Article

Synergy, Not Substitution. Responsible Human–AI Collaboration in Academic Research

Marco V. Crivellaro *

Posted Date: 16 September 2025

doi: 10.20944/preprints202509.1249.v1

Keywords: artificial intelligence (AI); academic research; human–AI collaboration; research ethics; epistemic risks; research integrity; governance frameworks; transparency and disclosure; literature review automation; responsible AI adoption

Article

Synergy, Not Substitution. Responsible Human–AI Collaboration in Academic Research

Marco V. Crivellaro

American Community Schools of Athens; crivellarom@acs.gr

Abstract

Artificial intelligence (AI) is rapidly reshaping academic research, offering powerful tools for literature review, data analysis, and knowledge synthesis while raising pressing concerns about reliability, integrity, and ethics. This paper provides an integrative review combining thematic syntheses of scholarly literature with a comparative case analysis of five representative AI tools—Storm, AnswerThis, Coral AI, NotebookLM, and Zotero. The analysis highlights AI's capacity to accelerate research efficiency, broaden access to knowledge, and support collaborative workflows, while also surfacing risks such as fabricated or biased outputs, shallow synthesis, and threats to data privacy. Cross-cutting themes emphasize the importance of transparency, provenance, and human oversight, particularly through practices such as separating generation from validation, and disclosing AI involvement in scholarly outputs. The paper contributes a unified framework situating benefits, risks, and ethics in academic AI use; practical illustrations of human–AI complementarity across diverse tools; policy-relevant insights for governing high- versus low-stakes research applications. The findings converge on a central principle: AI is not a replacement for human scholarship, but a collaborative partner whose outputs require verification, contextualization, and ethical governance. By adopting risk-sensitive, transparent practices, the research community can move beyond polarized debates toward a pragmatic model of responsible human–AI synergy.

Keywords: artificial intelligence (AI); academic research; human–AI collaboration; research ethics; epistemic risks; research integrity; governance frameworks; transparency and disclosure; literature review automation; responsible AI adoption

---

1. Introduction

Artificial intelligence (AI) has become one of the most transformative forces in academic research, promising to reshape how knowledge is produced, validated, and disseminated. From accelerating literature reviews and supporting data analysis to aiding in experimental design, AI tools are increasingly embedded into scholarly workflows (Butson & Spronken-Smith, 2024; Cardona, Rodríguez, & Ishmael, 2023). Major technology companies and academic institutions alike now present AI as a "co-scientist," capable of automating time-consuming processes and freeing human researchers to focus on conceptual and interpretive work (Heikkilä, 2025a; *7 Practical Ways of Integrating AI Tools into Your Research Workflow*, n.d.).

Despite this optimism, deep concerns persist. Commentators have warned that AI may contribute to a "crutch effect," where researchers outsource critical thinking and diminish their own learning (Cummings, 2024; Maker, 2025). Others highlight epistemic risks, including the generation of convincing but flawed outputs that distort the scientific record (Messeri & Crockett, 2024; Brodsky, 2022). The potential for AI to undermine scholarly trust is especially pronounced in domains such as peer review, where fabricated or biased assessments can have long-lasting consequences (Farber, 2025). In parallel, broader societal worries about misinformation and manipulation spill over into academic contexts, where credibility and research integrity are paramount (Maslach et al., 2023; Noor, 2025).

 doi:10.20944/preprints202509.1249.v1

The tension between efficiency and integrity underscores the need for careful examination of AI's place in scholarly work. Researchers increasingly explore frameworks for human–AI complementarity, showing how algorithms and human judgment might be aligned to enhance, rather than replace, one another (Arias-Rosales, 2022; Fragiadakis et al., 2025). At the same time, institutions and policymakers have begun publishing ethical guidelines and governance frameworks to support responsible adoption of generative and analytical AI (Da Veiga, 2025; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025; *New Ethical Framework to Help Navigate Use of AI in Academic Research*, 2024). Yet, the literature remains fragmented, with discussions of benefits, risks, and ethics often occurring in parallel rather than in dialogue.

This paper contributes to bridging that gap by synthesizing existing research into three interrelated dimensions:

1. **Benefits and Opportunities** — the ways AI enhances research efficiency, knowledge discovery, and collaboration.
2. **Risks and Challenges** — the threats to scholarly integrity, learning, and epistemic reliability.
3. **Ethical and Governance Pathways** — the frameworks, guidelines, and principles proposed to regulate responsible use.

By structuring the discussion across these themes, the paper advances a balanced framework for what might be termed "responsible human–AI research synergy." The aim is not only to highlight the current state of knowledge but also to identify tensions, limitations, and future research directions necessary for ensuring that AI supports, rather than undermines, the advancement of science.

2. Literature Review

2.1. Benefits and Opportunities

The potential of AI to accelerate and enhance research is widely acknowledged across disciplines. Studies highlight its role in expediting literature reviews, enabling more efficient knowledge synthesis, and helping researchers identify research gaps with greater precision (Woldetsadik, 2024; Williams, 2025; Parker, 2024; Reimann, 2025). For instance, AI-powered platforms such as SciSpace and Web of Science integrations streamline the once time-intensive task of scanning and synthesizing vast bibliographies (Hennig, 2025; Reimann, 2025).

Beyond information retrieval, AI contributes to productivity gains in designing experiments, analyzing data, and supporting scholarly writing. França (2023) and Madanchian and Taherdoost (2025) argue that AI has the potential to reduce time-to-discovery, freeing human researchers to focus on higher-order reasoning. In a similar vein, Maslach (2023) positions generative AI as a "supercharger" of research efficiency, while institutional reports emphasize its value in creating novel avenues for teaching and learning (Cardona et al., 2023; Heikkilä, 2025a).

Human–AI collaboration is not only about speed but also about complementarity. Early studies of human–AI teaming in problem-solving environments demonstrate how AI can extend human perception and expand the scope of exploration (Arias-Rosales, 2022; Crandall et al., 2018). Recent systematic reviews confirm that hybrid human–AI systems frequently outperform either humans or machines alone, suggesting significant promise in co-creative scientific processes (Vaccaro, Almaatouq, & Malone, 2024; Fragiadakis et al., 2025).

2.2. Risks and Challenges

Despite these benefits, a parallel body of literature emphasizes the risks of AI in research. Critics warn of "doing more but learning less," whereby reliance on AI tools diminishes deep learning and critical thinking (Cummings, 2024; Zhai, Wibowo, & Li, 2024). Maker (2025) labels this the "GenAI crutch," cautioning against premature dependence on AI without sufficient human expertise.

Concerns also extend to the integrity of scientific knowledge. Brodsky (2022) and Maslach et al. (2023) raise alarms about AI's capacity to fabricate convincing but false research outputs, threatening

 doi:10.20944/preprints202509.1249.v1

to erode trust in the scientific process. Messeri and Crockett (2024) expand on this by showing how AI can create illusions of understanding, where researchers overestimate the validity of AI-generated insights.

In peer review, where scholarly standards are enforced, AI introduces further challenges. Farber (2025) highlights the limits of AI-assisted reviewing, suggesting hybrid approaches but noting persistent risks of bias and shallow evaluation. Noor (2025) illustrates the societal consequences of these risks, showing how AI can be exploited to amplify doubt in environmental science. Collectively, these findings point to the epistemic fragility of over-integrating AI into critical evaluative tasks.

2.3. Ethical and Governance Frameworks

In response to these risks, numerous frameworks and guidelines have emerged to regulate responsible AI use in research. Chetwynd (2024), Da Veiga (2025), and Granjeiro et al. (2025) examine ethical implications of AI in scholarly publishing, identifying issues such as authorship, accountability, and transparency. Institutional reports provide more structured frameworks: the University of Oxford (2024) introduced an ethical decision-making framework, while the European Commission (2025) published "Living Guidelines" for responsible generative AI in research. EDUCAUSE (Georgieva et al., 2025) and George Mason University (2025) likewise propose practical guidelines for educators and researchers.

At the practical level, universities and professional bodies have published toolkits and "do's and don'ts" lists for responsible adoption ("The Do's & Don'ts of Using Generative AI Tools Ethically in Academia," 2023; "Effective and Responsible Use of AI in Research," 2024; Upson, 2025). As at the time this is being written, the University of Virginia's Office of Research Integrity (2025) stresses the importance of safeguarding traditional research ethics while integrating new technologies. Collectively, these frameworks reveal a consensus around the need for transparency, disclosure of AI use, and preservation of human accountability.

2.4. Emerging Models of Human–AI Collaboration

A growing area of research focuses on developing conceptual and methodological models for human–AI teaming. Hemmer et al. (2024) and Holstein et al. (2023) explore frameworks of complementarity, where humans and AI each contribute distinct strengths, while Pyae (2025) proposes the "Human-AI Handshake" as a bidirectional model of collaboration.

Systematic reviews provide methodological lenses for evaluating human–AI synergy. Lou et al. (2025) and Fragiadakis et al. (2025) outline how collaboration can be systematically assessed, while Kong et al. (2025) apply hybrid intelligence models to educational environments. Oubou (2024) illustrates how human–AI hybrid workflows are already being adopted in content production and research workflows.

Finally, advances in "learning-to-defer" algorithms (Mao, Mohri, & Zhong, 2021; Raman & Yee, 2021; Paat & Shen, 2025) explore technical solutions to allocate decision-making dynamically between humans and AI. Together, these approaches move beyond binary debates about AI replacement and toward a vision of human–AI partnership optimized for reliability, fairness, and creativity.

3. Methodological Approach

This paper adopts a narrative literature review and thematic synthesis approach to examine the role of artificial intelligence in academic research. Rather than conducting a systematic review with narrowly defined inclusion and exclusion criteria, the aim here is to integrate insights across diverse strands of scholarship, policy guidelines, and commentary. This approach is well-suited for emerging research areas where the literature is heterogeneous and rapidly evolving (Malik & Terzidis, 2025; King's College London Libraries & Collections, n.d.).

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 September 2025 doi:10.20944/preprints202509.1249.v1

3.1. Source Selection

The bibliography underpinning this study was compiled from peer-reviewed journal articles, institutional reports, policy guidelines, preprints, and expert commentaries published between 2018 and 2025. Sources span multiple disciplines, including education, computer science, information systems, and ethics, ensuring a broad capture of perspectives. Emphasis was placed on works that:

1. Evaluate the benefits of AI in research efficiency, collaboration, and knowledge creation (e.g., França, 2023; Madanchian & Taherdoost, 2025).
2. Identify the risks and challenges to epistemic reliability, learning, and research integrity (e.g., Messeri & Crockett, 2024; Brodsky, 2022).
3. Propose or analyze ethical frameworks for responsible AI integration (e.g., Da Veiga, 2025; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025
4. Develop or test models of human–AI collaboration (e.g., Hemmer et al., 2024; Pyae, 2025).

One notable omission from this review is SciSpace, a widely publicized AI platform for literature exploration and PDF interaction. While its visibility in research practice is significant, it was excluded from this analysis for two reasons. First, the present study prioritizes tools that either (a) have a clear research lineage in peer-reviewed evaluations (*Storm*, *NotebookLM*), (b) are explicitly positioned as research-support applications (*AnswerThis*, *Coral AI*), or (c) represent long-established infrastructure with emerging AI extensions (*Zotero*). In contrast, SciSpace has thus far been assessed mainly in product reviews and blog posts, with limited independent scholarly evaluation. Second, the focus of this paper is on tools where transparency, grounding, and provenance can be systematically examined. Although SciSpace offers practical features (semantic search, conversational summaries), its closed retrieval pipeline and limited reproducibility of outputs made it less suitable for a detailed comparative analysis of epistemic and governance concerns. For these reasons, SciSpace was excluded, while acknowledging that it remains a significant platform in the broader ecosystem of AI tools for research and may warrant focused study in future work.

3.2. Analytical Strategy

The analysis proceeded in three stages:

- **Mapping:** References were first categorized into the thematic clusters of benefits and opportunities, risks and challenges, ethical and governance frameworks, and human–AI collaboration models.
- **Thematic Synthesis:** Within each cluster, works were compared to identify converging insights, tensions, and emerging patterns.
- **Integrative Framework:** The findings were then synthesized to highlight cross-cutting themes and to propose a conceptual model for “responsible human–AI research synergy.”

3.3. Limitations

This approach carries several limitations. First, while comprehensive, the bibliography is not exhaustive; relevant studies may have been omitted due to the rapid pace of AI-related research. Second, the inclusion of preprints and commentary introduces variability in peer-review standards. Third, the thematic synthesis approach privileges breadth over depth, which may limit fine-grained disciplinary insights. Nonetheless, the diversity of sources provides a robust foundation for identifying major trends and debates.

3.4. Use of Generative AI tools

Generative AI systems were used in a limited and supervised capacity. ChatGPT-5 (OpenAI, 2025) and Gemini (Google DeepMind, 2025) were employed to assist in refining prose, restructuring draft paragraphs for clarity, and synthesizing thematic connections in the literature. Their outputs were treated as secondary drafts and were cross-checked against the original sources to avoid

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 September 2025 doi:10.20944/preprints202509.1249.v1

fabricated or biased content. STORM (Stanford Open Virtual Assistant Lab, 2024/2025) was tested to scaffold literature review outlines and to generate structured summaries with citations; these outputs were used as prompts for further human-led analysis. Zotero 7.0 with optional AI-enabled plug-ins (Roy Rosenzweig Center for History and New Media, 2025) was used for reference management and bibliographic organization. At all times, the human author maintained responsibility for interpretation, verification, and final editing. No generative outputs were accepted without manual validation against primary literature.

4. Findings

4.1. Impact on Research Efficiency

Across disciplines, AI tools compress time-to-insight by automating search, screening, summarization, and triage for literature reviews and evidence synthesis (Hennig, 2025; Reimann, 2025; Parker, 2024; Williams, 2025; Woldetsadik, 2024; Wilkins, 2023). Empirical and conceptual accounts attribute these gains to automation of routine cognitive labor and augmentation of pattern recognition—freeing researchers for hypothesis generation and interpretation (França, 2023; Madanchian & Taherdoost, 2025; Lund, Lamba, & Oh, 2024). Institutional and media reports emphasize accelerating components of the research pipeline and even positioning AI as a "co-scientist" that supports experimental design and discovery (Cardona, Rodríguez, & Ishmael, 2023; Heikkilä, 2025a).

Importantly, the efficiency dividend is task- and field-dependent: gains concentrate where literatures are well-indexed and tasks are pattern-rich; they taper in sparse or idiosyncratic domains (Reimann, 2025; Hennig, 2025). Guidance from universities and graduate schools stresses pairing speed with traceability and disclosure—logging tool use, preserving provenance, and instituting human verification loops (*7 Practical Ways of Integrating AI Tools into Your Research Workflow*, n.d.; "Effective and Responsible Use of AI in Research," 2024).

4.2. Human–AI Synergy in Research Tasks

A growing body of work shows hybrid teams (humans + AI) can outperform either alone when collaboration is designed around complementary strengths (Vaccaro, Almaatouq, & Malone, 2024; Fragiadakis et al., 2025; Lou, Lu, Raghu, & Zhang, 2025). Conceptual and HCI frameworks detail where complementarity arises—capability differences, information asymmetries, and perspective diversity—and how to measure success beyond accuracy (e.g., calibration, workload, trust appropriateness) (Hemmer et al., 2024; Holstein, De-Arteaga, Tumati, & Cheng, 2023; Jiang, Sun, Fu, & Lv, 2024).

Task-level studies illustrate what AI adds: expanding early design exploration and perceptual reach (Arias-Rosales, 2022), enabling cooperative strategies (Crandall et al., 2018), and scaffolding reflection on "unobservables" that humans miss (Holstein et al., 2023). In educational and research settings, structured teaming improves problem solving when roles and feedback loops are explicit (Kong et al., 2025; Tong et al., 2025).

On the algorithmic side, learning-to-defer formalizes *who decides when*: systems route instances to humans or models based on uncertainty and expertise, now extended to multi-expert settings with conformal guarantees (Mao, Mohri, & Zhong, 2021; Raman & Yee, 2021; Paat & Shen, 2025). Alignment work further stabilizes collaborative behavior via instruction-following with human feedback (Ouyang et al., 2022). Design principles recur: make rationales legible, surface uncertainty, and preserve human authority over consequential judgments (Vössing, Kühl, Lind, & Satzger, 2022; Hemmer et al., 2024).

4.3. Emerging Risks

The most frequently cited epistemic risk is over-confidence—humans over-trusting fluent AI outputs that *appear* sound (Messeri & Crockett, 2024). In learning contexts, reviews document over-reliance and cognitive offloading to dialogue systems, with potential erosion of deep learning (Zhai, Wibowo, & Li, 2024) and warnings about the "GenAI crutch" in professional settings (Maker, 2025). Surveys suggest researchers welcome AI as a teammate but resist ceding authority, revealing trust asymmetries that can distort collaboration if unaddressed (Kelly, 2025).

The integrity threat surface spans peer review and scientific communication: plausible-sounding but low-substance reviews, bias amplification, and targeted misuse to sow doubt in contested domains (Farber, 2025; Noor, 2025). Commentary has long warned of fabrication pipelines that could "wreck the scientific process" if editorial safeguards lag behind generation capabilities (Brodsky, 2022; Maslach, ChatGPT Robot, Doshi, Mors, Puranam, & Seamans, 2023). In risky decision-making contexts, human factors engineering/management work recommends constraint-driven handoffs, provenance tracking, and explicit uncertainty displays to prevent automation surprises (Xiong, Fan, Ma, & Wang, 2022). University integrity offices echo disclosure, verification, and audit trails as minimal conditions for responsible use (University of Virginia Office of Research Integrity, n.d.)

4.4. Ethical and Governance Landscape

Ethics and policy sources converge on three pillars: transparency/disclosure, accountability/authorship, and proportional human oversight. Publishing-oriented analyses synthesize journal practices for acknowledging AI involvement and clarifying responsibility (Chetwynd, 2024; Da Veiga, 2025; Granjeiro et al., 2025).

Institutional frameworks translate principles into actionable governance. The European Commission's *Living Guidelines* provide modular, updatable recommendations for research settings; the University of Oxford outlines decision frameworks for researchers; EDUCAUSE and George Mason University provide operational checklists for departments and labs (*Living Guidelines on the Responsible Use of Generative AI in Research*, 2025; *New Ethical Framework to Help Navigate Use of AI in Academic Research*, 2024; Georgieva et al., 2025; George Mason University, 2025). Practitioner resources ("The Do's & Don'ts of Using Generative AI Tools Ethically in Academia," 2023; "Effective and Responsible Use of AI in Research," 2024 Upson, 2025) emphasize procedural safeguards: declare prompts and models used, preserve intermediate artifacts, double-source claims, and run human verification before publication.

4.5. Cross-Cutting Patterns and Tensions

The previous sections have highlighted how AI transforms research efficiency, collaboration, risks, and governance. Yet these insights do not exist in isolation: when considered together, they reveal cross-cutting tensions that shape the conditions for responsible human–AI synergy. Synthesizing across the literature, four recurrent patterns emerge.

Efficiency–Integrity Trade-off. While AI accelerates early-stage research tasks, this very speed risks compromising the depth and rigor of scholarly scrutiny unless accompanied by systematic verification and provenance-tracking practices (Reimann, 2025; "Effective and Responsible Use of AI in Research," 2024).

Complementarity over Substitution. Evidence consistently shows that durable gains arise not from replacing human researchers but from workflows that exploit complementary strengths, including structured deferral to humans in cases of uncertainty or value-laden judgment (Vaccaro et al., 2024; Mao et al., 2021; Holstein et al., 2023).

Governance as Enablement. Far from being merely restrictive, ethical guidelines and institutional frameworks play an enabling role by clarifying responsibilities around disclosure, authorship, and oversight, thus lowering barriers for safe experimentation and adoption (*Living Guidelines on the Responsible Use of Generative AI in Research*, 2025; Georgieva et al., 2025).

Skills Shift. The rapid proliferation of tools shifts the emphasis of researcher expertise: beyond disciplinary knowledge, scholars must now cultivate prompt literacy, critical evaluation of AI outputs, uncertainty reasoning, and meta-analytic validation. Whether efficiency gains translate into higher-quality scholarship depends on training practices and evolving lab norms (Hennig, 2025; Upson, 2025).

Taken together, these patterns frame the central challenge of AI in research: to balance speed with rigor, harness complementarity, embed governance as a scaffold for responsible innovation, and equip researchers with new literacies. The next subsection (4.6) builds on these tensions to translate them into practical implications for research workflows, offering concrete strategies for implementing responsible human–AI collaboration.

4.6. Practical Implications for Research Workflows

Building on these cross-cutting tensions, it becomes clear that translating principles into practice requires intentional workflow design. The efficiency–integrity trade-off, the emphasis on complementarity, the enabling role of governance, and the emerging skills shift all point toward concrete adjustments in how research pipelines are organized. The following recommendations operationalize these insights, offering practical strategies that align tool use with transparency, accountability, and human oversight.

Translating these patterns into daily research practice requires deliberate adjustments across the workflow. One foundational step is to **i**nstrument the research pipeline so that prompts, model versions, and checkpoints are carefully recorded, alongside human review logs. This practice not only enhances accountability and reproducibility but also aligns with governance frameworks that stress transparency and traceability (*Living Guidelines on the Responsible Use of Generative AI in Research*, 2025; Georgieva et al., 2025).

Equally important is the adoption of deferral protocols, whereby routine and low-stakes tasks can be automated with confidence, while novel, uncertain, or high-impact cases are escalated to human researchers. By triaging tasks according to uncertainty and risk, teams can preserve efficiency without sacrificing rigor—a principle well illustrated in recent work on "learning to defer" approaches (Mao et al., 2021; Raman & Yee, 2021; Paat & Shen, 2025).

Transparency in design is another critical pillar. Interfaces that surface rationales and uncertainties, or that deliberately prompt users to reflect on "unobservables," foster more trustworthy and collaborative engagement between humans and AI (Vössing et al., 2022; Holstein et al., 2023). Rather than treating AI as a black box, these designs create opportunities for human oversight and correction, which are essential for maintaining epistemic integrity.

At the validation stage, it is advisable to separate generation from verification. In practice, this means that different individuals—or at least different passes—should be responsible for drafting versus checking, with a requirement that AI-generated claims be double-sourced before inclusion in scholarly work. Such division of labor resonates strongly with institutional guidance, which frames disclosure and independent validation as minimal conditions for responsible adoption ("Effective and Responsible Use of AI in Research," 2024; "The Do's & Don'ts of Using Generative AI Tools Ethically in Academia," 2023).

Finally, researchers must disclose AI involvement in line with journal and institutional policies. Acknowledging the role of generative systems in producing drafts, summaries, or references not only fulfills ethical obligations around authorship and accountability but also ensures that peers can properly evaluate potential sources of bias or error (Chetwynd, 2024; Da Veiga, 2025; *New Ethical Framework to Help Navigate Use of AI in Academic Research*, 2024).

Together, these practices demonstrate that responsible integration of AI into research is not a matter of prohibition or uncritical embrace, but of careful workflow design. By embedding transparency, accountability, and human oversight into every stage, researchers can harness the efficiency gains of AI while safeguarding the integrity and reproducibility of scholarship.

5. Comparative Case Analysis: Five AI Tools in Research Workflows

While the preceding sections synthesized findings across the literature, it is equally important to examine how specific tools embody these dynamics in practice. This section presents comparative case analyses of five representative tools—Storm (Stanford), AnswerThis, Coral AI, NotebookLM, and Zotero—to illustrate how real products operationalize the benefits, risks, and governance considerations discussed earlier. For each, we examine typical workflows, distinctive strengths, design limitations, and policy/ethics implications.

5.1. Storm: structured, cited, Wikipedia-style synthesis

What it is. Storm is a Stanford research system for generating grounded, Wikipedia-like articles by explicitly modeling the pre-writing stage: it retrieves diverse sources, poses multi-perspective questions, simulates a writer–expert dialog to curate an outline, and then drafts a report with citations. A public demo and open-source code exist; the project also explores a collaborative variant (Co-STORM). (Shao et al., 2024; *STORM*, n.d.; *Stanford-Oval/Storm*, 2024/2025)

Strengths. Storm's emphasis on process transparency (outline + evidence before prose) directly targets common pain points in long-form synthesis: scoping breadth, structuring subsections, and surfacing references inline. The accompanying paper documents measurable gains on *FreshWiki* (e.g., editors judged Storm outputs +25% more organized and +10% broader vs. a baseline), while also identifying realistic failure modes (e.g., source-bias transfer). (Shao et al., 2024)

Limitations & governance. As a research prototype, coverage and stability can vary; prose can be template-like and still requires human editing. Storm does not eliminate the need to open sources and verify claims—particularly where the retrieval stack or perspective selection may skew coverage. For scholarly use, pair Storm with a provenance log (saved outline, source list) and standard AI-use disclosure consistent with the EU's *Living Guidelines* (traceability, human accountability). (Shao et al., 2024; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025)

5.2. AnswerThis: Question-Driven Literature Scoping with Surfaced Citations

What it is. AnswerThis is a hosted research assistant positioned for fast, conversational querying over a large academic corpus; the platform advertises ~200–250M papers indexed, direct citations in responses, and workflow features (bibliometrics, structured literature reviews, BibTeX/RIS export, and a data-handling guide). These claims are platform-reported and should be verified for your use case. (AnswerThis, n.d.; *AnswerThis Data Handling Guide*, n.d.)

Strengths. The platform's citation-forward UI lowers the friction of exploratory Q&A: it can quickly assemble a seed set of candidate papers, map topical connections, and export references for downstream work in a manager like Zotero. For research teams, it's well-suited to the earliest scoping pass, when the aim is breadth and hypothesis surfacing rather than fine-grained appraisal. (AnswerThis, n.d.)

Limitations & governance. Like all aggregator-LLM stacks, AnswerThis can produce shallow synthesis and occasional misattributed or incomplete citations; its own data-handling guide advises users to verify outputs against the source PDFs/HTML. Treat coverage, user counts, and accuracy claims as marketing until validated in context. Recommended practice: (1) open and check every cited source; (2) keep a verification log; (3) disclose tool use per institutional policy. (*AnswerThis Data Handling Guide*, n.d., *FAQs for the Policy on Generative AI in Research*, n.d.; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025)

5.3. Coral AI: Document-Centric Assistant with Per-Answer Citations

What it is. Coral AI is a closed-corpus assistant: you upload PDFs, notes, and even meeting audio/video; it returns answers grounded in your files with page-level, clickable citations, plus utilities for mind maps and presentations. The provider emphasizes "a citation with every response" and multi-file chat ("100+ files at once"). (*Coral AI*, n.d.))

 doi:10.20944/preprints202509.1249.v1

Strengths. For lab groups, Coral excels at internal corpus triage—rapid cross-document Q&A, traceable snippets, and quick communication artifacts (maps, slide decks). Because it relies on your materials, it avoids some pitfalls of open-web retrieval and is strong for compliance-sensitive workflows if access controls and storage policies are configured properly. (*Coral AI*, n.d.)

Limitations & governance. Risks concentrate around privacy/IP (uploaded content, retention, access to shared workspaces). Even with page-level citations, users should spot-check against the original PDFs and align usage with data-management plans and IRB requirements. At minimum: enforce role-based access, maintain versioned exports with document manifests, and include AI-use acknowledgments in downstream outputs per EU guidance. (*Coral AI*, n.d.; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025)

5.4. NotebookLM (Google): Source-Grounded Notebooks with Multimodal Overviews

What it is. NotebookLM is Google's source-grounded research/learning notebook: users load their own sources (PDFs, Google Docs/Slides, web links, YouTube, etc.) and chat to get answers with inline citations. Recent updates add Video Overviews, a redesigned Studio that stores multiple outputs, and Mind Maps; there's expanding support for public/featured notebooks for sharing. (*Learn about NotebookLM*, n.d.; Wang & Bin Shafqat, 2025; Diaz, 2025)

Strengths. NotebookLM's grounding to user-provided sources plus multimodal outputs (audio/video overviews, mind maps) make it a strong fit for teaching, lab onboarding, and team synthesis. Google's help pages explicitly emphasize clear inline citations to support accuracy and transparency. (*Learn about NotebookLM*, n.d.)

Limitations & governance. Overviews can shape interpretation—helpful for comprehension, but a risk for framing bias if users don't open the sources. Teams should manage sharing scope (public vs. private notebooks), track source manifests, and include AI-use statements on derived materials. As with other assistants, critical claims require human verification against primary texts. (Wang & Bin Shafqat, 2025, *Learn about NotebookLM*, n.d.)

5.5. Zotero: Evidence Backbone with Optional AI via Plugins

What it is. Zotero is a free, open-source reference manager and evidence hub with desktop/mobile apps, a modernized Zotero 7 UI, built-in PDF/ePUB readers, and robust syncing. While Zotero itself isn't generative, its plugin ecosystem (e.g., *Zotero-GPT*, *PDF-Translate*, ARIA) enables optional AI-adjacent workflows. By default, your library lives locally (with optional sync). (*Zotero. About*, n.d.; *Zotero*, n.d.)

Strengths. Zotero is uniquely strong on provenance and auditability: you own the PDFs/notes, can export annotations, and can trace every citation back to a library item. Its open architecture (connectors, "translators," group libraries) makes it the evidence backbone that other tools feed into; AI add-ons are opt-in. (*Zotero. About*, n.d.)

Limitations & governance. Generative features depend on third-party plugins of varying maturity, so labs should vet add-ons (security, maintenance) and document versions/models used. For sensitive projects, confirm sync/storage settings (local vs. cloud) and follow institutional data policies; Zotero's own privacy docs outline server locations and third-party requests used for metadata retrieval. ((*Zotero. About*, n.d.)

5.6. Cross-Tool Patterns and Implications

- **Grounding & traceability.** Storm and NotebookLM make source grounding explicit in their designs; Coral promises page-level citations; AnswerThis surfaces citations but still requires manual PDF checks; Zotero anchors the whole pipeline with first-party custody of sources. These patterns align with governance guidance emphasizing transparency, provenance, and human verification. (Shao et al., 2024; *Living Guidelines on the Responsible Use of Generative AI in Research*, 2025)

 doi:10.20944/preprints202509.1249.v1

- **Best-fit use cases.** Storm: scaffolded topic briefs with documented research steps; AnswerThis: rapid scoping across large literatures; Coral: closed-corpus analysis and artifact generation; NotebookLM: teamable synthesis and teaching artifacts; Zotero: reference & annotation backbone for reproducibility.
- **Risk hot-spots.** Over-templated overviews (Storm), hallucinated/misattributed citations (AnswerThis), privacy/IP exposure on uploads (Coral/NotebookLM), and plugin variability (Zotero). Mitigations: open the source, keep verification logs, set access controls, and disclose AI involvement per EU/Oxford policies.

5.7. The Emerging Norm of Citing AI Tools in Research

As AI systems become increasingly embedded in research workflows, a pressing issue is how their use should be acknowledged in scholarly communication. Unlike traditional software packages which have long been routinely cited in methods sections, generative and assistant-style AI tools blur the boundary between infrastructure and authorship. Their influence can extend beyond mechanical support to shaping interpretation, synthesis, and even phrasing.

Recent guidelines from publishers and research bodies emphasize transparency in AI use. For example, *Science Editing* and *COPE* recommend that AI tools may not be listed as authors but should be cited or acknowledged when they contribute substantively to a manuscript's development (Da Veiga, 2025; Granjeiro et al., 2025). The European Commission's *Living Guidelines* (2025) and the University of Oxford (2024) likewise stress disclosure of generative AI involvement as a matter of research integrity. Failing to disclose such use risks both reputational harm and epistemic opacity: readers may overestimate the originality of interpretations or underestimate the automation embedded in findings.

The comparative case analysis reinforces this imperative. For instance, if Storm contributes to outlining or scoping, or AnswerThis generates seed references, these contributions should be documented in the methods section. Similarly, when Coral AI or NotebookLM are used to summarize source materials, or when Zotero plug-ins assist in annotation, such use should be cited to provide an audit trail. Explicit citation practices make workflows reproducible, allow peers to evaluate potential biases, and align with established scholarly norms around tool acknowledgment.

Going forward, researchers and journals alike will need to standardize citation formats for AI tools, possibly drawing from existing models for citing software and datasets (e.g., DOI-based software citations, GitHub release citations). Such practices will not only preserve transparency but also contribute to building a shared understanding of how AI tools are shaping knowledge production across disciplines.

Example: Methods Citation Template for AI Tools

When acknowledging the use of AI tools in research, a clear statement in the Methods section should:

1. Name the tool and version used.
2. Specify the role it played (e.g., literature scoping, summarization, reference management).
3. Provide a formal citation (to the tool's website, developer documentation, or DOI if available).
4. State limitations (e.g., human verification of outputs, disclosure of AI use).

Sample Wording in Methods Section

- "We used Storm (v.2023.10), a Stanford research prototype for AI-assisted synthesis, to generate an initial outline and candidate references for our literature review. All AI-generated outputs were subsequently verified against original sources."
- "Preliminary literature scoping was supported by AnswerThis (*https://answerthis.ai*), which produced conversational summaries with surfaced citations. References identified through the tool were manually checked against source PDFs before inclusion."

 doi:10.20944/preprints202509.1249.v1

- "Uploaded source materials were summarized with NotebookLM (Google, 2025) to generate Audio and Video Overviews. These outputs were used as secondary aids for comprehension and were verified by the authors against the original documents."
- "For reference management, we used Zotero 7.0 (Roy Rosenzweig Center for History and New Media, 2024) with the Zotero-GPT plugin to assist in annotation and PDF summarization. Only human-verified notes were included in analysis."

Sample APA-style References

- Stanford HAI. (2023). *Storm: Structured, grounded report generation system* [Computer software]. Stanford University. https://storm.genie.stanford.edu/
- AnswerThis. (2024). *AnswerThis: Conversational literature assistant* [Web application]. https://answerthis.ai
- Google. (2025). *NotebookLM* [Web application]. https://notebooklm.google/
- Coral AI. (2024). *Coral AI: Document-centric assistant* [Web application]. https://getcoralai.com/
- Zotero. (2024). *Zotero (Version 7.0)* [Computer software]. Roy Rosenzweig Center for History and New Media. https://www.zotero.org

6. Discussion

6.1. Synthesizing Benefits and Risks

The literature review and findings highlight the dual-use character of AI in academic research: accelerating knowledge work while simultaneously introducing risks to integrity. AI systems clearly deliver productivity gains in literature review, data analysis, and research communication (França, 2023; Hennig, 2025; Reimann, 2025). However, these gains are consistently offset by epistemic risks such as overconfidence, diminished learning, and the propagation of fabricated or biased outputs (Messeri & Crockett, 2024; Zhai et al., 2024; Brodsky, 2022).

The case studies in Section 6 exemplify this trade-off. Storm demonstrates how structured synthesis with transparent pre-writing scaffolds can improve research efficiency, but still produces templated prose that requires verification. AnswerThis highlights the convenience of conversational scoping with citations, while also exposing the risk of shallow synthesis and hallucinated references. These tools thus reinforce the need for generation–validation separation: AI may accelerate the assembly of candidate sources or drafts, but human review remains indispensable.

6.2. Human–AI Complementarity in Practice

The evidence suggests that the greatest value emerges not from automation alone, but from designed complementarity between human and machine. As hybrid-intelligence studies show, distributed decision-making improves accuracy and breadth (Vaccaro et al., 2024; Mao et al., 2021). Case illustrations mirror this principle: Coral AI excels in closed-corpus analysis by grounding answers in user-uploaded materials, yet requires human oversight to ensure data privacy and correct interpretation. NotebookLM demonstrates how multimodal summaries (audio, video, mind maps) can scaffold understanding and collaboration, but also how AI framing can subtly steer interpretations if unchecked.

These examples underscore a broader point: collaboration models matter. Tools that foreground transparency, grounding, and clear attribution (Storm, NotebookLM, Zotero) better support human–AI synergy than those optimized purely for speed (AnswerThis).

6.3. Ethical and Governance Imperatives

Across both the literature and the tool examples, ethics and governance frameworks emerge as enablers rather than obstacles. Institutions and policymakers stress transparency, disclosure, and accountability (*Living Guidelines on the Responsible Use of Generative AI in Research*, 2025; *New Ethical Framework to Help Navigate Use of AI in Academic Research*, 2024). The case analyses illustrate why:

 doi:10.20944/preprints202509.1249.v1

Zotero, long a cornerstone of reference management, remains valuable precisely because it provides an auditable evidence backbone that aligns with governance principles of provenance and traceability. By contrast, newer tools such as Coral AI and NotebookLM, which require document uploads or collaborative sharing, raise urgent questions about data privacy and IP governance.

The comparison suggests that best practice should follow a risk-based approach: low-stakes exploratory uses may tolerate lighter oversight, but high-stakes tasks (peer review, policy advising, student assessment) demand stricter controls, provenance documentation, and explicit human accountability.

A further implication emerging from both the literature and the tool cases is the need to explicitly cite AI tools in scholarly outputs. Just as statistical packages (e.g., SPSS, R) and reference managers (e.g., Zotero) are routinely cited in methods sections, the use of generative and assistant-style AI tools should also be documented. This practice aligns with guidelines from the European Commission (2025), the University of Oxford (2024), and scholarly publishing bodies (Da Veiga, 2025; Granjeiro et al., 2025), all of which stress disclosure of AI involvement as a condition of research integrity. Failing to cite such tools risks epistemic opacity—readers cannot evaluate the reliability of claims if the role of AI in shaping sources, interpretations, or phrasing is hidden. The case studies underscore this point: if Storm contributes to structuring an outline, or AnswerThis to identifying seed references, such use should be reported; if NotebookLM or Coral AI assist with summarization, their contribution should be acknowledged; and when Zotero's AI-enabled plug-ins are used for annotation, these should also be disclosed. The adoption of standardized AI citation practices would help normalize transparency and ensure reproducibility in academic research.

6.4. Contributions of this Paper

By integrating a thematic synthesis of academic literature with a comparative case analysis of widely used AI tools, this paper makes three contributions:

1. An integrative framework situating benefits, risks, and ethics within a single analysis.
2. Practical grounding through real-world examples, showing how principles such as complementarity, transparency, and verification apply across distinct classes of tools (e.g., synthesis engines, Q&A assistants, collaborative platforms, reference managers).
3. Policy-relevant implications, highlighting how governance frameworks should differentiate between tools that are exploratory and those embedded in high-stakes research workflows.

6.5. Limitations and Future Directions

This review is limited by the inclusion of preprints and platform-reported claims, which may not always be peer-reviewed or empirically validated. Moreover, tools evolve rapidly; features described in Section 6 may change within months. Future research should therefore include longitudinal studies of tool adoption in labs and classrooms, comparative benchmarking of AI accuracy and reliability across platforms, and policy analyses of how institutional guidelines are being implemented in practice.

6.6. Toward Responsible Human–AI Research Synergy

Both the literature and the case illustrations converge on a central theme: AI should be framed not as a replacement for human scholarship but as a collaborative partner, one whose outputs demand verification, contextualization, and ethical oversight. Whether in the templated outputs of Storm, the conversational summaries of AnswerThis, the collaborative affordances of Coral AI, the multimodal synthesis of NotebookLM, or the evidence management backbone of Zotero, the lesson is the same: responsible adoption requires transparency, provenance, and human accountability.

By aligning tool use with these principles, the research community can move beyond polarized debates about AI as threat or panacea and toward a pragmatic model of responsible human–AI synergy—a model that sustains both the pace and the integrity of scientific discovery.

7. Conclusions

Artificial intelligence is rapidly reshaping academic research, offering clear benefits while also raising complex risks. As the literature synthesis demonstrated, AI tools can accelerate literature reviews, streamline data analysis, and scaffold scholarly writing, yet they also risk generating shallow synthesis, hallucinated outputs, and epistemic overconfidence. The efficiency–integrity trade-off remains a central tension: productivity gains must not come at the expense of scholarly rigor.

The case analyses of Storm, AnswerThis, Coral AI, NotebookLM, and Zotero provide concrete illustrations of these dynamics. Storm exemplifies how structured synthesis can improve efficiency while still requiring careful human verification. AnswerThis demonstrates the promise of rapid question-driven exploration, alongside the risks of shallow or incorrect citations. Coral AI and NotebookLM highlight how grounding to user-provided sources can enhance collaboration and reflection, but also bring new challenges around data privacy and framing bias. Finally, Zotero underscores the enduring importance of provenance, transparency, and researcher control—qualities that align closely with emerging governance frameworks. This paper deliberately excluded SciSpace, despite its prominence as an AI platform for literature exploration and PDF interaction. As noted in the methodology, the exclusion was based on its limited evaluation in peer-reviewed scholarship and the difficulty of systematically assessing transparency and reproducibility in its closed retrieval pipeline. While this choice kept the comparative analysis focused on tools with clearer documentation and research lineage, it also represents a limitation: SciSpace is widely used in practice and may influence researcher workflows in ways not captured here. Future work should therefore include dedicated studies of SciSpace and similar commercial platforms to ensure that their impact is assessed alongside more openly documented systems.

Taken together, the findings above suggest that responsible AI use in research must follow three guiding principles:

1. **Complementarity, not substitution** — AI should support, not replace, human expertise.
2. **Transparency and provenance** — every AI-assisted output must be traceable to its sources.
3. **Human accountability** — researchers retain final responsibility for interpretation, validation, and ethical disclosure.

For researchers, this means using AI as a collaborator while maintaining critical oversight. For institutions and journals, it means developing enforceable disclosure policies, training programs in prompt and critique literacy, and infrastructure for provenance tracking. For policymakers, it means harmonizing fragmented frameworks into coherent governance structures that encourage innovation without compromising integrity.

A final implication is the emerging norm of citing AI tools in research outputs. Just as methods sections routinely document software, databases, and analytical frameworks, so too should the use of systems such as Storm, AnswerThis, Coral AI, NotebookLM, or Zotero (with AI plug-ins) be acknowledged. Standardizing these citation practices will not only ensure transparency and reproducibility but also normalize responsible human–AI collaboration as an accepted component of the research process.

Ultimately, the future of AI in academic research will be shaped less by the tools themselves than by the norms, practices, and policies surrounding their use. The examples analyzed here show that AI can indeed accelerate discovery and foster creativity—but only if embedded within workflows that preserve the epistemic integrity of scholarship. Cultivating such responsible human–AI synergy is essential for ensuring that the next phase of research is both faster and more trustworthy.

Acknowledgments: The author acknowledges the use of the following generative AI tools in the preparation of this manuscript: ChatGPT-5 (OpenAI, 2025), used for language refinement and thematic synthesis; Gemini (Google DeepMind, 2025), used for structural editing and draft summarization; STORM (Stanford Open Virtual

Assistant Lab, 2024/2025), used for generating structured outlines and exploratory synthesis; Zotero (Version 7.0, Roy Rosenzweig Center for History and New Media, 2025), including AI plug-ins for annotation and reference management.
All outputs were independently reviewed and validated by the author. Responsibility for accuracy and interpretation rests solely with the human contributor.

References

4. 7 Practical Ways of Integrating AI Tools into Your Research Workflow. (n.d.). Stanford Medicine. Retrieved August 21, 2025, from https://lanecal.stanford.edu/event/14875184
5. 46 AI tools for research: How to choose the best tool for your needs. (2025, June 16). Zendy. https://zendy.io/blog/ai-tools-for-research-comprehensive-list
6. AI Guidelines for Research. (n.d.). George Mason University. Retrieved August 21, 2025, from https://www.gmu.edu/ai-guidelines/ai-guidelines-research
7. AI Tools for Literature Review in 2025: What's Working and What's Not. (2025, May 19). Elite Academic Brokers. https://www.linkedin.com/pulse/ai-tools-literature-review-2025-whats-working-olkpf
8. Almanza-Ferro, A. (n.d.). Using AI tools in Research: Ethical AI. Retrieved August 21, 2025, from https://guides.library.tamucc.edu/AI/ethicalAI
9. AnswerThis. (n.d.). AnswerThis. AnswerThis. Retrieved August 22, 2025, from https://answerthis.io
10. AnswerThis Data Handling Guide. (n.d.). AnswerThis. Retrieved August 22, 2025, from https://answerthis.io/docs/data_privacy/
11. Apata, O. E., Kwok, O.-M., & Lee, Y.-H. (2025). The Use of Generative Artificial Intelligence (AI) in Academic Research: A Review of the Consensus App. Cureus. https://doi.org/10.7759/cureus.87297
12. Arangüena, I. (2024, April 5). AI in Research: Its Uses and Limitations. Research to Action. https://www.researchtoaction.org/2024/04/ai-in-research-its-uses-and-limitations/
13. Arias-Rosales, A. (2022). The perceived value of human-AI collaboration in early shape exploration: An exploratory assessment. PLOS ONE, 17(9), e0274496. https://doi.org/10.1371/journal.pone.0274496
14. Bouchrika, I. (2024, November 27). The Role of AI in Academic Research: Benefits and Ethical Considerations. Research.Com. https://research.com/research/the-role-of-ai-in-academic-research
15. Bridgers, S., Jain, R., Greig, R., & Shah, R. (2025, June 4). Human-AI Complementarity: A Goal for Amplified Oversight. Medium. https://deepmindsafetyresearch.medium.com/human-ai-complementarity-a-goal-for-amplified-oversight-0ad8a44cae0a
16. Brodsky, S. (2022, December 1). How AI Could Fake Papers and Wreck the Scientific Process. Lifewire. https://www.lifewire.com/how-ai-could-fake-papers-and-wreck-the-scientific-process-6834621
17. Butler, J., Jaffe, S., Baym, N., Czerwinski, M., Iqbal, S., Nowak, K., Rintel, R., Sellen, A., Vorvoreanu, M., Hecht, B., & Teevan, J. (Eds.). (2023). Microsoft New Future of Work Report 2023 (No. MSRTR-2023-34). Microsoft Research. https://aka.ms/nfw2023
18. Butson, R., & Spronken-Smith, R. (2024). AI and its implications for research in higher education: A critical dialogue. Higher Education Research & Development, 43(3), 563–577. https://doi.org/10.1080/07294360.2023.2280200
19. Cardona, M. A., Rodríguez, R. J., & Ishmael, K. (2023). Artificial Intelligence and the Future of Teaching and Learning, Insights and Recommendations. U.S. Department of Education, Office of Educational Technology. https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf
20. Chetwynd, E. (2024). Ethical Use of Artificial Intelligence for Scientific Writing: Current Trends. Journal of Human Lactation, 40(2), 211–215. https://doi.org/10.1177/08903344241235160
21. Chust, E. (2025, March 5). STORM AI: The Ultimate Tool for Content Research and Curation. https://www.cyberclick.net/numericalblogen/storm-ai
22. Coral AI: The AI Assistant For Your Documents. (n.d.). Coral AI. Retrieved August 22, 2025, from https://www.getcoralai.com/
23. Crandall, J. W., Oudah, M., Tennom, Ishowo-Oloko, F., Abdallah, S., Bonnefon, J.-F., Cebrian, M., Shariff, A., Goodrich, M. A., & Rahwan, I. (2018). Cooperating with machines. Nature Communications, 9(1), 233. https://doi.org/10.1038/s41467-017-02597-8

24. Cummings, M. (2024, March 7). Doing more, but learning less: The risks of AI in research. Yale News. https://news.yale.edu/2024/03/07/doing-more-learning-less-risks-ai-research
25. Da Veiga, A. (2025). Ethical guidelines for the use of generative artificial intelligence and artificial intelligence-assisted tools in scholarly publishing: A thematic analysis. Science Editing, 12(1), 28–34. https://doi.org/10.6087/kcse.352
26. Diaz, N. (2025, July 14). Google stocks “featured notebooks” on NotebookLM shelves for curated topics. Android Central. https://www.androidcentral.com/apps-software/ai/google-stocks-featured-notebooks-notebooklm-curated-research-topics-rollout
27. Durgumahanthi, S. (2024, December 11). What is Scispace? Detailed Review of Features, Pricing, and Alternatives. Paperpal. https://paperpal.com/blog/news-updates/scispace-review-features-pricing-and-alternatives
28. Effective and Responsible Use of AI in Research. (2024, December 10). UW Graduate School. https://grad.uw.edu/advice/effective-and-responsible-use-of-ai-in-research/?utm_source=PANTHEON_STRIPPED
29. Ekundayo, T., Khan, Z., & Nuzhat, S. (2024). Evaluating the Influence of Artificial Intelligence on Scholarly Research: A Study Focused on Academics. Human Behavior and Emerging Technologies, 2024(1), 8713718. https://doi.org/10.1155/2024/8713718
30. EL-Zanfaly, D., Huang, Y., & Dong, Y. (2022). Sand Playground: Designing Human-AI physical Interface for Co-creation in Motion. Creativity and Cognition, 49–55. https://doi.org/10.1145/3527927.3532791
31. FAQs for the policy on generative AI in research. (n.d.). University of Oxford. Retrieved August 22, 2025, from https://www.ox.ac.uk/research/support-researchers/research-practice/faqs-policy-generative-ai-research?utm_source=chatgpt.com
32. Farber, S. (2025). Comparing human and AI expertise in the academic peer review process: Towards a hybrid approach. Higher Education Research & Development, 44(4), 871–885. https://doi.org/10.1080/07294360.2024.2445575
33. Fitzpatrick, D. (2024, December 20). The Truth About AI-Generated Research And Its Impact On Education. Forbes. https://www.forbes.com/sites/danfitzpatrick/2024/12/20/the-truth-about-ai-generated-research-and-its-impact-on-education/
34. Fragiadakis, G., Diou, C., Kousiouris, G., & Nikolaidou, M. (2025). Evaluating Human-AI Collaboration: A Review and Methodological Framework (No. arXiv:2407.19098). arXiv. https://doi.org/10.48550/arXiv.2407.19098
35. França, C. (2023). AI empowering research: 10 ways how science can benefit from AI (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2307.10265
36. Georgieva, M., Webb, J., Stuart, J., Bell, J., Crawford, S., & Ritter-Guth, B. (2025). AI Ethical Guidelines. EDUCAUSE. https://library.educause.edu/resources/2025/6/ai-ethical-guidelines
37. Granjeiro, J. M., Cury, A. A. D. B., Cury, J. A., Bueno, M., Sousa-Neto, M. D., & Estrela, C. (2025). The Future of Scientific Writing: AI Tools, Benefits, and Ethical Implications. Brazilian Dental Journal, 36, e25-6471. https://doi.org/10.1590/0103-644020256471
38. Heikkilä, M. (2025a, February 19). Google builds AI ‘co-scientist’ tool to speed up research. Financial Times. https://www.ft.com/content/6e53cc55-9031-4ba4-9e7c-e5e9c02b3203
39. Heikkilä, M. (2025b, August 14). The art of persuasion: How top AI chatbots can change your mind. Financial Times. https://www.ft.com/content/31e528b3-9800-4743-af0a-f5c3b80032d0
40. Hemmer, P., Schemmer, M., Kühl, N., Vössing, M., & Satzger, G. (2024). Complementarity in Human-AI Collaboration: Concept, Sources, and Evidence (No. arXiv:2404.00029). arXiv. https://doi.org/10.48550/arXiv.2404.00029
41. Hennig, N. (n.d.). Integration of AI tools into your research: SciSpace. The University of Arizona. University Libraries. Retrieved August 19, 2025, from https://libguides.library.arizona.edu/ai-researchers/scispace
42. Holstein, K., De-Arteaga, M., Tumati, L., & Cheng, Y. (2023). Toward Supporting Perceptual Complementarity in Human-AI Collaboration via Reflection on Unobservables. Proceedings of the ACM on Human-Computer Interaction, 7(CSCW1), 1–20. https://doi.org/10.1145/3579628

43. Huff, C. H. (2024, October 1). The promise and perils of using AI for research and writing. American Psychological Association. https://www.apa.org/topics/artificial-intelligence-machine-learning/ai-research-writing
44. Ibrahim, L., Huang, S., Bhatt, U., Ahmad, L., & Anderljung, M. (2025, June 23). Towards Interactive Evaluations for Interaction Harms in Human-AI Systems. Knight First Amendment Institute. http://knightcolumbia.org/content/towards-interactive-evaluations-for-interaction-harms-in-human-ai-systems?utm_source=chatgpt.com
45. Jain, S. J., Sibbu, K., & Kuri, R. (2023). Conducting Effective Research using SciSpace: A Practical Approach. https://doi.org/10.22541/au.170111059.99508682/v1
46. Jiang, T., Sun, Z., Fu, S., & Lv, Y. (2024). Human-AI interaction research agenda: A user-centered perspective. Data and Information Management, 8(4), 100078. https://doi.org/10.1016/j.dim.2024.100078
47. Kelly, R. (2025, August 13). Workers view agents as 'important teammates' – but the prospect of an AI "boss" is a step too far. IT Pro. https://www.itpro.com/technology/artificial-intelligence/workers-view-agents-as-important-teammates-but-the-prospect-of-an-ai-boss-is-a-step-too-far
48. Khalifa, M., & Albadawy, M. (2024). Using artificial intelligence in academic writing and research: An essential productivity tool. Computer Methods and Programs in Biomedicine Update, 5, 100145. https://doi.org/10.1016/j.cmpbup.2024.100145
49. Khlaif, Z. N., Mousa, A., Hattab, M. K., Itmazi, J., Hassan, A. A., Sanmugam, M., & Ayyoub, A. (2023). The Potential and Concerns of Using AI in Scientific Research: ChatGPT Performance Evaluation. JMIR Medical Education, 9, e47049. https://doi.org/10.2196/47049
50. King's College London Libraries & Collections. (n.d.). Searching for Systematic Reviews & Evidence Synthesis: AI tools in evidence synthesis. Retrieved August 21, 2025, from https://libguides.kcl.ac.uk/systematicreview/ai
51. Kong, X., Fang, H., Chen, W., Xiao, J., & Zhang, M. (2025). Examining human–AI collaboration in hybrid intelligence learning environments: Insight from the Synergy Degree Model. Humanities and Social Sciences Communications, 12(1), 821. https://doi.org/10.1057/s41599-025-05097-z
52. Learn about NotebookLM. (n.d.). Google. Retrieved August 22, 2025, from https://support.google.com/notebooklm/answer/16164461?co=GENIE.Platform%3DDesktop&hl=en&utm_source=chatgpt.com
53. Living guidelines on the responsible use of generative AI in research. (2025, April). European Commission. Directorate-General for Research and Innovation. https://research-and-innovation.ec.europa.eu/document/download/2b6cf7e5-36ac-41cb-aab5-0d32050143dc_en?filename=ec_rtd_ai-guidelines.pdf
54. Lou, B., Lu, T., Raghu, T. S., & Zhang, Y. (2025). Unraveling Human-AI Teaming: A Review and Outlook (No. arXiv:2504.05755). arXiv. https://doi.org/10.48550/arXiv.2504.05755
55. Lund, B., Lamba, M., & Oh, S. H. (2024). The Impact of AI on Academic Research and Publishing (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2406.06009
56. Lund, B., Wang, T., Mannuru, N. R., Nie, B., Shimray, S., & Wang, Z. (2023). ChatGPT and a New Academic Reality: Artificial Intelligence-Written Research Papers and the Ethics of the Large Language Models in Scholarly Publishing. https://doi.org/10.48550/ARXIV.2303.13367
57. Madanchian, M., & Taherdoost, H. (2025). The impact of artificial intelligence on research efficiency. Results in Engineering, 26, 104743. https://doi.org/10.1016/j.rineng.2025.104743
58. Maker, P. (2025, August 14). The GenAI "crutch": Why teams must learn before they lean. TechRadar. https://www.techradar.com/pro/the-genai-crutch-why-teams-must-learn-before-they-lean
59. Malik, F. S., & Terzidis, O. (2025). A hybrid framework for creating artificial intelligence-augmented systematic literature reviews. Management Review Quarterly. https://doi.org/10.1007/s11301-025-00522-8
60. Mao, A., Mohri, M., & Zhong, Y. (2021). Principled Approaches for Learning to Defer with Multiple Experts (No. arXiv:ISAIM2024_Mao_etal). arXiv. https://isaim2024.cs.ou.edu/papers/ISAIM2024_Mao_etal.pdf
61. Maslach, D. (2023, December 14). Generative AI Can Supercharge Your Academic Research. Harvard Business Impact. https://hbsp.harvard.edu/inspiring-minds/generative-ai-can-supercharge-your-academic-research?utm_source=chatgpt.com

62. Maslach, D., Robot, C., Doshi, A. R., Mors, M. L., Puranam, P., & Seamans, R. C. (2023). Academic Research is Dead. The Robots Won. Academy of Management Proceedings, 2023(1), 13106. https://doi.org/10.5465/AMPROC.2023.13106symposium
63. Messeri, L., & Crockett, M. J. (2024). Artificial intelligence and illusions of understanding in scientific research. Nature, 627(8002), 49–58. https://doi.org/10.1038/s41586-024-07146-0
64. National Academy of Sciences, National Academy of Engineering, & National Academy of Medicine. (2022). Context for Automated Research Workflows. In Automated Research Workflows for Accelerated Discovery: Closing the Knowledge Discovery Loop. NAP.edu. https://doi.org/10.17226/26532
65. New ethical framework to help navigate use of AI in academic research. (2024, November 13). University of Oxford. https://www.ox.ac.uk/news/2024-11-13-new-ethical-framework-help-navigate-use-ai-academic-research
66. Nguyen, T., & Elbanna, A. (2025). Understanding Human-AI Augmentation in the Workplace: A Review and a Future Research Agenda. Information Systems Frontiers. https://doi.org/10.1007/s10796-025-10591-5
67. Nikolic, S., Sandison, C., Haque, R., Daniel, S., Grundy, S., Belkina, M., Lyden, S., Hassan, G. M., & Neal, P. (2024). ChatGPT, Copilot, Gemini, SciSpace and Wolfram versus higher education assessments: An updated multi-institutional study of the academic integrity impacts of Generative Artificial Intelligence (GenAI) on assessment, teaching and learning in engineering. Australasian Journal of Engineering Education, 29(2), 126–153. https://doi.org/10.1080/22054952.2024.2372154
68. Noor, D. (2025, June 27). How a new AI tool could amplify doubt in pollution science. The Guardian. https://www.theguardian.com/technology/ng-interactive/2025/jun/27/tony-cox-epidemiology-risk-assessment-chatgpt-ai
69. Oubou, I. (2024, December 19). Building an AI-Human Hybrid Content Team and Workflows. Vocable Blog. https://blog.vocable.ai/building-an-ai-human-hybrid-content-team/
70. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2203.02155
71. Paat, H., & Shen, G. (2025). Conformal Set-based Human-AI Complementarity with Multiple Experts (No. arXiv:2508.06997). arXiv. https://doi.org/10.48550/arXiv.2508.06997
72. Parker, J. (2024, February 20). Best AI Tools to Write a Literature Review. Moxie. https://moxielearn.ai/blog/best-ai-tools-to-write-a-literature-review
73. Pyae, A. (2025). The Human-AI Handshake Framework: A Bidirectional Approach to Human-AI Collaboration (No. arXiv:2502.01493). arXiv. https://doi.org/10.48550/arXiv.2502.01493
74. Raman, N., & Yee, M. (2021). Improving Learning-to-Defer Algorithms Through Fine-Tuning (No. arXiv:2112.10768). arXiv. https://arxiv.org/pdf/2112.10768
75. Reimann, P. (2025, April 10). Streamlining AI Literature Reviews in Web of Science. https://clarivate.com/academia-government/blog/streamlining-literature-review-with-agentic-ai-in-the-web-of-science/
76. Saluna, K. (2025, April 1). Beyond AI Detection: A Research Expert's Guide to Ethical AI Integration in Academic Workflows. Moxie. https://moxielearn.ai/blog/beyond-ai-detection-a-research-experts-guide-to-ethical-ai-integration-in-academic-workflows
77. Shabanov, I. (2024, February 11). SciSpace: An all-in-one AI tool for literature reviews. The Effortless Academic. https://effortlessacademic.com/scispace-an-all-in-one-ai-tool-for-literature-reviews/
78. Shao, Y., Jiang, Y., Kanell, T. A., Xu, P., Khattab, O., & Lam, M. S. (2024). Assisting in Writing Wikipedia-like Articles From Scratch with Large Language Models (No. arXiv:2402.14207). arXiv. https://doi.org/10.48550/arXiv.2402.14207
79. Stanford-oval/storm. (2025). [Python]. Stanford Open Virtual Assistant Lab. https://github.com/stanford-oval/storm (Original work published 2024)
80. Stokel-Walker, C. (n.d.). The AI Nobel Prizes Could Change the Focus of Research. Wired. Retrieved August 21, 2025, from https://www.wired.com/story/how-the-ai-nobel-prizes-could-change-the-focus-of-research/

81. STORM. (n.d.). STORM. Retrieved August 22, 2025, from https://storm.genie.stanford.edu/
82. The Do's & Don'ts of Using Generative AI Tools Ethically in Academia. (2023, December 21). Paperpal. https://paperpal.com/blog/academic-writing-guides/the-dos-donts-of-using-generative-ai-tools-ethically-in-academia
83. Tong, D., Jin, B., Tao, Y., Ren, H., Atiquil Islam, A. Y. M., & Bao, L. (2025). Exploring the role of human-AI collaboration in solving scientific problems. Physical Review Physics Education Research, 21(1), 010149. https://doi.org/10.1103/PhysRevPhysEducRes.21.010149
84. University of Virginia Office of Research Integrity. (n.d.). Reminder of the Importance of Research Integrity & Use of AI. University of Virginia. Retrieved August 21, 2025, from https://compliance.research.virginia.edu/about/integrity-and-ethics/reminder-importance-research-integrity-use-ai
85. Upson, M. (n.d.). Literature Reviews: AI Tools for Academic Research & Writing. Retrieved August 18, 2025, from https://info.library.okstate.edu/literaturereviews/ai-tools
86. Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour, 8(12), 2293–2303. https://doi.org/10.1038/s41562-024-02024-1
87. Vieriu, A. M., & Petrea, G. (2025). The Impact of Artificial Intelligence (AI) on Students' Academic Development. Education Sciences, 15(3), 343. https://doi.org/10.3390/educsci15030343
88. Vössing, M., Kühl, N., Lind, M., & Satzger, G. (2022). Designing Transparency for Effective Human-AI Collaboration. Information Systems Frontiers, 24(3), 877–895. https://doi.org/10.1007/s10796-022-10284-3
89. Wang, S., & Bin Shafqat, U. (2025, July 29). What's new in NotebookLM: Video Overviews and an upgraded Studio. Google. https://blog.google/technology/google-labs/notebooklm-video-overviews-studio-upgrades/
90. Wilkins, D. (2023). Automated title and abstract screening for scoping reviews using the GPT-4 Large Language Model (No. arXiv:2311.07918). arXiv. https://doi.org/10.48550/arXiv.2311.07918
91. Willett, R. (n.d.). How AI is transforming scientific research: [Broadcast]. Retrieved August 21, 2025, from https://news.uchicago.edu/how-ai-transforming-scientific-research
92. Williams, B. (2025, February 5). Best 7 AI Tools To Identify Research Gaps. Insight7 - AI Tool For Call Analytics & Evaluation. https://insight7.io/best-7-ai-tools-to-identify-research-gaps/
93. Woldetsadik, M., Community, K. A. A. | L. | L. A. C. C.-E. |, & Networks, Y. E. L. | L. S. A. H. | C. of G. L. (2024, February 19). The Best 8 AI-Powered Tools for Literature Review. INFOLIBNET – INFOrmation and LIBraries NETwork. https://mylibrarianship.wordpress.com/2024/02/19/the-best-8-ai-powered-tools-for-literature-review-2/
94. Xiong, W., Fan, H., Ma, L., & Wang, C. (2022). Challenges of human—Machine collaboration in risky decision-making. Frontiers of Engineering Management, 9(1), 89–103. https://doi.org/10.1007/s42524-021-0182-0
95. Xu, Z. (2025). Patterns and Purposes: A Cross-Journal Analysis of AI Tool Usage in Academic Writing (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2502.00632
96. Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. Smart Learning Environments, 11(1), 28. https://doi.org/10.1186/s40561-024-00316-7
97. Zotero. About. (n.d.). Zotero. Retrieved August 22, 2025, from https://www.zotero.org/about/
98. Zotero Privacy Policy. (n.d.). Zotero. Retrieved August 22, 2025, from https://www.zotero.org/support/privacy?utm_source=chatgpt.com
99. Zotero (Version 8.0). (n.d.). [Computer software]. Retrieved August 22, 2025, from https://github.com/zotero/zotero
100. Zotero/zotero-connectors. (2025). [JavaScript]. Zotero. https://github.com/zotero/zotero-connectors (Original work published 2011)

doi:10.20944/preprints202509.1249.v1

disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.