang, Y. Zhao, and X. Li. RATE: Reviewer profiling and annotation-free training for expertise ranking in peer review systems. *arXiv preprint arXiv:2601.19637*, 2026.
- [119] Y. Liu, Y. Wu, D. Zhang, and L. Sun. Agentic AutoSurvey: Let LLMs survey LLMs. *arXiv preprint arXiv:2509.18661*, 2025.
- [120] Y. Liu, Z. Yang, T. Xie, J. Ni, B. Gao, Y. Li, S. Tang, W. Ouyang, E. Cambria, and D. Zhou. ResearchBench: Benchmarking LLMs in scientific discovery via inspiration-based task decomposition. In *Findings of the Association for Computational Linguistics: ACL 2026*, 2026. URL <https://arxiv.org/abs/2503.21248>.
- [121] Z. Liu, X. Bao, P. Li, J. Zhou, Z. Liao, Y. He, K. Jiang, C.-W. Xie, Y. Zheng, and H. Xie. ShowTable: Unlocking creative table visualization with collaborative reflection and refinement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.
- [122] C. Lu, C. Lu, R. T. Lange, J. Foerster, J. Clune, and D. Ha. The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- [123] C. Lu, C. Lu, R. T. Lange, Y. Yamada, S. Hu, J. Foerster, D. Ha, and J. Clune. Towards end-to-end automation of AI research. *Nature*, 651:914–919, 2026.
- [124] Y. Luo, R. Wang, P. Gam, J. Cui, circlestarzero, S. Ni, J. Quanta, Q. Fu, and S. Hou. ChatPaper: Use LLM to summarize papers. <https://github.com/kaixindelele/ChatPaper>, 2023.
- [125] Z. Luo, A. Kasirzadeh, and N. B. Shah. The more you automate, the less you see: Hidden pitfalls of AI scientist systems. *arXiv preprint arXiv:2509.08713*, 2025.
- [126] Z. Luo, Z. Yang, Z. Xu, W. Yang, and X. Du. LLM4SR: A survey on large language models for scientific research. *arXiv preprint arXiv:2501.04306*, 2025.
- [127] Y. Lyu, X. Zhang, X. Yi, Y. Zhao, S. Guo, W. Hu, J. Piotrowski, J. Kaliski, J. Urbani, Z. Meng, L. Zhou, and X. Yan. EvoScientist: Towards multi-agent evolving AI scientists for end-to-end scientific discovery. *arXiv preprint arXiv:2603.08127*, 2026.
- [128] J. Lala, O. O’Donoghue, A. Shtedritski, S. Cox, S. G. Rodrigues, and A. D. White. PaperQA: Retrieval-augmented generative agent for scientific research. *arXiv preprint arXiv:2312.07559*, 2023.
- [129] Q. Ma, C. Guo, Z. Tian, S. Wang, J. Xiao, Y. Yue, and Z. Zhang. Paper2Rebuttal: A multi-agent framework for transparent author response assistance. *arXiv preprint arXiv:2601.14171*, 2026.
- [130] B. P. Majumder, H. Surana, D. Agarwal, B. D. Mishra, A. Meena, A. Prakhar, T. Vora, T. Khot, A. Sabharwal, and P. Clark. DiscoveryBench: Towards data-driven discovery with large language models. In *International Conference on Learning Representations*, 2025.
- [131] METR Team. Measuring AI ability to complete long tasks. <https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks>, 2025.
- [132] J. Miao, J. R. Davis, Y. Zhang, J. K. Pritchard, and J. Zou. Paper2Agent: Reimagining research papers as interactive and reliable AI agents. *arXiv preprint arXiv:2509.06917*, 2025.
- [133] L. Mitchener, A. Yiu, B. Chang, M. Bourdenx, T. Nadolski, A. Sulovari, E. C. Landsness, D. L. Barabasi, S. Narayanan, N. Evans, S. Reddy, M. Foiani, A. Kamal, L. P. Shriver, F. Cao, A. T. Wassie, J. M. Laurent, E. Melville-Green, M. Caldas, A. Bou, K. F. Roberts, S. Zagorac, T. C. Orr, M. E. Orr, K. J. Zvezdaryk, A. E. Ghareeb, L. McCoy, B. Gomes, E. A. Ashley, K. E. Duff, T. Buonassisi, T. Rainforth, R. J. Bateman, M. Skarlinski, S. G. Rodrigues, M. M. Hinks, and A. D. White. Kosmos: An AI scientist for autonomous discovery. *arXiv preprint arXiv:2511.02824*, 2025.
- [134] M. Naddaf. More than half of researchers now use AI for peer review — often against guidance. *Nature News*, 649: 273–274, 2026. doi: 10.1038/d41586-025-04066-5. URL <https://www.nature.com/articles/d41586-025-04066-5>.
- [135] D. Nathani, L. Madaan, N. Roberts, N. Bashlykov, A. Menon, V. Moens, A. Budhiraja, D. Magka, V. Vorotilov, G. Chaurasia, D. Hupkes, R. S. Cabral, T. Shavrina, J. Foerster, Y. Bachrach, W. Y. Wang, and R. Raileanu. MLGym: A new framework and benchmark for advancing AI research agents. *arXiv preprint arXiv:2502.14499*, 2025.

- [136] B. Newman, Y. Lee, A. Naik, P. Siangliulue, R. Fok, J. Kim, D. S. Weld, J. C. Chang, and K. Lo. ArxivDIGESTables: Synthesizing scientific literature into tables using language models. In *Conference on Empirical Methods in Natural Language Processing*, pages 9612–9631, 2024.
- [137] M.-A. Nguyen, M.-D. Nguyen, H. L. N.T., K. H. Dang, N. T. Dong, and D. D. Le. SurveyG: A multi-agent LLM framework with hierarchical citation graph for automated survey generation. *arXiv preprint arXiv:2510.07733*, 2025.
- [138] S. Ni. ChatReviewer: ChatGPT-based paper reviewing and response generation. <https://github.com/nishiwen1214/ChatReviewer>, 2023.
- [139] Y. Ni, P. Nie, K. Zou, X. Yue, and W. Chen. VisCoder: Fine-tuning LLMs for executable python visualization code generation. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 2956–2983, 2025. doi: 10.18653/v1/2025.findings-emnlp.160. URL <https://aclanthology.org/2025.findings-emnlp.160/>.
- [140] A. Novikov, N. Vu, M. Eisenberger, E. Dupont, P.-S. Huang, A. Z. Wagner, S. Shirobokov, B. Kozlovskii, F. J. R. Ruiz, A. Mehrabian, M. P. Kumar, A. See, S. Chaudhuri, G. Holland, A. Davies, S. Nowozin, P. Kohli, and M. Balog. AlphaEvolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.
- [141] O. O’Donoghue, A. Shtedritski, J. Ginger, R. Abboud, A. E. Ghareeb, J. Booth, and S. G. Rodrigues. BioPlanner: Automatic evaluation of LLMs on protocol planning. *arXiv preprint arXiv:2310.10632*, 2023.
- [142] J. Ou, W. G. Walden, K. Sanders, Z. Jiang, K. Sun, J. Cheng, W. Jurayj, M. Wanner, S. Liang, C. Morgan, S. Han, W. Wang, C. May, H. Recknor, D. Khashabi, and B. V. Durme. ClaimCheck: How grounded are LLM critiques of scientific papers? In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 21712–21735, 2025. doi: 10.18653/v1/2025.findings-emnlp.1185. URL <https://aclanthology.org/2025.findings-emnlp.1185/>.
- [143] A. Ouyang, S. Guo, S. Arora, A. L. Zhang, W. Hu, C. Ré, and A. Mirhoseini. KernelBench: Can LLMs write efficient GPU kernels? In *International Conference on Machine Learning*, 2025.
- [144] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744, 2022.
- [145] B. Pan, Y. Fu, K. Wang, J. Lu, L. Pan, Z. Qian, Y. Chen, G. Wang, Y. Zhou, L. Zheng, Y. Tang, Z. Wen, Y. Wu, J. Lu, B. Zhu, M. Zhu, B. Zhang, and W. Chen. VIS-Shepherd: Constructing critic for LLM-based data visualization generation. *arXiv preprint arXiv:2506.13326*, 2025.
- [146] W. Pang, K. Q. Lin, X. Jian, X. He, and P. Torr. Paper2Poster: Towards multimodal poster automation from scientific papers. In *NeurIPS 2025 Datasets and Benchmarks Track*, 2025. URL <https://openreview.net/forum?id=p0E74lpRBD>.
- [147] S. S. Panigrahi, J. Videnovic, and M. Brbic. HeurekaBench: A benchmarking framework for AI co-scientist. In *International Conference on Learning Representations*, 2026.
- [148] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. BLEU: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL)*, pages 311–318, 2002. URL <https://aclanthology.org/P02-1040/>.
- [149] L. Patel, N. Arabzadeh, H. Gupta, A. Sundar, I. Stoica, M. Zaharia, and C. Guestrin. DeepScholar-Bench: A live benchmark and automated evaluation for generative research synthesis. *arXiv preprint arXiv:2508.20033*, 2025.
- [150] R. A. Poldrack. AI-Peer-Review: Multi-LLM peer review with meta-review. <https://github.com/poldrack/ai-peer-review>, 2024.
- [151] O. Press, A. Hochlehnert, A. Prabhu, V. Udandarao, O. Press, and M. Bethge. CiteME: Can language models accurately cite scientific claims? *arXiv preprint arXiv:2407.12861*, 2024.
- [152] K. Pu, K. J. K. Feng, T. Grossman, T. Hope, B. D. Mishra, M. Latzke, J. Bragg, J. C. Chang, and P. Siangliulue. IdeaSynth: Iterative research idea development through evolving and composing idea facets with literature-grounded feedback. In *CHI Conference on Human Factors in Computing Systems*, pages 1–31, 2025.
- [153] Y. Pu, T. Lin, and H. Chen. PiFlow: Principle-aware scientific discovery with multi-agent collaboration. *arXiv preprint arXiv:2505.15047*, 2025.

- [154] Y. Qiu, H. Zhang, Z. Xu, M. Li, D. Song, Z. Wang, and K. Zhang. AI Idea Bench 2025: AI research idea generation benchmark. *arXiv preprint arXiv:2504.14191*, 2025.
- [155] A. Qu, H. Zheng, Z. Zhou, Y. Yan, Y. Tang, S. Y. Ong, F. Hong, K. Zhou, C. Jiang, M. Kong, J. Zhu, X. Jiang, S. Li, C. Wu, B. K. H. Low, J. Zhao, and P. P. Liang. CORAL: Towards autonomous multi-agent evolution for open-ended discovery. *arXiv preprint arXiv:2604.01658*, 2026.
- [156] Y. Qu, K. Huang, M. Yin, K. Zhan, D. Liu, D. Yin, H. C. Cousins, W. A. Johnson, X. Wang, M. Shah, R. B. Altman, D. Zhou, M. Wang, and L. Cong. CRISPR-GPT for agentic automation of gene-editing experiments. *arXiv preprint arXiv:2404.18021*, 2024.
- [157] V. Raina, A. Liusie, and M. Gales. Is LLM-as-a-judge robust? investigating universal adversarial attacks on zero-shot LLM assessment. In *Conference on Empirical Methods in Natural Language Processing*, pages 7499–7517, 2024.
- [158] B. Rank, H. Bhatnagar, A. Prabhu, S. Eisenberg, K. Nguyen, M. Bethge, and M. Andriushchenko. PostTrainBench: Can LLM agents automate LLM post-training? *arXiv preprint arXiv:2603.08640*, 2026.
- [159] V. Rao, A. Kumar, H. Lakkaraju, and N. B. Shah. Detecting LLM-generated peer reviews. *arXiv preprint arXiv:2503.15772*, 2025.
- [160] J. A. Rodriguez, A. Puri, S. Agarwal, I. H. Laradji, P. Rodriguez, S. Rajeswar, D. Vazquez, C. Pal, and M. Pedersoli. StarVector: Generating scalable vector graphics code from images and text. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16175–16186, 2025.
- [161] B. Romera-Paredes, M. Barekatin, A. Novikov, M. Balog, M. P. Kumar, E. Dupont, F. J. R. Ruiz, J. S. Ellenberg, P. Wang, O. Fawzi, P. Kohli, and A. Fawzi. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475, 2024.
- [162] K. Ruan, X. Wang, J. Hong, P. Wang, Y. Liu, and H. Sun. LiveIdeaBench: Evaluating LLMs’ scientific creativity and idea generation with minimal context. *arXiv preprint arXiv:2412.17596*, 2024.
- [163] Q. Ruan and I. Gurevych. Author-in-the-loop response generation and evaluation: Integrating author expertise and intent in responses to peer review. *arXiv preprint arXiv:2602.11173*, 2026.
- [164] G. Russo Latona, M. H. Ribeiro, T. R. Davidson, V. Veselovsky, and R. West. The AI review lottery: Widespread AI-assisted peer reviews boost paper scores and acceptance rates. *arXiv preprint arXiv:2405.02150*, 2024.
- [165] R. Saha, G. Juneja, D. Chaudhuri, N. Sajeewan, N. B. Shah, and D. Pruthi. Policies permitting LLM use for polishing peer reviews are currently not enforceable. *arXiv preprint arXiv:2603.20450*, 2026.
- [166] D. Sahoo, M. Prasad, V. Majhi, J. Singh, V. Chamola, Y. Sinha, M. Mandal, and D. Kumar. When reject turns into accept: Quantifying the vulnerability of LLM-based scientific reviewers to indirect prompt injection. *arXiv preprint arXiv:2512.10449*, 2025.
- [167] G. Sahu, H. Larochelle, L. Charlin, and C. Pal. ReviewerToo: Should AI join the program committee? a look at the future of peer review. *arXiv preprint arXiv:2510.08867*, 2025.
- [168] A. Sanyal, S. Schapiro, S. Shashidhar, R. Moon, L. R. Varshney, and D. Hakkani-Tur. Spark: A system for scientifically creative idea generation. In *International Conference on Computational Creativity*, 2025.
- [169] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, and T. Scialom. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551, 2023.
- [170] S. Schmidgall and M. Moor. AgentRxiv: Towards collaborative autonomous research. *arXiv preprint arXiv:2503.18102*, 2025.
- [171] S. Schmidgall, Y. Su, Z. Wang, X. Sun, J. Wu, X. Yu, J. Liu, M. Moor, Z. Liu, and E. Barsoum. Agent laboratory: Using LLM agents as research assistants. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5977–6043, 2025. doi: 10.18653/v1/2025.findings-emnlp.320. URL <https://aclanthology.org/2025.findings-emnlp.320/>.
- [172] N. Schuster, A. N. Salcedo, S. Bouchard, D. Frei, A. Pisani, J. E. Bautista, J. Zoubian, S. Escoffier, W. Liu, G. Valogiannis, and P. Zarrouk. Setting SAIL: Leveraging scientist-AI-loops for rigorous visualization tools. *arXiv preprint arXiv:2603.18145*, 2026.

- [173] SenseTime. Office raccoon: Ai-powered office agent. <https://www.sensetime.com/en/news-detail/51170569>, 2026. News and Blog, March 4, 2026.
- [174] M. Seo, J. Baek, S. Lee, and S. J. Hwang. Paper2Code: Automating code generation from scientific papers in machine learning. *arXiv preprint arXiv:2504.17192*, 2025.
- [175] F. Shahhosseini, A. Marioriyad, A. Momen, M. S. Baghshah, M. H. Rohban, and S. H. Javanmard. Large language models for scientific idea generation: A creativity-centered survey. *arXiv preprint arXiv:2511.07448*, 2025.
- [176] C. Shao, Y. Li, and F. Xu. SciNetBench: A relation-aware benchmark for scientific literature retrieval agents. *arXiv preprint arXiv:2601.03260*, 2026.
- [177] E. Shao, Y. Wang, Y. Qian, Z. Pan, H. Liu, and D. Wang. SciSciGPT: advancing human–AI collaboration in the science of science. *Nature Computational Science*, 6:301–315, 2026.
- [178] Y. Shao, Y. Jiang, T. A. Kanell, P. Xu, O. Khattab, and M. S. Lam. Assisting in writing Wikipedia-like articles from scratch with large language models. *arXiv preprint arXiv:2402.14207*, 2024.
- [179] H. Shen, H. Yang, Z. Gu, and W. Han. ScholarGym: Benchmarking large language model capabilities in the information-gathering stage of deep research. *arXiv preprint arXiv:2601.21654*, 2026.
- [180] C. Shi, Q. Cai, Z. Chen, L. Zeng, Y. Zhao, J. Yu, J. Yu, and X. Li. APEX: Academic poster editing agentic expert. *arXiv preprint arXiv:2601.04794*, 2026.
- [181] J. Shi, Z. Zhang, B. Wu, Y. Liang, M. Fang, L. Chen, and Y. Zhao. PresentAgent: Multimodal agent for presentation video generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 760–773, 2025. doi: 10.18653/v1/2025.emnlp-demos.58. URL <https://aclanthology.org/2025.emnlp-demos.58/>.
- [182] N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652, 2023.
- [183] C. Si, D. Yang, and T. Hashimoto. Can LLMs generate novel research ideas? A large scale human study with 100+ NLP researchers. *arXiv preprint arXiv:2409.04109*, 2024.
- [184] C. Si, T. Hashimoto, and D. Yang. The ideation-execution gap: Execution outcomes of LLM-generated versus human research ideas. *arXiv preprint arXiv:2506.20803*, 2025.
- [185] C. Si, Z. Yang, Y. Choi, E. Candes, D. Yang, and T. Hashimoto. Towards execution-grounded automated AI research. *arXiv preprint arXiv:2601.14525*, 2026.
- [186] M. Siddiqui, R. Pea, and H. Subramonyam. Script&shift: A layered interface paradigm for integrating content development and rhetorical strategy with LLM writing assistants. In *CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2025.
- [187] M. N. Siddiqui, V. Feliciano, R. Pea, and H. Subramonyam. AI in the writing process: How purposeful AI support fosters student writing. In *International Conference on Artificial Intelligence in Education*, pages 190–203, 2025.
- [188] M. N. Siddiqui, N. Nasser, A. Coscia, R. Pea, and H. Subramonyam. DraftMarks: Enhancing transparency in human-AI co-writing through interactive skeuomorphic process traces. In *arXiv preprint arXiv:2509.23505*, 2025.
- [189] M. D. Skarlinski, S. Cox, J. M. Laurent, J. D. Braza, M. Hinks, M. J. Hammerling, M. Ponnampati, S. G. Rodrigues, and A. D. White. Language agents achieve superhuman synthesis of scientific knowledge. *arXiv preprint arXiv:2409.13740*, 2024.
- [190] Skywork AI. DeepResearchAgent: A hierarchical multi-agent system for deep research. <https://github.com/SkyworkAI/DeepResearchAgent>, 2025.
- [191] Y. Song, Y. Song, T. Pfister, and J. Yoon. PaperOrchestra: A multi-agent framework for automated AI research paper writing. *arXiv preprint arXiv:2604.05018*, 2026.
- [192] G. Starace, O. Jaffe, D. Sherburn, J. Aung, J. S. Chan, L. Maksin, R. Dias, E. Mays, B. Kinsella, W. Thompson, J. Heidecke, A. Glaese, and T. Patwardhan. PaperBench: Evaluating AI’s ability to replicate AI research. In *International Conference on Machine Learning*, 2025.

- [193] H. Su, R. Chen, S. Tang, Z. Yin, X. Zheng, J. Li, B. Qi, Q. Wu, H. Li, W. Ouyang, P. Torr, B. Zhou, and N. Dong. Many heads are better than one: Improved scientific idea generation by a LLM-based multi-agent system. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 28201–28240, 2025. doi: 10.18653/v1/2025.acl-long.1368. URL <https://aclanthology.org/2025.acl-long.1368/>.
- [194] T. Sun, E. Pan, Z. Yang, K. Sui, J. Shi, X. Cheng, T. Li, G. Zhang, W. Huang, J. Yang, and Z. Li. P2P: Automated paper-to-poster generation and fine-grained benchmark. In *International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=JoJyT9niJL>.
- [195] SWE-bench Team. SWE-bench Leaderboards. <https://www.swebench.com>, 2026. Accessed: 2026-04-24.
- [196] P. Taechoyotin and D. Acuna. REMOR: Automated peer review generation with LLM reasoning and multi-objective reinforcement learning. *arXiv preprint arXiv:2505.11718*, 2025.
- [197] C. Tan, D. Lyu, S. Li, Z. Gao, J. Wei, S. Ma, Z. Liu, and S. Z. Li. Peer review as a multi-turn and long-context dialogue with role-based interactions. *arXiv preprint arXiv:2406.05688*, 2024.
- [198] J. Tang, T. Fan, and C. Huang. AutoAgent: A fully-automated and zero-code framework for LLM agents. *arXiv preprint arXiv:2502.05957*, 2025.
- [199] J. Tang, L. Xia, Z. Li, and C. Huang. AI-Researcher: Autonomous scientific innovation. *arXiv preprint arXiv:2505.18705*, 2025.
- [200] X. Tang, X. Duan, and Z. G. Cai. Large language models for automated literature review. *arXiv preprint arXiv:2412.13612*, 2024.
- [201] M. V. T. Thai, T. Le, D. N. Manh, H. P. Nhat, and N. D. Q. Bui. SWE-EVO: Benchmarking coding agents in long-horizon software evolution scenarios. *arXiv preprint arXiv:2512.18470*, 2025.
- [202] N. Thakkar, M. Yuksekgonul, J. Silberg, A. Garg, N. Peng, F. Sha, R. Yu, C. Vondrick, and J. Zou. A large-scale randomized study of large language model feedback in peer review. *Nature Machine Intelligence*, 8:326–336, 2026. doi: 10.1038/s42256-026-01188-x. URL <https://www.nature.com/articles/s42256-026-01188-x>.
- [203] M. Tian, L. Gao, S. D. Zhang, X. Chen, C. Fan, X. Guo, R. Haas, P. Ji, K. Krongchon, Y. Li, S. Liu, D. Luo, Y. Ma, H. Tong, K. Trinh, C. Tian, Z. Wang, B. Wu, Y. Xiong, S. Yin, M. Zhu, K. Lieret, Y. Lu, G. Liu, Y. Du, T. Tao, O. Press, J. Callan, E. Huerta, and H. Peng. SciCode: A research coding benchmark curated by scientists. *arXiv preprint arXiv:2407.13168*, 2024.
- [204] Y. Tian, W. Cui, D. Deng, X. Yi, Y. Yang, H. Zhang, and Y. Wu. ChartGPT: Leveraging LLMs to generate charts from abstract natural language. *arXiv preprint arXiv:2311.01920*, 2023.
- [205] G. Tie, P. Zhou, and L. Sun. A survey of AI scientists. *arXiv preprint arXiv:2510.23045*, 2025.
- [206] K. Ueda, W. Hirota, T. Asakura, T. Omi, K. Takahashi, K. Arima, and T. Ishigaki. Exploring design of multi-agent LLM dialogues for research ideation. In *Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 322–337, 2025.
- [207] R. van de Schoot, J. de Bruin, R. Schram, P. Zahedi, J. de Boer, F. Weijdemans, B. Kramer, M. Huijts, M. Hoogerwerf, G. Ferdinands, A. Harkema, J. Willemsen, Y. Ma, Q. Fang, S. Hindriks, L. Tummers, and D. L. Oberski. An open source machine learning framework for efficient and transparent systematic reviews. *Nature Machine Intelligence*, 3:125–133, 2021.
- [208] M. Wang, R. Lin, K. Hu, J. Jiao, N. Chowdhury, E. Chang, and T. Patwardhan. FrontierScience: Evaluating AI’s ability to perform expert-level scientific tasks. *arXiv preprint arXiv:2601.21165*, 2026.
- [209] Q. Wang, D. Downey, H. Ji, and T. Hope. SciMON: Scientific inspiration machines optimized for novelty. In *Annual Meeting of the Association for Computational Linguistics*, pages 279–299, 2024.
- [210] Q. Wang, H. Wang, L. Chen, Z. Yang, G. Chen, H. Alinejad-Rokny, H. Li, Y. Lin, and M. Yang. FlowPIE: Test-time scientific idea evolution with flow-guided literature exploration. *arXiv preprint arXiv:2603.29557*, 2026.
- [211] W. Wang, L. Gu, L. Zhang, Y. Luo, Y. Dai, C. Shen, L. Xie, B. Lin, X. He, and J. Ye. SciPIP: An LLM-based scientific paper idea proposer. *arXiv preprint arXiv:2410.23166*, 2024.
- [212] X. Wang, B. Li, Y. Song, F. F. Xu, X. Tang, M. Zhuge, J. Pan, Y. Song, B. Li, J. Singh, H. H. Tran, F. Li, R. Ma, M. Zheng, B. Qian, Y. Shao, N. Muennighoff, Y. Zhang, B. Hui, J. Lin, R. Brennan, H. Peng, H. Ji,

- and G. Neubig. OpenHands: An open platform for AI software developers as generalist agents. *arXiv preprint arXiv:2407.16741*, 2024.
- [213] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi. Self-Instruct: Aligning language models with self-generated instructions. In *Annual Meeting of the Association for Computational Linguistics*, pages 13484–13508, 2023.
- [214] Y. Wang, Q. Guo, W. Yao, H. Zhang, X. Zhang, Z. Wu, M. Zhang, X. Dai, M. Zhang, Q. Wen, W. Ye, S. Zhang, and Y. Zhang. AutoSurvey: Large language models can automatically write surveys. *arXiv preprint arXiv:2406.10252*, 2024.
- [215] Y. Wang, X. Ma, P. Nie, H. Zeng, Z. Lyu, Y. Zhang, B. Schneider, Y. Lu, X. Yue, and W. Chen. ScholarCopilot: Training large language models for academic writing with accurate citations. *arXiv preprint arXiv:2504.00824*, 2025.
- [216] Z. Wang, H. Zhang, C.-L. Li, J. M. Eisenschlos, V. Perot, Z. Wang, L. Miculicich, Y. Fujii, J. Shang, C.-Y. Lee, and T. Pfister. Chain-of-table: Evolving tables in the reasoning chain for table understanding. In *International Conference on Learning Representations*, 2024.
- [217] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837, 2022.
- [218] J. Wei, C. Tan, Q. Chen, G. Wu, S. Li, Z. Gao, L. Sun, B. Yu, and R. Guo. From words to structured visuals: A benchmark and framework for text-to-diagram generation and editing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13315–13325, 2025.
- [219] Z. Wen, J. Cao, Z. Wang, B. Guo, R. Yang, and S. Liu. InteractiveSurvey: An LLM-based personalized and interactive survey paper generation system. *arXiv preprint arXiv:2504.08762*, 2025.
- [220] Y. Weng, M. Zhu, G. Bao, H. Zhang, J. Wang, Y. Zhang, and L. Yang. CycleResearcher: Improving automated research via automated review. In *International Conference on Learning Representations*, 2025.
- [221] H. Wijk, T. Lin, J. Becker, S. Jawhar, N. Parikh, T. Broadley, L. Chan, M. Chen, J. Clymer, J. Dhyani, E. Elicheva, K. Garcia, B. Goodrich, N. Jurkovic, H. Karnofsky, M. Kinniment, A. Lajko, S. Nix, L. Sato, W. Saunders, M. Taran, B. West, and E. Barnes. RE-Bench: Evaluating frontier AI R&D capabilities of language model agents against human experts. *arXiv preprint arXiv:2411.15114*, 2024.
- [222] H. Wu, B. Zheng, D. Song, Y. Jiang, J. Gao, L. Xing, L. Sun, and Y. Yuan. Towards a medical AI scientist. *arXiv preprint arXiv:2603.28589*, 2026.
- [223] Y. C. Wu, J. Jian, Z. Zhao, L. Cong, and M. Wang. LabClaw: Operating layer for labos. <https://labclaw-ai.github.io>, 2026.
- [224] Y. Xiang, H. Yan, S. Ouyang, L. Gui, and Y. He. SciReplicate-Bench: Benchmarking LLMs in agent-driven algorithmic reproduction from research papers. *arXiv preprint arXiv:2504.00255*, 2025.
- [225] W. Xu, Y. Zhou, Y. Zhou, Q. Cao, S. Li, J. Bu, B. Liu, Y. Chen, X. He, X. Zhao, X. Zhuang, F. Wang, Z. Zhou, Q. Feng, W. Huang, J. Wei, H. Wu, Y. Yang, G. Wang, S. Xu, Z. Huang, X. Liu, J. Liu, C. Tang, W. Li, Y. Chen, J. Ning, P. Jiang, C. Ma, Y. Du, C. Ji, H. Xu, M. Hu, J. Zheng, X. Chen, Y. Wu, F. Jiang, X. Chen, X. Tang, Y. Fu, Y. Lu, Y. Zhang, L. Sun, C. Li, J. Ma, W. Liu, Y. Liu, K.-C. Wu, S. Chai, Y. Wang, O. Zhangjin, C. Tang, S. Zhang, W. Cao, J. Ren, T. Cui, Z. Yao, J. Deng, Y. Sun, F. Liu, W. Wei, J. Xu, Z. Li, J. Gong, Z. Guo, Z. Yao, Z. Chen, T. Peng, F. Yu, B. Zhang, D. Zhou, S. Tang, J. Liu, F. Ling, Y. Lu, Y. Ren, B. Fei, Z. Zhao, X. Gu, R. Su, X.-M. Wu, W. Si, Y. Liu, H. Chen, X. Yan, X. Yang, J. Yan, J. Wu, Q. Zheng, C. Li, Z. Gao, H. Kong, J. He, M. Su, T. Fu, P. Ye, C. Song, N. Dong, Y. Li, H. Fu, S. Sun, L. Cheng, J. Lin, W. Ouyang, B. Zhou, W. Zhang, and L. Bai. Probing scientific general intelligence of LLMs with scientist-aligned workflows. *arXiv preprint arXiv:2512.16969*, 2025.
- [226] W. Xu, Y. Zhou, Y. Zhou, Q. Cao, S. Li, J. Bu, B. Liu, Y. Chen, X. He, X. Zhao, et al. Probing scientific general intelligence of llms with scientist-aligned workflows. *arXiv preprint arXiv:2512.16969*, 2025. URL <https://arxiv.org/abs/2512.16969>.
- [227] W. Xu, T. Mi, Y. Liu, Y. Nan, Z. Zhou, L. Ye, L. Zhang, Y. Qiao, and P. Liu. ASI-Evolve: AI accelerates AI. <https://github.com/GAIR-NLP/ASI-Evolve>, 2026.
- [228] X. Xu, C. Feng, C. Zha, W. He, M. He, B. Xiao, and X. Gao. ProteinMCP: An agentic AI framework for autonomous protein engineering. *Protein Science*, 35(4):e70547, 2026.

- [229] Y. Yamada, R. T. Lange, C. Lu, S. Hu, C. Lu, J. Foerster, J. Clune, and D. Ha. The AI Scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *arXiv preprint arXiv:2504.08066*, 2025.
- [230] X. Yan, S. Feng, J. Yuan, R. Xia, B. Wang, B. Zhang, and L. Bai. SurveyForge: On the outline heuristics, memory-driven generation, and multi-dimensional evaluation for automated survey writing. *arXiv preprint arXiv:2503.04629*, 2025.
- [231] J. Yang, C. E. Jimenez, A. Wettig, K. Lieret, S. Yao, K. Narasimhan, and O. Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *Advances in Neural Information Processing Systems*, volume 37, pages 50528–50652, 2024.
- [232] M. Yang. ResearchClaw: Personal AI research assistant with extensible skills. <https://github.com/ymx10086/ResearchClaw>, 2025.
- [233] R. Yang, Y. Li, and S. Li. ARIS: Fully autonomous research via adversarial multi-agent collaboration. <https://github.com/wanshuiyin/Auto-claude-code-research-in-sleep>, 2026.
- [234] X. Yang, X. Yang, S. Fang, Y. Zhang, J. Wang, B. Xian, Q. Li, J. Li, M. Xu, Y. Li, H. Pan, Y. Zhang, W. Liu, Y. Shen, W. Chen, and J. Bian. R&D-Agent: An LLM-agent framework towards autonomous data science. *arXiv preprint arXiv:2505.14738*, 2025.
- [235] Y. Yang, W. Jiang, Y. Wang, Y. Song, Y. Wang, and C. Zhang. Auto-Slides: An interactive multi-agent system for creating and customizing research presentations. *arXiv preprint arXiv:2509.11062*, 2025.
- [236] Z. Yang, Z. Zhou, S. Wang, X. Cong, X. Han, Y. Yan, Z. Liu, Z. Tan, P. Liu, D. Yu, Z. Liu, X. Shi, and M. Sun. MatPlotAgent: Method and evaluation for LLM-based agentic scientific data visualization. *arXiv preprint arXiv:2402.11453*, 2024.
- [237] Z. Yang, W. Liu, B. Gao, Y. Liu, W. Li, T. Xie, L. Bing, W. Ouyang, E. Cambria, and D. Zhou. MOOSE-Chem2: Exploring LLM limits in fine-grained scientific hypothesis discovery via hierarchical search. In *Advances in Neural Information Processing Systems*, 2025. URL <https://nips.cc/virtual/2025/poster/118171>.
- [238] Z. Yang, W. Liu, B. Gao, T. Xie, Y. Li, W. Ouyang, S. Poria, E. Cambria, and D. Zhou. MOOSE-Chem: Large language models for rediscovering unseen chemistry scientific hypotheses. In *International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=X9OfMNNepI>.
- [239] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023.
- [240] Y. Yao, H. Zhu, P. Wang, J. Ren, X. Yang, Q. Chen, X. Li, D. Shi, J. Li, Q. Wang, S. Wang, X. Liu, J. Wu, M. Liu, and W. Zhou. O-Researcher: An open ended deep research model via multi-agent distillation and agentic RL. *arXiv preprint arXiv:2601.03743*, 2026.
- [241] J. Ye, Y. Wang, Y. Huang, D. Chen, Q. Zhang, N. Moniz, T. Gao, W. Geyer, C. Huang, P.-Y. Chen, N. V. Chawla, and X. Zhang. Justice or prejudice? quantifying biases in LLM-as-a-Judge. *arXiv preprint arXiv:2410.02736*, 2024. URL <https://arxiv.org/abs/2410.02736>.
- [242] R. Ye, X. Pang, J. Chai, J. Chen, Z. Yin, Z. Xiang, X. Dong, J. Shao, and S. Chen. Are we there yet? revealing the risks of utilizing large language models in scholarly peer review. *arXiv preprint arXiv:2412.01708*, 2024.
- [243] H. Yu, Z. Hong, Z. Cheng, K. Zhu, K. Xuan, J. Yao, T. Feng, and J. You. ResearchTown: Simulator of human research community. In *International Conference on Machine Learning*, pages 73051–73096. PMLR, 2025.
- [244] S. Yu, M. Luo, A. Madasu, V. Lal, and P. Howard. Is your paper being reviewed by an LLM? benchmarking AI text detection in peer review. *arXiv preprint arXiv:2502.19614*, 2025.
- [245] J. Yuan, X. Yan, B. Zhang, T. Chen, B. Shi, W. Ouyang, Y. Qiao, L. Bai, and B. Zhou. Dolphin: Moving towards closed-loop auto-research through thinking, practice, and feedback. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21768–21789, 2025. doi: 10.18653/v1/2025.acl-long.1056. URL <https://aclanthology.org/2025.acl-long.1056/>.
- [246] M. Yuksekgonul, D. Kocejca, X. Li, F. Bianchi, J. McCaleb, X. Wang, J. Kautz, Y. Choi, J. Zou, C. Guestrin, and Y. Sun. Learning to discover at test time. *arXiv preprint arXiv:2601.16175*, 2026.
- [247] S. Zeng, K. Tian, K. Zhang, Y. wang, J. Gao, R. Liu, S. Yang, J. Li, X. Long, J. Ma, B. Qi, and B. Zhou. ReviewRL: Towards automated scientific review with RL. In *Conference on Empirical Methods in Natural Language Processing*, pages 16931–16943, 2025. doi: 10.18653/v1/2025.emnlp-main.857. URL <https://aclanthology.org/2025.emnlp-main.857/>.

- [248] W. Zeng, M. Ouyang, L. Cui, and H. T. Ng. SlideTailor: Personalized presentation slide generation for scientific papers. In *AAAI Conference on Artificial Intelligence*, volume 40, pages 34584–34592, 2026.
- [249] B. Zhang, S. Feng, X. Yan, J. Yuan, R. Ma, Y. Hu, Z. Yu, X. He, S. Huang, S. Hou, Z. Nie, Z. Wang, J. Liu, T. Peng, P. Ye, D. Zhou, S. Zhang, X. Wang, Y. Zhang, M. Li, Z. Tu, X. Yue, W. Ouyang, B. Zhou, and L. Bai. InternAgent: When agent becomes the scientist – building closed-loop system from hypothesis to verification. *arXiv preprint arXiv:2505.16938*, 2025.
- [250] D. Zhang, Z. Bao, S. Du, Z. Zhao, K. Zhang, D. Bao, and Y. Yang. Re2: A consistency-ensured dataset for full-stage peer review and multi-turn rebuttal discussions. *arXiv preprint arXiv:2505.07920*, 2025.
- [251] H. Zhang, H. Cui, Y. Wang, Y. Tian, Q. Guo, C. Wang, J. Wu, C. Song, and Y. Zhang. IterSurvey: Deep literature survey automation with an iterative workflow. *arXiv preprint arXiv:2510.21900*, 2025.
- [252] J. Zhang, S. Hu, C. Lu, R. Lange, and J. Clune. Darwin godel machine: Open-ended evolution of self-improving agents. *arXiv preprint arXiv:2505.22954*, 2025.
- [253] J. Zhang, J. Zhang, Z. Cui, J. Yang, L. Zhang, B. Hui, Q. Liu, Z. Wang, L. Wang, and J. Lin. PlotCraft: Pushing the limits of LLMs for complex and interactive data visualization. *arXiv preprint arXiv:2511.00010*, 2025.
- [254] T. Zhang and H. Sun. SciNav: A general agent framework for scientific coding tasks. *arXiv preprint arXiv:2603.20256*, 2026.
- [255] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. BERTScore: Evaluating text generation with BERT. In *International Conference on Learning Representations (ICLR)*, 2020. URL <https://arxiv.org/abs/1904.09675>.
- [256] T. Zhang, H. Lin, Z. Liu, C. Chen, and W. Zhang. SciFlow-Bench: Evaluating structure-aware scientific diagram generation via inverse parsing. *arXiv preprint arXiv:2602.09809*, 2026.
- [257] Z. Zhang, X. Zhang, J. Wei, Y. Xu, and C. You. PosterGen: Aesthetic-aware multi-modal paper-to-poster generation via multi-agent LLMs. *arXiv preprint arXiv:2508.17188*, 2025.
- [258] B. Zhao, J. Zhang, C. Whitehouse, M. Jiang, M. Shvartsman, A. Charnalia, D. Magka, T. Shavrina, D. Dunfield, O. M. Aodha, and Y. Bachrach. APRES: An agentic paper revision and evaluation system. *arXiv preprint arXiv:2603.03142*, 2026.
- [259] K. Zhao, W. Lin, Q. Zheng, F. Xu, and Y. Li. Deep ideation: Designing LLM agents to generate novel research ideas on scientific concept network. *arXiv preprint arXiv:2511.02238*, 2025.
- [260] X. Zhao, Z. Sang, Y. Li, Q. Shi, W. Zhao, S. Wang, D. Zhang, X. Han, Z. Liu, and M. Sun. AutoReproduce: Automatic AI experiment reproduction with paper lineage. *arXiv preprint arXiv:2505.20662*, 2025.
- [261] Y. Zhao, W. Chen, Z. Xu, M. Patwardhan, C. Wang, Y. Liu, L. Vig, and A. Cohan. AbGen: Evaluating large language models in ablation study design and evaluation for scientific research. In *Annual Meeting of the Association for Computational Linguistics*, pages 12479–12491, 2025.
- [262] H. Zheng, X. Guan, H. Kong, W. Zhang, J. Zheng, W. Zhou, H. Lin, Y. Lu, X. Han, and L. Sun. PPTAgent: Generating and evaluating presentations beyond text-to-slides. In *Conference on Empirical Methods in Natural Language Processing*, pages 14402–14418, 2025.
- [263] H. Zheng, G. Mo, X. Yan, Q. Yuan, W. Zhang, X. Chen, Y. Lu, H. Lin, X. Han, and L. Sun. DeepPresenter: Environment-grounded reflection for agentic presentation generation. *arXiv preprint arXiv:2602.22839*, 2026.
- [264] T. Zheng, Z. Deng, H. T. Tsang, W. Wang, J. Bai, Z. Wang, and Y. Song. From automation to autonomy: A survey on large language models for scientific discovery. *arXiv preprint arXiv:2505.13259*, 2025.
- [265] Y. Zheng, S. Sun, L. Qiu, D. Ru, C. Jiayang, X. Li, J. Lin, B. Wang, Y. Luo, R. Pan, Y. Xu, Q. Min, Z. Zhang, Y. Wang, W. Li, and P. Liu. OpenResearcher: Unleashing AI for accelerated scientific research. In *Conference on Empirical Methods in Natural Language Processing*, pages 209–218, 2024.
- [266] Q. Zhou, Z. Zhang, Z. Li, and L. Sun. “Give a positive review only”: An early investigation into in-paper prompt injection attacks and defenses for AI reviewers. *arXiv preprint arXiv:2511.01287*, 2025.
- [267] C. Zhu, J. Xiong, R. Ma, Z. Lu, Y. Liu, and L. Li. When your reviewer is an LLM: Biases, divergence, and prompt injection risks in peer review. *arXiv preprint arXiv:2509.09912*, 2025.
- [268] D. Zhu, R. Meng, Y. Song, X. Wei, S. Li, T. Pfister, and J. Yoon. PaperBanana: Automating academic illustration for AI scientists. *arXiv preprint arXiv:2601.23265*, 2026.

- [269] M. Zhu, Y. Weng, L. Yang, and Y. Zhang. DeepReview: Improving LLM-based paper review with human-like deep thinking process. *arXiv preprint arXiv:2503.08569*, 2025.
- [270] M. Zhu, Z. Lin, Y. Weng, P. Lu, Q. Xie, Y. Wei, S. Liu, Q. Sun, and Y. Zhang. AutoFigure: Generating and refining publication-ready scientific illustrations. In *International Conference on Learning Representations*, 2026.
- [271] Z. Zhu, K. Q. Lin, and M. Z. Shou. Paper2Video: Automatic video generation from scientific papers. *arXiv preprint arXiv:2510.05096*, 2025.
- [272] Z. Zhuang, J. Chen, H. Xu, Y. Jiang, and J. Lin. Large language models for automated scholarly paper review: A survey. *Information Fusion*, 124:103332, 2025. doi: 10.1016/j.inffus.2025.103332. URL <https://www.sciencedirect.com/science/article/abs/pii/S1566253525004051>.