

# From Hypothesis to Publication: A Comprehensive Survey of AI-Driven Research Support Systems

Anonymous ACL submission

## Abstract

Research is a fundamental process driving the advancement of human civilization, yet it demands substantial time and effort from researchers. In recent years, the rapid development of artificial intelligence (AI) technologies has inspired researchers to explore how AI can accelerate and enhance research. To monitor relevant advancements, this paper presents a systematic review of the progress in this domain. Specifically, we organize the relevant studies into three main categories: hypothesis formulation, hypothesis validation, and manuscript publication. Hypothesis formulation involves knowledge synthesis and hypothesis generation. Hypothesis validation includes the verification of scientific claims, theorem proving, and experiment validation. Manuscript publication encompasses manuscript writing and the peer review process. Furthermore, we identify and discuss the current challenges faced in these areas, as well as potential future directions for research. Finally, we also offer a comprehensive overview of existing benchmarks and tools across various domains that support the integration of AI into the research process. We hope this paper serves as an introduction for beginners and fosters future research.

## 1 Introduction

Research is creative and systematic work aimed at expanding knowledge and driving civilization's development (Eurostat, 2018). Researchers typically identify a topic, review relevant literature, synthesize existing knowledge, and formulate hypothesis, which are validated through theoretical and experimental methods. Findings are then documented in manuscripts that undergo peer review before publication (Benos et al., 2007; Boyko et al., 2023). However, this process is resource-intensive, requiring specialized expertise and posing entry barriers for researchers (Blaxter et al., 2010).

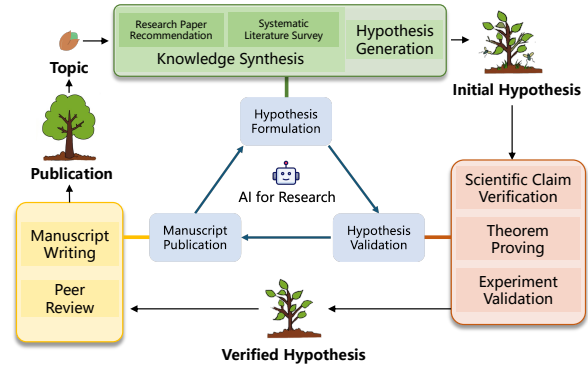


Figure 1: Overview of AI for research. The framework consists of three stages: hypothesis formulation, hypothesis validation, and manuscript publication. In the hypothesis formulation stage, knowledge integration leads to the proposal of an initial hypothesis after a topic is identified. The hypothesis validation stage involves verifying the hypothesis from three perspectives to ensure its correctness and validity. Finally, the manuscript publication stage focuses on drafting and publishing the validated hypothesis.

In recent years, artificial intelligence (AI) technologies, represented by large language models (LLMs), have experienced rapid development (Brown et al., 2020; OpenAI, 2023; Dubey et al., 2024; Yang et al., 2024a; DeepSeek-AI et al., 2024; Guo et al., 2025). These models exhibit exceptional capabilities in text understanding, reasoning, and generation (Schaeffer et al., 2023). In this context, AI is increasingly involving the entire research pipeline (Messerli and Crockett, 2024), sparking extensive discussion about its implications for research (Hutson, 2022; Williams et al., 2023; Morris, 2023; Fecher et al., 2023). Moreover, following the release of ChatGPT, approximately 20% of academic papers and peer-reviewed texts in certain fields have been modified by LLMs (Liang et al., 2024a,b). A study also reveals that 81% of researchers integrate LLMs into their workflows (Liao et al., 2024).

As the application of AI in research attracts increasing attention, a significant body of related studies has begun to emerge. To systematically synthesize existing research, we present comprehensive survey that emulates human researchers by using the research process as an organizing framework. Specifically, as depicted in Figure 1, the research process is divided into three key stages: (1) Hypothesis Formulation, involving knowledge synthesis and hypothesis generation; (2) Hypothesis Validation, encompassing scientific claim verification, theorem proving, and experiment validation; (3) Manuscript Publication, which focuses on academic publications and is further divided into manuscript writing and peer review.

**Comparing with Existing Surveys** Although Luo et al. (2025) reviews the application of AI in research, it predominantly focuses on LLMs, while neglecting the knowledge synthesis that precedes hypothesis generation and the theoretical validation of hypothesis. Other surveys concentrate on more specific areas, such as paper recommendation (Beel et al., 2016; Bai et al., 2019; Kreutz and Schenkel, 2022), scientific literature review (Altmami and Menai, 2022), hypothesis generation (Kulkarni et al., 2025), scientific claim verification (Vladika and Matthes, 2023; Dmonte et al., 2024), theorem proving (Li et al., 2024e), manuscript writing (Li and Ouyang, 2024), and peer review (Lin et al., 2023a; Kousha and Thelwall, 2024). Additionally, certain surveys emphasize the application of AI in scientific domains (Zheng et al., 2023b; Zhang et al., 2024d; Gridach et al., 2025).

**Contributions** Our contributions can be summarized as follows: (1) We align the relevant fields with the research process of human researchers, systematically integrating and extending these aspects while primarily focusing on the research process itself. (2) We introduce a meticulous taxonomy (shown in Figure 2). (3) We provide a summary of tools that can assist in the research process. (4) We discuss new frontiers, outline their challenges, and shed light on future research.

**Survey Organization** We first elaborate hypothesis formulation (§2), followed by hypothesis validation (§3) and manuscript publication (§4). Additionally, we present benchmarks (§5), and tools (§6) that utilized in research. Finally, we outline challenges as well as future directions (§7). In the Appendix, we provide further discussion on open questions (§A), challenges faced in different do-

main (§B), discussion about relevant ethical considerations (§C), and a comparison of capabilities among different methods (§D).

## 2 Hypothesis Formulation

This stage centers on the process of hypothesis formulation. As illustrated in Figure 3, it commences with developing a comprehensive understanding of the domain, followed by identifying a specific aspect and generating pertinent hypothesis. This section is further structured into two key components: Knowledge Synthesis and Hypothesis Generation.

### 2.1 Knowledge Synthesis

Knowledge synthesis constitutes the foundational step in the research process. During this phase, researchers are required to identify and critically evaluate existing literature to establish a thorough understanding of the field. This step is pivotal for uncovering new research directions, refining methodologies, and supporting evidence-based decision-making (Asai et al., 2024). In this section, the process of knowledge synthesis is divided into two modules: Research Paper Recommendation and Systematic Literature Review.

#### 2.1.1 Research Paper Recommendation

Research paper recommendation (RPR) identifies and recommends novel and seminal articles aligned with researchers’ interests. Prior studies have shown that recommendation systems outperform keyword-based search engines in terms of efficiency and reliability when extracting valuable insights from large-scale datasets (Bai et al., 2019). Existing methodologies are broadly categorized into four paradigms: content-based filtering, collaborative filtering, graph-based approaches, and hybrid systems (Beel et al., 2016; Li and Zou, 2019; Bai et al., 2019; Shahid et al., 2020). Recent advancements propose multi-dimensional classification frameworks based on data source utilization (Kreutz and Schenkel, 2022).

Recent trends in research suggest a decline in publication volumes related to RPR (Sharma et al., 2023), alongside an increasing focus on user-centric optimizations. Existing studies emphasize the limitations of traditional paper-centric interaction models and advocate for more effective utilization of author relationship graphs (Kang et al., 2023). Multi-stage recommendation architectures that integrate diverse methodologies have been shown to achieve superior performance (Pinedo

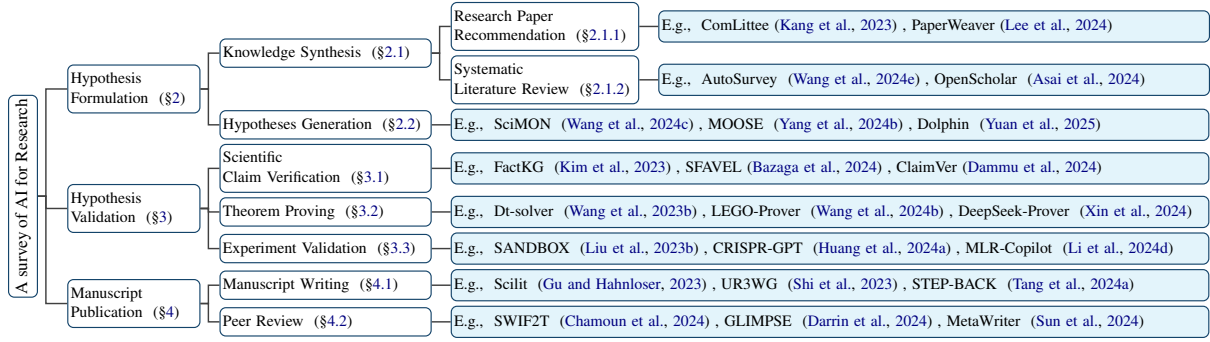


Figure 2: Taxonomy of Hypothesis Formulation, Hypothesis Validation and Manuscript Publication (This is a simplified version, full version in Figure 6).

et al., 2024; Stergiopoulos et al., 2024). Visualization techniques that link recommended papers to users’ publication histories via knowledge graphs (Kang et al., 2022) and LLMs-powered comparative analysis frameworks (Lee et al., 2024) represent emerging directions for enhancing interpretability and contextual relevance.

### 2.1.2 Systematic Literature Review

Systematic literature review (SLR) constitutes a rigorous and structured methodology for evaluating and integrating prior research on a specific topic (Webster and Watson, 2002; Zhu et al., 2023; Bolaños et al., 2024). In contrast to single-document summaries (Elhadad et al., 2005), SLR entails synthesizing information across multiple related scientific documents (Altmami and Menai, 2022). SLR can further be divided into two stages: outline generation and full-text generation (Shao et al., 2024; Agarwal et al., 2024b; Block and Kuckertz, 2024).

**Outline generation**, especially structured outline generation, is highlighted by recent studies as a pivotal factor in enhancing the quality of surveys. Zhu et al. (2023) demonstrated that hierarchical frameworks substantially enhance survey coherence. AutoSurvey (Wang et al., 2024e) extended conventional outline generation by recommending both sub-chapter titles and detailed content descriptions, ensuring comprehensive topic coverage. Additionally, multi-level topic generation via clustering methodologies has been proposed as an effective strategy for organizing survey structures (Katz et al., 2024). Advanced systems such as STORM (Shao et al., 2024) employed LLM-driven outline drafting combined with multi-agent discussion cycles to iteratively refine the generated outlines. Tree-based hierarchical architectures have gained increasing adoption in this domain. For instance, CHIME (Hsu et al., 2024) optimized

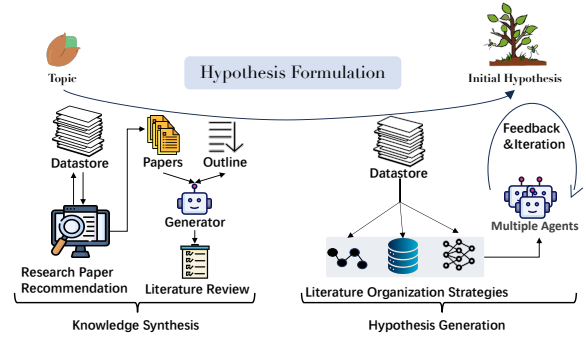


Figure 3: This figure illustrates the hypothesis formulation process, consisting of two stages: knowledge synthesis and hypothesis generation, which together produce an initial hypothesis related to a specific topic.

LLM-generated hierarchies through human-AI collaboration, while HiReview (Hu et al., 2024b) demonstrated the efficacy of multi-layer tree representations for systematic knowledge organization.

**Full-text generation** follows the outline generation stage. AutoSurvey and Lai et al. (2024) utilized LLMs with carefully designed prompts to construct comprehensive literature reviews step-by-step. PaperQA2 (Skarlinski et al., 2024) introduced an iterative agent-based approach for generating reviews, while STORM employed multi-agent conversation data for this purpose. LitLLM (Agarwal et al., 2024a) and Agarwal et al. (2024b) adopted a plan-based search enhancement strategy. KGSUM (Wang et al., 2022a) integrated knowledge graph information into paper encoding and used a two-stage decoder for summary generation. Bio-SIEVE (Robinson et al., 2023) and Susnjak et al. (2024) fine-tuned LLMs for automatic review generation. OpenScholar (Asai et al., 2024) developed a pipeline that trained a new model without relying on a dedicated survey-generation model.



## 2.2 Hypothesis Generation

Hypothesis generation, known as idea generation, refers to the process of coming up with new concepts, solutions, or approaches. It is the most important step in driving the progress of the entire research (Qi et al., 2023).

Early work focused more on predicting relationships between concepts, because researchers believed that new concepts come from links with old concepts (Henry and McInnes, 2017; Krenn et al., 2022). As language models became more powerful (Zhao et al., 2023a), researchers are beginning to focus on open-ended idea generation (Girotra et al., 2023; Si et al., 2024; Kumar et al., 2024). Recent advancements in AI-driven hypothesis generation highlight diverse approaches to research conceptualization. For instance, MOOSE-Chem (Yang et al., 2024c) and IdeaSynth (Pu et al., 2024) used LLMs to bridge inspiration-to-hypothesis transformation via interactive frameworks. The remaining research primarily falls into two areas: enhancing input data quality and improving the quality of generated hypothesis.

**Input data quality improvement** is demonstrated by Majumder et al. (2024a); Liu et al. (2024a), who showed that LLMs can generate comprehensive hypothesis from existing academic data. Literature organization strategies have evolved through various methodologies, including triplet representations (Wang et al., 2024c), chain-based architectures (Li et al., 2024a), and complex database systems (Wang et al., 2024d). Knowledge graphs emerge as critical infrastructure (Hogan et al., 2021), enabling semantic relationship mapping via subgraph identification (Buehler, 2024; Ghafarollahi and Buehler, 2024). Notably, SciMuse (Gu and Krenn, 2024) pioneered researcher-specific hypothesis generation by constructing personalized knowledge graphs.

**Hypothesis quality improvement** has been addressed through feedback and iteration (Rabby et al., 2025), as proposed by HypoGeniC (Zhou et al., 2024) and MOOSE (Yang et al., 2024b). Feedback mechanisms include direct responses to hypothesis (Baek et al., 2024), experimental outcome evaluations (Ma et al., 2024; Yuan et al., 2025), comparison with the existing literature (Schmidgall and Moor, 2025), and automated peer review comments (Lu et al., 2024). FunSearch (Romera-Paredes et al., 2024) generates solutions by iteratively combining the innovative

capabilities of LLM with the verification capabilities of an evaluator. Beyond iterative feedback, collaborative efforts among researchers have also been recognized, leading to multi-agent hypothesis generation approaches (Nigam et al., 2024; Ghafarollahi and Buehler, 2024). VIRSCI (Su et al., 2024) further optimized this process by customizing knowledge for each agent. Additionally, the Nova framework (Hu et al., 2024a) was introduced to refine hypothesis by leveraging outputs from other research as input.

Knowledge synthesis and hypothesis generation comprise the hypothesis formulation phase. Research paper recommendation supports knowledge acquisition, while systematic literature review aids organization within knowledge synthesis. Recent advances integrate LLMs (de la Torre-López et al., 2023) to enhance knowledge integration (Huang and Tan, 2023; Gupta et al., 2023; Kacena et al., 2024; Tang et al., 2024b). By developing a deep understanding of a domain through knowledge synthesis, researchers can identify research directions and use hypothesis generation techniques to formulate hypothesis. Additionally, the distinction between scientific discovery and hypothesis generation is discussed in §A.

## 3 Hypothesis Validation

In scientific research, any proposed hypothesis must undergo rigorous validation to establish its validity. In some studies, this process is also referred to as 'falsification' (Liu et al., 2024d; Huang et al., 2025). As illustrated in Figure 4, this section explores the application of AI in verifying scientific hypothesis through three methodological components: Scientific Claim Verification, Theorem Proving, and Experiment Validation.

### 3.1 Scientific Claim Verification

Scientific claim verification, also referred to as scientific fact-checking or scientific contradiction detection, aims to assess the veracity of claims related to scientific knowledge. This process assists scientists in verifying research hypothesis, discovering evidence, and advancing scientific work (Wadden et al., 2020; Vladika and Matthes, 2023; Skarlinski et al., 2024). Research on scientific claim verification primarily focuses on three key elements: the claim, the evidence, and the validity of the claim (Dmonte et al., 2024).

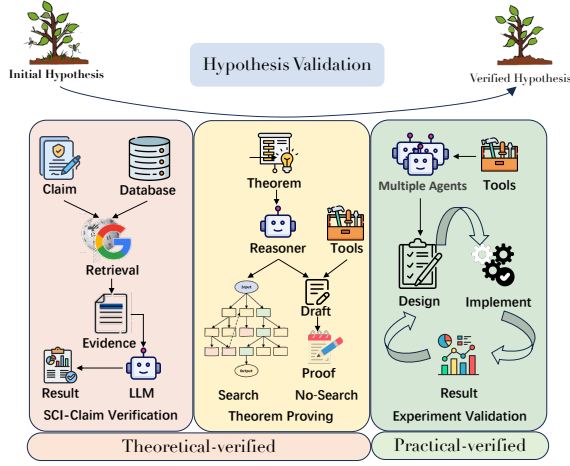


Figure 4: This figure illustrates the various perspectives for hypothesis validation during the hypothesis validation stage. A hypothesis is typically divided into scientific claims and theorems, with SCI-claim verification (scientific claim verification) and theorem proving ensuring theoretical correctness, while experiment validation assesses practical feasibility.

**Claim** Studies have highlighted that certain claims lack supporting evidence (Wühl et al., 2024a), while others have demonstrated the ability to perform claim-evidence alignment without annotated data (Bazaga et al., 2024). Additionally, methods such as HiSS (Zhang and Gao, 2023) and ProToCo (Zeng and Gao, 2023) proposed generating multiple claim variants to enhance verification.

**Evidence** Existing research has explored various aspects related to evidence, including evidentiary sources (Vladika and Matthes, 2024a), retrieval configurations (Vladika and Matthes, 2024b), strategies for identifying and mitigating flawed evidence (Glockner et al., 2022; Wühl et al., 2024b; Glockner et al., 2024a), and approaches for processing sentence-level (Pan et al., 2023b) versus document-level indicators (Wadden et al., 2022b).

**Verification** In the verification results generation phase, MAGIC (Kao and Yen, 2024) and SERIf (Cao et al., 2024b) proposed utilizing LLMs to synthesize evidence into more comprehensive information. FactKG (Kim et al., 2023) and Muharram and Purwarianti (2024) structured evidence into knowledge graphs, enabling claim attribution (Dammu et al., 2024; Wu et al., 2023). Furthermore, Atanasova et al. (2020); Krishna et al. (2022); Pan et al. (2023a); Eldifrawi et al. (2024); Zhang et al. (2024b) advocated for generating explanatory annotations alongside experimental out-

comes during the verification process. Meanwhile, Das et al. (2023); Altuncu et al. (2023) emphasized the critical role of domain expertise in ensuring accurate verification.

### 3.2 Theorem Proving

Theorem proving constitutes a subtask of logical reasoning, aimed at reinforcing the validity of the underlying theory within a hypothesis (Pease et al., 2019; Yang et al., 2023c; Li et al., 2024e).

Following the proposal of GPT-f (Polu and Sutskever, 2020) to utilize generative language models for theorem proving, researchers initially combined search algorithms with language models (Lample et al., 2022; Wang et al., 2023b). However, a limitation in search-based approaches was later identified by Wang et al. (2024a), who highlighted their tendency to explore insignificant intermediate conjectures. This led some teams to abandon search algorithms entirely. Subsequently, alternative methods emerged, such as the two-stage framework proposed by Jiang et al. (2023) and Lin et al. (2024), which prioritized informal conceptual generation before formal proof construction. Thor (Jiang et al., 2022a) introduced theorem libraries to accelerate proof generation, an approach enhanced by Logo-power (Wang et al., 2024b) through dynamic libraries. Studies like Baldur (First et al., 2023), Mustard (Huang et al., 2024c), and DeepSeek-Prover (Xin et al., 2024) demonstrated improvements via targeted data synthesis and fine-tuning, though COPRA (Thakur et al., 2024) questioned their generalizability and proposed an environment-agnostic alternative. Complementary strategies included theoretical decomposition into sub-goals (Zhao et al., 2023b) and leveraging LLMs as collaborative assistants in interactive environments (Song et al., 2024).

### 3.3 Experiment Validation

Experiment validation involves designing and conducting experiments based on the hypothesis. The empirical validity of the hypothesis is then determined through analysis of the experimental results (Huang et al., 2024b).

Experiment validation represents a time-consuming component of scientific research. Recent advancements in LLMs have enhanced their ability to plan and reason about experimental tasks (Kambhampati et al., 2024), prompting researchers to use these models for designing and implementing experiments (Ruan et al., 2024b).

To ensure accuracy, studies such as Zhang et al. (2023) and Arlt et al. (2024) imposed input/output constraints, though this reduced generalizability. To address this, Boiko et al. (2023); Bran et al. (2024); Huang et al. (2024a) integrated tools to expand model capabilities. Full automation was achieved by Ni and Buehler (2023); Li et al. (2024a); Lu et al. (2024) through prompt-guided multi-agent collaboration. Madaan et al. (2023); Yuan et al. (2025) further highlighted that the integration of feedback mechanisms demonstrated potential for enhancing design quality, while Zhang et al. (2024a); Liu et al. (2024c); Ni et al. (2024) employed experimental outcomes to refine hyperparameter configurations, and Szymanski et al. (2023); Li et al. (2024d); Baek et al. (2024) leveraged agent-generated analytical insights to facilitate iterative hypothesis refinement. In contrast, social science research often uses LLMs as experimental subjects to simulate human participants (Liu et al., 2023b; Manning et al., 2024; Mou et al., 2024).

A hypothesis can be conceptualized as consisting of two key components: claims and theorems. Scientific claim verification and theorem proving offer theoretical validation of hypothesis through formal reasoning and logical deduction, whereas experiment validation provides comprehensive practical validation via empirical testing.

## 4 Manuscript Publication

Upon validating a hypothesis as feasible, researchers generally progress to the publication stage. As depicted in Figure 5, this section categorizes Manuscript Publication into two primary components: Manuscript Writing and Peer Review.

### 4.1 Manuscript Writing

Manuscript writing, also referred to as scientific or research writing. At this stage, researchers articulate the hypothesis they have formulated and the results they have validated in the form of a scholarly paper. This process is crucial, as it not only disseminates findings but also deepens researchers' understanding of their work (Colyar, 2009).

Early shared tasks focused on assisting researchers in writing or analyzing linguistic features (Dale and Kilgarriff, 2010; Daudaravicius, 2015). Recent advances have led to three main directions: citation text generation, related work generation, and complete manuscript generation.

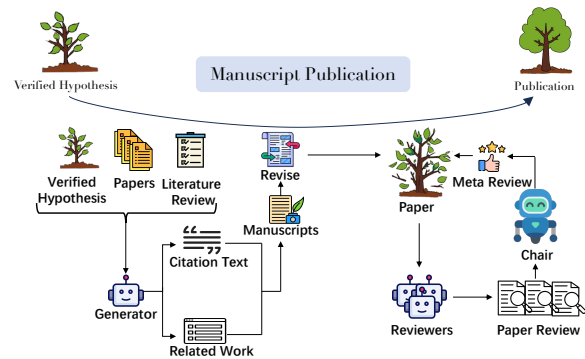


Figure 5: This figure shows the transformation of a validated hypothesis into a publication, leveraging outputs from the hypothesis formulation and validation stages.

**Citation Text Generation (Sentence Level)** A subset of research on AI in scientific writing has focused on citation text generation, which addresses the academic need for referencing prior work while mitigating model inaccuracies (Gao et al., 2023b; Gu and Hahnloser, 2023). For instance, Wang et al. (2022b) developed an automated citation generation system by integrating manuscript content with citation graphs. However, its reliance on rigid template-based architectures led to inflexible citation formats. This limitation motivated subsequent studies to propose incorporating citation intent as a control parameter during text generation, aiming to improve contextual relevance and rhetorical adaptability (Yu et al., 2022; Jung et al., 2022; Koo et al., 2023; Gu and Hahnloser, 2024).

**Related Work Generation (Paragraph Level)** In contrast to citation text generation, several studies have focused on related work generation in scholarly writing, emphasizing the production of multiple citation texts and the systematic analysis of inter-citation relationships (Li and Ouyang, 2022, 2024). The ScholaCite framework (Martin-Boyle et al., 2024) leveraged GPT-4 to cluster citation sources and generate draft literature review sections, although it required manually provided reference lists. By contrast, the UR3WG system (Shi et al., 2023) adopted a retrieval-augmented architecture integrated with large language models to autonomously acquire relevant references. To improve the quality of generated related work sections, Yu et al. (2024b) utilized GNNs to model complex relational dynamics between target manuscripts and cited literature, while Nishimura et al. (2024) initiative advocated for explicit novelty assertions regarding referenced publications.



## Complete Manuscripts Generation (Full-text Level)

The aforementioned investigations primarily focused on specific components of scientific writing, while a study by [Lai et al. \(2024\)](#) explored the progressive generation of complete manuscripts via structured workflows. The AI-Scientist system ([Lu et al., 2024](#)) further introduced section-wise self-reflection mechanisms to enhance compositional coherence. Several studies emphasized human-AI collaborative frameworks for improving writing efficiency ([Lin, 2024](#); [Feng et al., 2024](#); [Ifargan et al., 2024](#)), whereas [Tang et al. \(2024a\)](#) concentrated on enabling personalized content generation in multi-author collaborative environments. Following initial manuscript drafting, subsequent text revision processes were systematically examined ([Du et al., 2022b](#); [Jourdan et al., 2023](#); [Dang et al., 2025](#)). The OREO system ([Li et al., 2022](#)) utilized attribute classification for iterative in-situ editing, while [Du et al. \(2022a\)](#); [Pividori and Greene \(2024\)](#) incorporated researcher feedback loops for progressive text optimization. Notably, [Kim et al. \(2022\)](#); [Chamoun et al. \(2024\)](#); [D’Arcy et al. \(2024b\)](#) proposed replacing manual feedback with automated evaluation metrics.

## 4.2 Peer Review

Peer review serves as a critical mechanism for improving the quality of academic manuscripts through feedback and evaluation, forming the cornerstone of quality control in scientific research. However, the process is hindered by its slow pace, high time consumption, and increasing strain due to the growing academic workload ([Lin et al., 2023a](#); [Kousha and Thelwall, 2024](#); [Thelwall and Yaghi, 2024](#)). To address these challenges and enhance manuscript quality, researchers have investigated the application of AI in peer review ([Yuan et al., 2022](#); [Liu and Shah, 2023](#); [Niu et al., 2023](#); [Kuznetsov et al., 2024](#); [Thakkar et al., 2025](#)). Peer review can be categorized into two main types: paper review generation and meta-review generation.

**Paper Review Generation** In paper review generation, reviewers provide both scores and evaluations for manuscripts. For instance, [Setio and Tsuchiya \(2022\)](#) formulated score prediction as a regression task, [Muangkammuen et al. \(2022\)](#) utilized semi-supervised learning, and [Couto et al. \(2024\)](#) treated the task as a classification problem to evaluate the alignment between manuscripts and review criteria. While these approaches focused

on label prediction for paper reviews, [Yuan and Liu \(2022\)](#) extended the scope by directly generating reviews through the construction of a concept graph integrated with a citation graph.

Subsequently, a pilot study conducted by [Robertson \(2023\)](#) demonstrated the capability of GPT-4 to generate paper reviews. Further investigations, such as those by AI-Scientist ([Lu et al., 2024](#)) and [Liang et al. \(2023\)](#), evaluated its performance as a review agent. Additionally, systems like MARG ([D’Arcy et al., 2024a](#)) and SWIF2T ([Chamoun et al., 2024](#)) employed multi-agent frameworks to generate reviews via internal discussions and task decomposition. In contrast, AgentReview ([Jin et al., 2024](#)) and [Tan et al. \(2024\)](#) modeled the review process as a dynamic, multi-turn dialogue. Furthermore, CycleResearcher ([Weng et al., 2024](#)) and OpenReviewer ([Idahl and Ahmadi, 2024](#)) fine-tuned models for comparative reviews and structured outputs aligned with conference guidelines.

**Meta-Review Generation** In meta-review generation, chairs are tasked with identifying a paper’s core contributions, strengths, and weaknesses while synthesizing expert opinions on manuscript quality. Meta-reviews are conceptualized as abstractions of comments, discussions, and paper abstracts ([Li et al., 2023](#)). [Santu et al. \(2024\)](#) investigated the use of LLMs for automated meta-review generation, while [Zeng et al. \(2023\)](#) proposed a guided, iterative prompting approach. MetaWriter ([Sun et al., 2024](#)) utilized LLMs to extract key reviewer arguments, whereas GLIMPSE ([Darrin et al., 2024](#)) and [Kumar et al. \(2023\)](#) focused on reconciling conflicting statements to ensure fairness. Additionally, [Li et al. \(2024b\)](#) introduced a three-layer sentiment consolidation framework for meta-review generation, and PeerArg ([Sukpanichnant et al., 2024](#)) integrated LLMs with knowledge representation to address subjectivity and bias via a multiparty argumentation framework (MPAF). DeepReview ([Zhu et al., 2025](#)) generates a comprehensive meta-review by simulating expert evaluation across multiple dimensions.

During the Manuscript Publication phase, researchers can leverage AI to systematically complete manuscript writing by incorporating validated hypothesis, related papers, and literature reviews. The manuscript is subsequently subjected to peer review, involving iterative revisions before culminating in its final publication.

## 5 Benchmarks

Given that AI for research spans multiple disciplines, the tasks addressed within each domain vary significantly. To facilitate cross-domain exploration, we provide a summary of benchmarks associated with various areas, including research paper recommendation, systematic literature review, hypothesis generation, scientific claim verification, theorem proving, experiment verification, manuscript writing, and peer review. An overview of these benchmarks is presented in Table 9.

## 6 Tools

To accelerate the research workflow, we have curated a collection of tools designed to support various stages of the research process, with their applicability specified for each stage. To ensure practical relevance, our selection criteria emphasize tools that are publicly accessible or demonstrate significant influence on GitHub. A comprehensive overview of these tools is presented in Table 10.

## 7 Challenges

We identify several intriguing and promising avenues for future research.

### 7.1 Integration of Diverse Research Tasks

The research process is an integrated pipeline of interdependent stages. Paper recommendation and literature review provide an AI tool with a field overview and relevant works, ensuring that hypothesis generation is informed and of higher quality. Hypothesis validation assesses feasibility both logically and practically, with results feeding back to refine the hypothesis. In manuscript writing, validated hypotheses and prior outputs serve as key inputs. Peer review evaluates the manuscript and offers feedback across modules, enabling the hypothesis generator to adjust content accordingly (Lu et al., 2024). In addition, combinations can also be made between some small fields, for instance, meta-review generation could be integrated with scientific claim verification, experiment verification could be linked with hypothesis formulation (Jansen et al., 2025; Yuan et al., 2025; Liu et al., 2024d), and research paper recommendation systems could be connected with manuscript writing processes (Gu and Hahnloser, 2023). Furthermore, some studies have begun to emphasize the development of systems capable of covering mul-

multiple stages of the research process (Jansen et al., 2024; Weng et al., 2024; Yu et al., 2024a).

### 7.2 Integration with Reasoning-Oriented Language Models

Research is a process that places a significant emphasis on logic and reasoning. Theorem proving serves as a subtask within logical reasoning (Li et al., 2024e), while hypothesis generation is widely recognized as the primary form of reasoning employed by scientists when observing the world and proposing hypothesis to explain these observations (Yang et al., 2024b). Experiment verification, in turn, demands a high degree of planning capability from models (Kambhampati et al., 2024). Recent advances in reasoning-oriented language models, such as OpenAI-o1 (Jaech et al., 2024) and DeepSeek-R1 (Guo et al., 2025), have substantially enhanced the reasoning abilities of these models. Consequently, we posit that integrating reasoning language models with reasoning tasks is a promising future direction. This prediction was validated by experiments conducted by Schmidgall et al. (2025) using o1-Preview.

Furthermore, in Appendix §B, we provide a summary of the challenges in hypothesis formulation, validation, and manuscript publication.

## 8 Conclusion

This paper provides a systematic survey of existing research on AI for research, offering a comprehensive review of the advancements in the field. Within each category, we offer detailed descriptions of the associated subfields. In addition, we examine current challenges, ethical considerations, and potential directions for future research. To support researchers in exploring AI-driven research applications and enhancing workflow efficiency, we also summarize existing benchmarks and tools, accompanied by a comparative analysis of representative methods and their capabilities.

Furthermore, in the course of investigating various subfields within AI for research, we observed that this domain remains in its infancy. Research in numerous directions remains at an experimental stage, and substantial progress is necessary before these approaches can be effectively applied in practical scenarios. We hope that this survey serves as an introduction to the field for researchers and contributes to its continued advancement.



## Limitation

This study presents a comprehensive survey of AI for research, based on the framework of the research process conducted by human researchers.

We have made our best effort, but there may still be some limitations. Due to space constraints, we provide only concise summaries of each method without detailed technical elaboration. Given the rapid progress in AI and the expanding research landscape, we primarily focus on works published after 2022, with earlier studies receiving less attention. To emphasize areas that closely mimic the human research process, some topics are excluded from the main text but briefly discussed in Appendix §A. Moreover, as AI for Research is still an emerging field, the lack of standardized benchmarks and evaluation metrics hinders direct comparison. Nonetheless, we offer a comparative analysis of representative methods across domains using attribute graphs in Appendix §D.

## References

Shubham Agarwal, Issam H. Laradji, Laurent Charlin, and Christopher Pal. 2024a. [Litlm: A toolkit for scientific literature review](#). *CoRR*, abs/2402.01788.

Shubham Agarwal, Gaurav Sahu, Abhay Puri, Issam H Laradji, Krishnamurthy DJ Dvijotham, Jason Stanley, Laurent Charlin, and Christopher Pal. 2024b. Lims for literature review: Are we there yet? *arXiv preprint arXiv:2412.15249*.

Microsoft Research AI4Science and Microsoft Azure Quantum. 2023. [The impact of large language models on scientific discovery: a preliminary study using GPT-4](#). *CoRR*, abs/2311.07361.

Fadi Aljamaan, Mohamad-Hani Tamsah, Ibraheem Altamimi, Ayman Al-Eyadhy, Amr Jamal, Khalid Alhasan, Tamer A Mesallam, Mohamed Farahat, Khalid H Malki, and 1 others. 2024. Reference hallucination score for medical artificial intelligence chatbots: development and usability study. *JMIR Medical Informatics*, 12(1):e54345.

Nouf Ibrahim Altmami and Mohamed El Bachir Menai. 2022. [Automatic summarization of scientific articles: A survey](#). *J. King Saud Univ. Comput. Inf. Sci.*, 34(4):1011–1028.

Enes Altuncu, Jason R. C. Nurse, Meryem Bagriacik, Sophie Kaleba, Haiyue Yuan, Lisa Bonheme, and Shujun Li. 2023. [aedfact: Scientific fact-checking made easier via semi-automatic discovery of relevant expert opinions](#). *CoRR*, abs/2305.07796.

Sören Arlt, Haonan Duan, Felix Li, Sang Michael Xie, Yuhuai Wu, and Mario Krenn. 2024. [Meta-designing](#)

[quantum experiments with language models](#). *CoRR*, abs/2406.02470.

Akari Asai, Jacqueline He, Rulin Shao, Weijia Shi, Amanpreet Singh, Joseph Chee Chang, Kyle Lo, Luca Soldaini, Sergey Feldman, Mike D’Arcy, David Wadden, Matt Latzke, Minyang Tian, Pan Ji, Shengyan Liu, Hao Tong, Bohao Wu, Yanyu Xiong, Luke Zettlemoyer, and 6 others. 2024. [Openscholar: Synthesizing scientific literature with retrieval-augmented lms](#). *CoRR*, abs/2411.14199.

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [Generating fact checking explanations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7352–7364. Association for Computational Linguistics.

Sai Anirudh Athaluri, Sandeep Varma Manthena, VSR Krishna Manoj Kesapragada, Vineel Yarlagadda, Tirth Dave, and Rama Tulasi Siri Duddumpudi. 2023. Exploring the boundaries of reality: investigating the phenomenon of artificial intelligence hallucination in scientific writing through chatgpt references. *Cureus*, 15(4).

Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2024. [Researchagent: Iterative research idea generation over scientific literature with large language models](#). *CoRR*, abs/2404.07738.

Xiaomei Bai, Mengyang Wang, Ivan Lee, Zhuo Yang, Xiangjie Kong, and Feng Xia. 2019. [Scientific paper recommendation: A survey](#). *IEEE Access*, 7:9324–9339.

André Bauer, Simon Trapp, Michael Stenger, Robert Leppich, Samuel Kounev, Mark Leznik, Kyle Chard, and Ian Foster. 2024. Comprehensive exploration of synthetic data generation: A survey. *arXiv preprint arXiv:2401.02524*.

Adrián Bazaga, Pietro Lio, and Gos Micklem. 2024. [Unsupervised pretraining for fact verification by language model distillation](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Joeran Beel, Bela Gipp, Stefan Langer, and Corinna Breitinger. 2016. Paper recommender systems: a literature survey. *International Journal on Digital Libraries*, 17:305–338.

Dale J Benos, Edlira Bashari, Jose M Chaves, Amit Gagar, Niren Kapoor, Martin LaFrance, Robert Mans, David Mayhew, Sara McGowan, Abigail Polter, and 1 others. 2007. The ups and downs of peer review. *Advances in physiology education*, 31(2):145–152.

Loraine Blaxter, Christina Hughes, and Malcolm Tight. 2010. *How to research*. McGraw-Hill Education (UK).

787	Joern Block and Andreas Kuckertz. 2024. What is the future of human-generated systematic literature reviews in an age of artificial intelligence? <i>Management Review Quarterly</i> , pages 1–6.	
788		
789		
790		
791	Ben Bogin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2024. <a href="#">SUPER: evaluating agents on setting up and executing tasks from research repositories</a> . In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024</i> , pages 12622–12645. Association for Computational Linguistics.	
792		
793		
794		
795		
796		
797		
798		
799		
800	Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. 2023. <a href="#">Autonomous chemical research with large language models</a> . <i>Nat.</i> , 624(7992):570–578.	
801		
802		
803		
804	Francisco Bolaños, Angelo A. Salatino, Francesco Osborne, and Enrico Motta. 2024. <a href="#">Artificial intelligence for literature reviews: opportunities and challenges</a> . <i>Artif. Intell. Rev.</i> , 57(9):259.	
805		
806		
807		
808	James Boyko, Joseph Cohen, Nathan Fox, Maria Han Veiga, Jennifer I-Hsiu Li, Jing Liu, Bernardo Modenesi, Andreas H. Rauch, Kenneth N. Reid, Soumi Tribedi, Anastasia Visseratina, and Xin Xie. 2023. <a href="#">An interdisciplinary outlook on large language models for scientific research</a> . <i>CoRR</i> , abs/2311.04929.	
809		
810		
811		
812		
813		
814	Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D. White, and Philippe Schwaller. 2024. <a href="#">Augmenting large language models with chemistry tools</a> . <i>Nat. Mac. Intell.</i> , 6(5):525–535.	
815		
816		
817		
818	Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. <i>Advances in neural information processing systems</i> , 33:1877–1901.	
819		
820		
821		
822		
823		
824	Markus J. Buehler. 2024. <a href="#">Accelerating scientific discovery with generative knowledge extraction, graph-based representation, and multimodal intelligent graph reasoning</a> . <i>Mach. Learn. Sci. Technol.</i> , 5(3):35083.	
825		
826		
827		
828		
829	Ruisheng Cao, Fangyu Lei, Haoyuan Wu, Jixuan Chen, Yeqiao Fu, Hongcheng Gao, Xinzhuang Xiong, Hanchong Zhang, Wenjing Hu, Yuchen Mao, Tianbao Xie, Hongshen Xu, Danyang Zhang, Sida I. Wang, Ruoxi Sun, Pengcheng Yin, Caiming Xiong, Ansong Ni, Qian Liu, and 4 others. 2024a. <a href="#">Spider2-v: How far are multimodal agents from automating data science and engineering workflows?</a> In <i>Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024</i> .	
830		
831		
832		
833		
834		
835		
836		
837		
838		
839		
840		
841	Yupeng Cao, Aishwarya Muralidharan Nair, Elyon Eymife, Nastaran Jamalipour Soofi, K. P. Subbalakshmi, John R. Wullert II, Chumki Basu, and David Shallcross. 2024b. <a href="#">Can large language models detect misinformation in scientific news reporting?</a> <i>CoRR</i> , abs/2402.14268.	844
842		845
843		846
	Eric Chamoun, Michael Sejr Schlichtkrull, and Andreas Vlachos. 2024. <a href="#">Automated focused feedback generation for scientific writing assistance</a> . In <i>Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024</i> , pages 9742–9763. Association for Computational Linguistics.	847
		848
		849
		850
		851
		852
		853
	Tzeng-Ji Chen. 2023. Chatgpt and other artificial intelligence applications speed up scientific writing. <i>Journal of the Chinese Medical Association</i> , 86(4):351–353.	854
		855
		856
		857
	Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao, Chen Wei, Zitong Lu, Vishal Dey, Mingyi Xue, Frazier N. Baker, Benjamin Burns, Daniel Adu-Ampratwum, Xuhui Huang, Xia Ning, Song Gao, Yu Su, and Huan Sun. 2025. <a href="#">Scienceagentbench: Toward rigorous assessment of language agents for data-driven scientific discovery</a> . In <i>The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025</i> . OpenReview.net.	858
		859
		860
		861
		862
		863
		864
		865
		866
		867
	Gautam Choudhary, Natwar Modani, and Nitish Maurya. 2021. <a href="#">React: A review comment dataset for actionability (and more)</a> . In <i>Web Information Systems Engineering - WISE 2021 - 22nd International Conference on Web Information Systems Engineering, WISE 2021, Melbourne, VIC, Australia, October 26-29, 2021, Proceedings, Part II</i> , volume 13081 of <i>Lecture Notes in Computer Science</i> , pages 336–343. Springer.	868
		869
		870
		871
		872
		873
		874
		875
		876
	Julia Colyar. 2009. Becoming writing, becoming writers. <i>Qualitative Inquiry</i> , 15(2):421–436.	877
		878
	Paulo Henrique Couto, Quang Phuoc Ho, Nageeta Kumari, Benedictus Kent Rachmat, Thanh Gia Hieu Khuong, Ihsan Ullah, and Lisheng Sun-Hosoya. 2024. <a href="#">Relevai-reviewer: A benchmark on AI reviewers for survey paper relevance</a> . <i>CoRR</i> , abs/2406.10294.	879
		880
		881
		882
		883
	Robert Dale and Adam Kilgarrieff. 2010. Helping our own: Text massaging for computational linguistics as a new shared task. In <i>Proceedings of the 6th International Natural Language Generation Conference</i> .	884
		885
		886
		887
	Preetam Prabhu Srikar Dammu, Himanshu Naidu, Mouly Dewan, YoungMin Kim, Tanya Roosta, Aman Chadha, and Chirag Shah. 2024. <a href="#">Claimver: Explainable claim-level verification and evidence attribution of text through knowledge graphs</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024</i> , pages 13613–13627. Association for Computational Linguistics.	888
		889
		890
		891
		892
		893
		894
		895
		896
	Hai Dang, Chelse Swoopes, Daniel Buschek, and Elena L. Glassman. 2025. <a href="#">Corpusstudio: Surfacing emergent patterns in A corpus of prior work while writing</a> . In <i>Proceedings of the 2025 CHI Conference</i>	897
		898
		899
		900

901	on Human Factors in Computing Systems, CHI 2025, YokohamaJapan, 26 April 2025- 1 May 2025, pages 1211:1–1211:19. ACM.	956
902		957
903		958
904	Mike D’Arcy, Tom Hope, Larry Birnbaum, and Doug Downey. 2024a. <a href="#">MARG: multi-agent review generation for scientific papers</a> . <i>CoRR</i> , abs/2401.04259.	959
905		960
906		961
907	Mike D’Arcy, Alexis Ross, Erin Bransom, Bailey Kuehl, Jonathan Bragg, Tom Hope, and Doug Downey. 2024b. <a href="#">ARIES: A corpus of scientific paper edits made in response to peer reviews</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 6985–7001. Association for Computational Linguistics.	962
908		963
909		964
910		965
911		966
912		967
913		968
914		969
915		970
916	Maxime Darrin, Ines Arous, Pablo Piantanida, and Jackie Chi Kit Cheung. 2024. <a href="#">GLIMPSE: pragmatically informative multi-document summarization for scholarly reviews</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 12737–12752. Association for Computational Linguistics.	971
917		972
918		973
919		974
920		975
921		976
922		977
923		978
924		979
925	Anubrata Das, Houjiang Liu, Venelin Kovatchev, and Matthew Lease. 2023. <a href="#">The state of human-centered NLP technology for fact-checking</a> . <i>Inf. Process. Manag.</i> , 60(2):103219.	980
926		981
927		982
928		983
929	Vidas Daudaravicius. 2015. Automated evaluation of scientific writing: Aesw shared task proposal. In <i>Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications</i> , pages 56–63.	984
930		985
931		986
932		987
933		988
934	José de la Torre-López, Aurora Ramírez, and José Raúl Romero. 2023. <a href="#">Artificial intelligence to automate the systematic review of scientific literature</a> . <i>Computing</i> , 105(10):2171–2194.	989
935		990
936		991
937		992
938	DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 81 others. 2024. <a href="#">Deepseek-v3 technical report</a> . <i>CoRR</i> , abs/2412.19437.	993
939		994
940		995
941		996
942		997
943		998
944		999
945	Alphaeus Dmonte, Roland Oruche, Marcos Zampieri, Prasad Calyam, and Isabelle Augenstein. 2024. <a href="#">Claim verification in the age of large language models: A survey</a> . <i>CoRR</i> , abs/2408.14317.	1000
946		1001
947		1002
948		1003
949	Iddo Drori and Dov Te’eni. 2024. <a href="#">Human-in-the-loop AI reviewing: Feasibility, opportunities, and risks</a> . <i>J. Assoc. Inf. Syst.</i> , 25(1):7.	1004
950		1005
951		1006
952	Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022a. <a href="#">Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision</a> . <i>CoRR</i> , abs/2204.03685.	1007
953		1008
954		1009
955		1010
		1011
	Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, and Dongyeop Kang. 2022b. <a href="#">Understanding iterative revision from human-written text</a> . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 3573–3590. Association for Computational Linguistics.	
	Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 82 others. 2024. <a href="#">The llama 3 herd of models</a> . <i>CoRR</i> , abs/2407.21783.	
	Nils Dycke, Ilia Kuznetsov, and Iryna Gurevych. 2023. <a href="#">Nlpeer: A unified resource for the computational study of peer review</a> . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2023, Toronto, Canada, July 9-14, 2023, pages 5049–5073. Association for Computational Linguistics.	
	Steffen Eger, Yong Cao, Jennifer D’Souza, Andreas Geiger, Christian Greisinger, Stephanie Gross, Yufang Hou, Brigitte Krenn, Anne Lauscher, Yizhi Li, Chenghua Lin, Nafise Sadat Moosavi, Wei Zhao, and Tristan Miller. 2025. <a href="#">Transforming science with large language models: A survey on ai-assisted scientific discovery, experimentation, content generation, and evaluation</a> . <i>CoRR</i> , abs/2502.05151.	
	Islam Eldifrawi, Shengrui Wang, and Amine Trabelsi. 2024. <a href="#">Automated justification production for claim veracity in fact checking: A survey on architectures and approaches</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 6679–6692. Association for Computational Linguistics.	
	Noemie Elhadad, M-Y Kan, Judith L Klavans, and Kathleen R McKeown. 2005. Customization in a unified framework for summarizing medical literature. <i>Artificial intelligence in medicine</i> , 33(2):179–198.	
	Eurostat. 2018. <i>The measurement of scientific, technological and innovation activities Oslo manual 2018 guidelines for collecting, reporting and using data on innovation</i> . OECD publishing.	
	Benedikt Fecher, Marcel Hebing, Melissa Laufer, Jörg Pohle, and Fabian Sofsky. 2023. <a href="#">Friend or foe? exploring the implications of large language models on the science system</a> . <i>CoRR</i> , abs/2306.09928.	
	K. J. Kevin Feng, Kevin Pu, Matt Latzke, Tal August, Pao Siangliulue, Jonathan Bragg, Daniel S. Weld, Amy X. Zhang, and Joseph Chee Chang. 2024. <a href="#">Co-coa: Co-planning and co-execution with AI agents</a> . <i>CoRR</i> , abs/2412.10999.	



Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. 2023. <a href="#">Baldur: Whole-proof generation and repair with large language models</a> . In <i>Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2023, San Francisco, CA, USA, December 3-9, 2023</i> , pages 1229–1241. ACM.	62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 4372–4405. Association for Computational Linguistics.	1069 1070 1071 1072 1073
Martin Funkquist, Ilia Kuznetsov, Yufang Hou, and Iryna Gurevych. 2023. <a href="#">Citebench: A benchmark for scientific citation text generation</a> . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 7337–7353. Association for Computational Linguistics.	Mourad Gridach, Jay Nanavati, Khaldoun Zine El Abidine, Lenon Mendes, and Christina Mack. 2025. Agentic ai for scientific discovery: A survey of progress, challenges, and future directions. <i>arXiv preprint arXiv:2503.08979</i> .	1074 1075 1076 1077 1078
Conner Ganjavi, Michael B Eppler, Asli Pekcan, Brett Biedermann, Andre Abreu, Gary S Collins, Inderbir S Gill, and Giovanni E Cacciamani. 2024. Publishers’ and journals’ instructions to authors on use of generative artificial intelligence in academic and scientific publishing: bibliometric analysis. <i>bmj</i> , 384.	Dritjon Gruda. 2024. Three ways chatgpt helps me in my academic writing. <i>Nature</i> , 10.	1079 1080
Catherine A. Gao, Frederick M. Howard, Nikolay S. Markov, Emma C. Dyer, Siddhi Ramesh, Yuan Luo, and Alexander T. Pearson. 2023a. <a href="#">Comparing scientific abstracts generated by chatgpt to real abstracts with detectors and blinded human reviewers</a> . <i>npj Digit. Medicine</i> , 6.	Nianlong Gu and Richard Hahnloser. 2024. Controllable citation sentence generation with language models. In <i>Proceedings of the Fourth Workshop on Scholarly Document Processing (SDP 2024)</i> , pages 22–37. Association for Computational Linguistics.	1081 1082 1083 1084 1085
Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023b. <a href="#">Enabling large language models to generate text with citations</a> . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 6465–6488. Association for Computational Linguistics.	Nianlong Gu and Richard H. R. Hahnloser. 2023. <a href="#">Scilit: A platform for joint scientific literature discovery, summarization and citation generation</a> . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2023, Toronto, Canada, July 10-12, 2023</i> , pages 235–246. Association for Computational Linguistics.	1086 1087 1088 1089 1090 1091 1092
Alireza Ghafarollahi and Markus J. Buehler. 2024. <a href="#">Sci-agents: Automating scientific discovery through multi-agent intelligent graph reasoning</a> . <i>CoRR</i> , abs/2409.05556.	Xuemei Gu and Mario Krenn. 2024. <a href="#">Generation and human-expert evaluation of interesting research ideas using knowledge graphs and large language models</a> . <i>CoRR</i> , abs/2405.17044.	1093 1094 1095 1096
Karan Girotra, Lennart Meincke, Christian Terwiesch, and Karl T Ulrich. 2023. Ideas are dimes a dozen: Large language models for idea generation in innovation. <i>Available at SSRN 4526071</i> .	Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiron Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. <i>arXiv preprint arXiv:2501.12948</i> .	1097 1098 1099 1100 1101 1102
Max Glockner, Yufang Hou, and Iryna Gurevych. 2022. <a href="#">Missing counter-evidence renders NLP fact-checking unrealistic for misinformation</a> . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022</i> , pages 5916–5936. Association for Computational Linguistics.	Sikun Guo, Amir Hassan Shariatmadari, Guangzhi Xiong, Albert Huang, Eric Xie, Stefan Bekiranov, and Aidong Zhang. 2024. <a href="#">Ideabench: Benchmarking large language models for research idea generation</a> . <i>CoRR</i> , abs/2411.02429.	1103 1104 1105 1106 1107
Max Glockner, Yufang Hou, Preslav Nakov, and Iryna Gurevych. 2024a. <a href="#">Grounding fallacies misrepresenting scientific publications in evidence</a> . <i>CoRR</i> , abs/2408.12812.	Rohun Gupta, Isabel Herzog, Joseph Weisberger, John Chao, Kongkrit Chaayasate, and Edward S Lee. 2023. Utilization of chatgpt for plastic surgery research: friend or foe? <i>Journal of Plastic, Reconstructive &amp; Aesthetic Surgery</i> , 80:145–147.	1108 1109 1110 1111 1112
Max Glockner, Yufang Hou, Preslav Nakov, and Iryna Gurevych. 2024b. <a href="#">Misisci: Reconstructing fallacies in misrepresented science</a> . In <i>Proceedings of the</i>	Tarun Gupta and Danish Pruthi. 2025. <a href="#">All that glitters is not novel: Plagiarism in AI generated research</a> . <i>CoRR</i> , abs/2502.16487.	1113 1114 1115
	Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. <a href="#">Augmenting scientific papers with just-in-time, position-sensitive definitions of terms and symbols</a> . In <i>CHI ’21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021</i> , pages 413:1–413:18. ACM.	1116 1117 1118 1119 1120 1121 1122 1123



1238	<i>NeurIPS 2022, New Orleans, LA, USA, November 28</i>	Dongyeop Kang, Waleed Ammar, Bhavana Dalvi,	1295
1239	<i>- December 9, 2022.</i>	Madeleine van Zuylen, Sebastian Kohlmeier, Ed-	1296
1240	Albert Qiaochu Jiang, Sean Welleck, Jin Peng Zhou,	uard H. Hovy, and Roy Schwartz. 2018. <a href="#">A dataset</a>	1297
1241	Timothée Lacroix, Jiacheng Liu, Wenda Li, Mateja	<a href="#">of peer reviews (peerread): Collection, insights and</a>	1298
1242	Jamnik, Guillaume Lample, and Yuhuai Wu. 2023.	<a href="#">NLP applications</a> . In <i>Proceedings of the 2018 Con-</i>	1299
1243	<a href="#">Draft, sketch, and prove: Guiding formal theorem</a>	<i>ference of the North American Chapter of the Associ-</i>	1300
1244	<a href="#">provers with informal proofs</a> . In <i>The Eleventh Inter-</i>	<i>ation for Computational Linguistics: Human Lan-</i>	1301
1245	<i>national Conference on Learning Representations,</i>	<i>guage Technologies, NAACL-HLT 2018, New Or-</i>	1302
1246	<i>ICLR 2023, Kigali, Rwanda, May 1-5, 2023.</i> Open-	<i>leans, Louisiana, USA, June 1-6, 2018, Volume 1</i>	1303
1247	<i>Review.net.</i>	<i>(Long Papers)</i> , pages 1647–1661. Association for	1304
1248	Chao Jiang, Wei Xu, and Samuel Stevens. 2022b. <a href="#">arX-</a>	Computational Linguistics.	1305
1249	<a href="#">ivedits: Understanding the human revision process in</a>	Dongyeop Kang, Andrew Head, Risham Sidhu, Kyle	1306
1250	<a href="#">scientific writing</a> . In <i>Proceedings of the 2022 Con-</i>	Lo, Daniel S. Weld, and Marti A. Hearst. 2020.	1307
1251	<i>ference on Empirical Methods in Natural Language</i>	<a href="#">Document-level definition detection in scholarly doc-</a>	1308
1252	<i>Processing, EMNLP 2022, Abu Dhabi, United Arab</i>	<a href="#">uments: Existing models, error analyses, and future</a>	1309
1253	<i>Emirates, December 7-11, 2022</i> , pages 9420–9435.	<a href="#">directions</a> . In <i>Proceedings of the First Workshop</i>	1310
1254	Association for Computational Linguistics.	<i>on Scholarly Document Processing, SDP@EMNLP</i>	1311
1255	Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kai-	<i>2020, Online, November 19, 2020</i> , pages 196–206.	1312
1256	jie Zhu, Yijia Xiao, and Jindong Wang. 2024. <a href="#">Agen-</a>	Association for Computational Linguistics.	1313
1257	<a href="#">treview: Exploring peer review dynamics with LLM</a>	Hyeonsu B. Kang, Rafal Kocielnik, Andrew Head,	1314
1258	<a href="#">agents</a> . In <i>Proceedings of the 2024 Conference on</i>	Jiangjiang Yang, Matt Latzke, Aniket Kittur,	1315
1259	<i>Empirical Methods in Natural Language Processing,</i>	Daniel S. Weld, Doug Downey, and Jonathan Bragg.	1316
1260	<i>EMNLP 2024, Miami, FL, USA, November 12-16,</i>	2022. <a href="#">From who you know to what you read: Aug-</a>	1317
1261	2024, pages 1208–1226. Association for Computa-	<a href="#">menting scientific recommendations with implicit</a>	1318
1262	tional Linguistics.	<a href="#">social networks</a> . In <i>CHI '22: CHI Conference on</i>	1319
1263	Léane Jourdan, Florian Boudin, Richard Dufour, and	<i>Human Factors in Computing Systems, New Orleans,</i>	1320
1264	Nicolas Hernandez. 2023. <a href="#">Text revision in sci-</a>	<i>LA, USA, 29 April 2022 - 5 May 2022</i> , pages 302:1–	1321
1265	<a href="#">entific writing assistance: An overview</a> . <i>CoRR</i> ,	302:23. ACM.	1322
1266	abs/2303.16726.	Hyeonsu B. Kang, Nouran Soliman, Matt Latzke,	1323
1267	Léane Jourdan, Nicolas Hernandez, Richard Dufour,	Joseph Chee Chang, and Jonathan Bragg. 2023. <a href="#">Com-</a>	1324
1268	Florian Boudin, and Akiko Aizawa. 2025. Pararev:	<a href="#">mittee: Literature discovery with personal elected au-</a>	1325
1269	Building a dataset for scientific paragraph revision	<a href="#">thor committees</a> . In <i>Proceedings of the 2023 CHI</i>	1326
1270	annotated with revision instruction. <i>arXiv preprint</i>	<i>Conference on Human Factors in Computing Systems,</i>	1327
1271	<i>arXiv:2501.05222</i> .	<i>CHI 2023, Hamburg, Germany, April 23-28, 2023,</i>	1328
1272	Léane Isabelle Jourdan, Florian Boudin, Nicolas Her-	pages 738:1–738:20. ACM.	1329
1273	nandez, and Richard Dufour. 2024. <a href="#">CASIMIR: A</a>	Ying Kang, Aiqin Hou, Zimin Zhao, and Daguang Gan.	1330
1274	<a href="#">corpus of scientific articles enhanced with multiple</a>	2021. A hybrid approach for paper recommendation.	1331
1275	<a href="#">author-integrated revisions</a> . In <i>Proceedings of the</i>	<i>IEICE TRANSACTIONS on Information and Systems,</i>	1332
1276	<i>2024 Joint International Conference on Computa-</i>	104(8):1222–1231.	1333
1277	<i>tional Linguistics, Language Resources and Evalua-</i>	Wei-Yu Kao and An-Zi Yen. 2024. <a href="#">MAGIC: multi-</a>	1334
1278	<i>tion, LREC/COLING 2024, 20-25 May, 2024, Torino,</i>	<a href="#">argument generation with self-refinement for domain</a>	1335
1279	<i>Italy</i> , pages 2883–2892. ELRA and ICCL.	<a href="#">generalization in automatic fact-checking</a> . In <i>Pro-</i>	1336
1280	Shing-Yun Jung, Ting-Han Lin, Chia-Hung Liao, Shyan-	<i>ceedings of the 2024 Joint International Conference</i>	1337
1281	Ming Yuan, and Chuen-Tsai Sun. 2022. Intent-	<i>on Computational Linguistics, Language Resources</i>	1338
1282	controllable citation text generation. <i>Mathematics</i> ,	<i>and Evaluation, LREC/COLING 2024, 20-25 May,</i>	1339
1283	10(10):1763.	2024, <i>Torino, Italy</i> , pages 10891–10902. ELRA and	1340
1284	Melissa A Kacena, Lilian I Plotkin, and Jill C Fehren-	ICCL.	1341
1285	bacher. 2024. The use of artificial intelligence in	Tetsu Kasanishi, Masaru Isonuma, Junichiro Mori, and	1342
1286	writing scientific review articles. <i>Current Osteoporo-</i>	Ichiro Sakata. 2023. <a href="#">Scireviewgen: A large-scale</a>	1343
1287	<i>sis Reports</i> , 22(1):115–121.	<a href="#">dataset for automatic literature review generation</a> . In	1344
1288	Subbarao Kambhampati, Karthik Valmeekam, Lin	<i>Findings of the Association for Computational Lin-</i>	1345
1289	Guan, Mudit Verma, Kaya Stechly, Siddhant Bham-	<i>guistics: ACL 2023, Toronto, Canada, July 9-14,</i>	1346
1290	bri, Lucas Saldyt, and Anil Murthy. 2024. <a href="#">Position:</a>	2023, pages 6695–6715. Association for Computa-	1347
1291	<a href="#">Llms can't plan, but can help planning in llm-modulo</a>	tional Linguistics.	1348
1292	<a href="#">frameworks</a> . In <i>Forty-first International Conference</i>	Uri Katz, Mosh Levy, and Yoav Goldberg. 2024. <a href="#">Knowl-</a>	1349
1293	<i>on Machine Learning, ICML 2024, Vienna, Austria,</i>	<a href="#">edge navigator: Llm-guided browsing framework for</a>	1350
1294	<i>July 21-27, 2024.</i> OpenReview.net.	<a href="#">exploratory search in scientific literature</a> . In <i>Find-</i>	1351
		<i>ings of the Association for Computational Linguis-</i>	1352
		<i>tics: EMNLP 2024, Miami, Florida, USA, November</i>	1353



1354	12-16, 2024, pages 8838–8855. Association for Computational Linguistics.	1408
1355		1409
1356	Jiho Kim, Sungjin Park, Yeonsu Kwon, Yohan Jo, James Thorne, and Edward Choi. 2023. <a href="#">Factkg: Fact verification via reasoning on knowledge graphs</a> . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2023, Toronto, Canada, July 9-14, 2023, pages 16190–16206. Association for Computational Linguistics.	1410
1357		1411
1358		1412
1359		1413
1360		1414
1361		1415
1362		1416
1363		1417
1364	Zae Myung Kim, Wanyu Du, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022. <a href="#">Improving iterative text revision by learning where to edit from other revision tasks</a> . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022</i> , pages 9986–9999. Association for Computational Linguistics.	1418
1365		1419
1366		1420
1367		1421
1368		1422
1369		1423
1370		1424
1371		1425
1372		1426
1373	Ryan Koo, Anna Martin, Linghe Wang, and Dongyeop Kang. 2023. <a href="#">Decoding the end-to-end writing trajectory in scholarly manuscripts</a> . <i>CoRR</i> , abs/2304.00121.	1427
1374		1428
1375		1429
1376	Kayvan Kousha and Mike Thelwall. 2024. <a href="#">Artificial intelligence to support publishing and peer review: A summary and review</a> . <i>Learn. Publ.</i> , 37(1):4–12.	1430
1377		1431
1378		1432
1379	Mario Krenn, Lorenzo Buffoni, Bruno C. Coutinho, Sagi Eppel, Jacob Gates Foster, Andrew Gritsevskiy, Harlin Lee, Yichao Lu, João P. Moutinho, Nima Sanjabi, Rishi Sonthalia, Ngoc Mai Tran, Francisco Valente, Yangxinyu Xie, Rose Yu, and Michael Kopp. 2022. <a href="#">Predicting the future of AI with AI: high-quality link prediction in an exponentially growing knowledge network</a> . <i>CoRR</i> , abs/2210.00881.	1433
1380		1434
1381		1435
1382		1436
1383		1437
1384		1438
1385		1439
1386		1440
1387	Christin Katharina Kreutz and Ralf Schenkel. 2022. <a href="#">Scientific paper recommendation systems: a literature review of recent publications</a> . <i>Int. J. Digit. Libr.</i> , 23(4):335–369.	1441
1388		1442
1389		1443
1390		1444
1391	Amrith Krishna, Sebastian Riedel, and Andreas Vlachos. 2022. <a href="#">Proofver: Natural logic theorem proving for fact verification</a> . <i>Trans. Assoc. Comput. Linguistics</i> , 10:1013–1030.	1445
1392		1446
1393		1447
1394		1448
1395	Adithya Kulkarni, Fatimah Alotaibi, Xinyue Zeng, Longfeng Wu, Tong Zeng, Barry Menglong Yao, Minqian Liu, Shuaicheng Zhang, Lifu Huang, and Dawei Zhou. 2025. Scientific hypothesis generation and validation: Methods, datasets, and future directions. <i>arXiv preprint arXiv:2505.04651</i> .	1449
1396		1450
1397		1451
1398		1452
1399		1453
1400		1454
1401	Sandeep Kumar, Tirthankar Ghosal, and Asif Ekbal. 2023. <a href="#">When reviewers lock horns: Finding disagreements in scientific peer reviews</a> . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 16693–16704. Association for Computational Linguistics.	1455
1402		1456
1403		1457
1404		1458
1405		1459
1406		1460
1407		1461
	Sandeep Kumar, Tirthankar Ghosal, Vinayak Goyal, and Asif Ekbal. 2024. <a href="#">Can large language models unlock novel scientific research ideas?</a> <i>CoRR</i> , abs/2409.06185.	1462
		1463
		1464
	Ilia Kuznetsov, Osama Mohammed Afzal, Koen Dercksen, Nils Dycke, Alexander Goldberg, Tom Hope, Dirk Hovy, Jonathan K. Kummerfeld, Anne Lauscher, Kevin Leyton-Brown, Sheng Lu, Mausam, Margot Mieskes, Aurélie Névéol, Danish Pruthi, Lizhen Qu, Roy Schwartz, Noah A. Smith, Thamar Solorio, and 5 others. 2024. <a href="#">What can natural language processing do for peer review?</a> <i>CoRR</i> , abs/2405.06563.	1465
		1466
		1467
		1468
		1469
		1470
		1471
		1472
		1473
		1474
		1475
		1476
		1477
		1478
		1479
		1480
		1481
		1482
		1483
		1484
		1485
		1486
		1487
		1488
		1489
		1490
		1491
		1492
		1493
		1494
		1495
		1496
		1497
		1498
		1499
		1500

1465	Jiang, Yifei Xin, Ronghao Dang, Deli Zhao, Yu Rong,	content at scale: A case study on the impact of chat-	1520
1466	Tian Feng, and Lidong Bing. 2024a. <a href="#">Chain of ideas: Revolutionizing research via novel idea development with LLM agents</a> . <i>CoRR</i> , abs/2410.13185.	gpt on AI conference peer reviews. In <i>Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024</i> . OpenReview.net.	1521
1467			1522
1468			1523
1469	Miao Li, Eduard H. Hovy, and Jey Han Lau. 2023. <a href="#">Summarizing multiple documents with conversational structure for meta-review generation</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 7089–7112. Association for Computational Linguistics.		1524
1470		Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D. Manning, and James Y. Zou. 2024b. <a href="#">Mapping the increasing use of llms in scientific papers</a> . <i>CoRR</i> , abs/2404.01268.	1525
1471			1526
1472			1527
1473			1528
1474			1529
1475			1530
1476	Miao Li, Jey Han Lau, and Eduard H. Hovy. 2024b. <a href="#">A sentiment consolidation framework for meta-review generation</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024</i> , pages 10158–10177. Association for Computational Linguistics.	Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu Wang, Daisy Ding, Xinyu Yang, Kailas Vodrahalli, Siyu He, Daniel Scott Smith, Yian Yin, Daniel A. McFarland, and James Zou. 2023. <a href="#">Can large language models provide useful feedback on research papers? A large-scale empirical analysis</a> . <i>CoRR</i> , abs/2310.01783.	1531
1477			1532
1478			1533
1479			1534
1480			1535
1481			1536
1482			1537
1483	Ruochen Li, Liqiang Jing, Chi Han, Jiawei Zhou, and Xinya Du. 2024c. <a href="#">Learning to generate research idea with dynamic control</a> . <i>CoRR</i> , abs/2412.14626.	Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X. Zhang. 2024. <a href="#">Llms as research tools: A large scale survey of researchers' usage and perceptions</a> . <i>CoRR</i> , abs/2411.05025.	1538
1484			1539
1485			1540
1486	Ruochen Li, Teerth Patel, Qingyun Wang, and Xinya Du. 2024d. <a href="#">Mlr-copilot: Autonomous machine learning research based on large language models agents</a> . <i>CoRR</i> , abs/2408.14033.		1541
1487			1542
1488			
1489			
1490	Weisheng Li, Chao Chang, Chaobo He, Zhengyang Wu, Jiongsheng Guo, and Bo Peng. 2020. <a href="#">Academic paper recommendation method combining heterogeneous network and temporal attributes</a> . In <i>Computer Supported Cooperative Work and Social Computing - 15th CCF Conference, ChineseCSCW 2020, Shenzhen, China, November 7-9, 2020, Revised Selected Papers</i> , volume 1330 of <i>Communications in Computer and Information Science</i> , pages 456–468. Springer.	Haohan Lin, Zhiqing Sun, Yiming Yang, and Sean Welleck. 2024. <a href="#">Lean-star: Learning to interleave thinking and proving</a> . <i>CoRR</i> , abs/2407.10040.	1543
1491			1544
1492			1545
1493		Jialiang Lin, Jiaxin Song, Zhangping Zhou, Yidong Chen, and Xiaodong Shi. 2023a. <a href="#">Automated scholarly paper review: Concepts, technologies, and challenges</a> . <i>Inf. Fusion</i> , 98:101830.	1546
1494			1547
1495			1548
1496			1549
1497		Jialiang Lin, Jiaxin Song, Zhangping Zhou, Yidong Chen, and Xiaodong Shi. 2023b. <a href="#">MOPRD: A multi-disciplinary open peer review dataset</a> . <i>Neural Comput. Appl.</i> , 35(34):24191–24206.	1550
1498			1551
1499			1552
1500	Xiangci Li and Jessica Ouyang. 2022. <a href="#">Automatic related work generation: A meta study</a> . <i>CoRR</i> , abs/2201.01880.	Zhicheng Lin. 2024. Techniques for supercharging academic writing with generative ai. <i>Nature Biomedical Engineering</i> , pages 1–6.	1553
1501			1554
1502			1555
1503	Xiangci Li and Jessica Ouyang. 2024. <a href="#">Related work and citation text generation: A survey</a> . In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024</i> , pages 13846–13864. Association for Computational Linguistics.	Chengwu Liu, Jianhao Shen, Huajian Xin, Zhengying Liu, Ye Yuan, Haiming Wang, Wei Ju, Chuanyang Zheng, Yichun Yin, Lin Li, Ming Zhang, and Qun Liu. 2023a. <a href="#">FIMO: A challenge formal dataset for automated theorem proving</a> . <i>CoRR</i> , abs/2309.04295.	1556
1504			1557
1505			1558
1506			1559
1507			1560
1508			1561
1509	Zhaoyu Li, Jialiang Sun, Logan Murphy, Qidong Su, Zenan Li, Xian Zhang, Kaiyu Yang, and Xujie Si. 2024e. <a href="#">A survey on deep learning for theorem proving</a> . <i>CoRR</i> , abs/2404.09939.	Haokun Liu, Yangqiaoyu Zhou, Mingxuan Li, Chenfei Yuan, and Chenhao Tan. 2024a. <a href="#">Literature meets data: A synergistic approach to hypothesis generation</a> . <i>CoRR</i> , abs/2410.17309.	1562
1510			1563
1511			1564
1512			1565
1513	Zhi Li and Xiaozhu Zou. 2019. <a href="#">A review on personalized academic paper recommendation</a> . <i>Comput. Inf. Sci.</i> , 12(1):33–43.	Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Denny Zhou, Andrew M. Dai, Diyi Yang, and Soroush Vosoughi. 2023b. <a href="#">Training socially aligned language models in simulated human society</a> . <i>CoRR</i> , abs/2305.16960.	1566
1514			1567
1515			1568
1516			1569
1517	Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou. 2024a. <a href="#">Monitoring ai-modified</a>	Ryan Liu and Nihar B. Shah. 2023. <a href="#">Reviewergpt? an exploratory study on using large language models for paper reviewing</a> . <i>CoRR</i> , abs/2306.00622.	1570
1518			1571
1519			1572
			1573

1574	Shengchao Liu, Jiong Xiao Wang, Yijin Yang, Cheng-	simulation as bilevel optimizers: A new paradigm	1631
1575	peng Wang, Ling Liu, Hongyu Guo, and Chaowei	to advance physical scientific discovery. In <i>Forty-</i>	1632
1576	Xiao. 2024b. <a href="#">Conversational drug editing using re-</a>	<i>first International Conference on Machine Learning,</i>	1633
1577	<a href="#">trieval and domain feedback</a> . In <i>The Twelfth Inter-</i>	<i>ICML 2024, Vienna, Austria, July 21-27, 2024.</i> Open-	1634
1578	<i>national Conference on Learning Representations,</i>	<i>Review.net.</i>	1635
1579	<i>ICLR 2024, Vienna, Austria, May 7-11, 2024.</i> Open-		
1580	<i>Review.net.</i>		
1581	Shuaiqi Liu, Jiannong Cao, Ruosong Yang, and Zhiyuan	Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler	1636
1582	Wen. 2022. <a href="#">Generating a structured summary of</a>	Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon,	1637
1583	<a href="#">numerous academic papers: Dataset and method</a> . In	Nouha Dziri, Shrimai Prabhumoye, Yiming Yang,	1638
1584	<i>Proceedings of the Thirty-First International Joint</i>	Shashank Gupta, Bodhisattwa Prasad Majumder,	1639
1585	<i>Conference on Artificial Intelligence, IJCAI 2022,</i>	Katherine Hermann, Sean Welleck, Amir Yazdan-	1640
1586	<i>Vienna, Austria, 23-29 July 2022, pages 4259–4265.</i>	bakhsh, and Peter Clark. 2023. <a href="#">Self-refine: Itera-</a>	1641
1587	<i>ijcai.org.</i>	<a href="#">tive refinement with self-feedback</a> . In <i>Advances in</i>	1642
1588		<i>Neural Information Processing Systems 36: Annual</i>	1643
1589	Siyi Liu, Chen Gao, and Yong Li. 2024c. <a href="#">Large lan-</a>	<i>Conference on Neural Information Processing Sys-</i>	1644
1590	<a href="#">guage model agent for hyper-parameter optimization.</a>	<i>tems 2023, NeurIPS 2023, New Orleans, LA, USA,</i>	1645
1591	<i>CoRR, abs/2402.01881.</i>	<i>December 10 - 16, 2023.</i>	1646
1592	Zijun Liu, Kaiming Liu, Yiqi Zhu, Xuanyu Lei, Zong-	Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv	1647
1593	han Yang, Zhenhe Zhang, Peng Li, and Yang Liu.	Agarwal, Sanchaita Hazra, Ashish Sabharwal, and	1648
1594	2024d. <a href="#">AIGS: generating science from ai-powered</a>	Peter Clark. 2024a. <a href="#">Position: Data-driven discovery</a>	1649
1595	<a href="#">automated falsification</a> . <i>CoRR, abs/2411.11910.</i>	<a href="#">with large generative models</a> . In <i>Forty-first Interna-</i>	1650
1596		<i>tional Conference on Machine Learning, ICML 2024,</i>	1651
1597	Kyle Lo, Joseph Chee Chang, Andrew Head, Jonathan	<i>Vienna, Austria, July 21-27, 2024.</i> OpenReview.net.	1652
1598	Bragg, Amy X. Zhang, Cassidy Trier, Chloe Anas-		
1599	tasiades, Tal August, Russell Authur, Danielle Bragg,	Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv	1653
1600	Erin Bransom, Isabel Cachola, Stefan Candra, Yo-	Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh	1654
1601	ganand Chandrasekhar, Yen-Sung Chen, Evie Yu-	Meena, Aryan Prakhar, Tirth Vora, Tushar Khot,	1655
1602	Yen Cheng, Yvonne Chou, Doug Downey, Rob	Ashish Sabharwal, and Peter Clark. 2024b. <a href="#">Discov-</a>	1656
1603	Evans, and 36 others. 2023. <a href="#">The semantic reader</a>	<a href="#">erybench: Towards data-driven discovery with large</a>	1657
1604	<a href="#">project: Augmenting scholarly documents through</a>	<a href="#">language models</a> . <i>CoRR, abs/2407.01725.</i>	1658
1605	<a href="#">ai-powered interactive reading interfaces</a> . <i>CoRR,</i>		
1606	<i>abs/2303.14334.</i>	Benjamin S Manning, Kehang Zhu, and John J Horton.	1659
1607	Renze Lou, Hanzi Xu, Sijia Wang, Jiangshu Du,	2024. Automated social science: Language models	1660
1608	Ryo Kamoi, Xiaoxin Lu, Jian Xie, Yuxuan Sun,	as scientist and subjects. Technical report, National	1661
1609	Yusen Zhang, Jihyun Janice Ahn, Hongchao Fang,	Bureau of Economic Research.	1662
1610	Zhuoyang Zou, Wenchao Ma, Xi Li, Kai Zhang, Con-		
1611	gying Xia, Lifu Huang, and Wenpeng Yin. 2024.	Anna Martin-Boyle, Aahan Tyagi, Marti A. Hearst,	1663
1612	<a href="#">AAAR-1.0: assessing ai’s potential to assist research.</a>	and Dongyeop Kang. 2024. <a href="#">Shallow synthesis of</a>	1664
1613	<i>CoRR, abs/2410.22394.</i>	<a href="#">knowledge in gpt-generated texts: A case study</a>	1665
1614		<a href="#">in automatic related work composition</a> . <i>CoRR,</i>	1666
1615		<i>abs/2402.12255.</i>	1667
1616	Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foer-	Rui Meng, Khushboo Thaker, Lei Zhang, Yue Dong,	1668
1617	ster, Jeff Clune, and David Ha. 2024. <a href="#">The AI scien-</a>	Xingdi Yuan, Tong Wang, and Daqing He. 2021.	1669
1618	<a href="#">tist: Towards fully automated open-ended scientific</a>	<a href="#">Bringing structure into summaries: a faceted sum-</a>	1670
1619	<a href="#">discovery</a> . <i>CoRR, abs/2408.06292.</i>	<a href="#">marization dataset for long scientific documents</a> . In	1671
1620		<i>Proceedings of the 59th Annual Meeting of the Asso-</i>	1672
1621	Xinyuan Lu, Liangming Pan, Qian Liu, Preslav Nakov,	<i>ciation for Computational Linguistics and the 11th</i>	1673
1622	and Min-Yen Kan. 2023. <a href="#">SCITAB: A challenging</a>	<i>International Joint Conference on Natural Language</i>	1674
1623	<a href="#">benchmark for compositional reasoning and claim</a>	<i>Processing, ACL/IJCNLP 2021, (Volume 2: Short</i>	1675
1624	<a href="#">verification on scientific tables</a> . In <i>Proceedings of the</i>	<i>Papers), Virtual Event, August 1-6, 2021, pages 1080–</i>	1676
1625	<i>2023 Conference on Empirical Methods in Natural</i>	<i>1089.</i> Association for Computational Linguistics.	1677
1626	<i>Language Processing, EMNLP 2023, Singapore, De-</i>		
1627	<i>cember 6-10, 2023, pages 7787–7813.</i> Association	Lisa Messeri and MJ Crockett. 2024. Artificial intel-	1678
1628	for Computational Linguistics.	ligence and illusions of understanding in scientific	1679
1629		research. <i>Nature</i> , 627(8002):49–58.	1680
1630	Ziming Luo, Zonglin Yang, Zexin Xu, Wei Yang, and	Meredith Ringel Morris. 2023. <a href="#">Scientists’ perspec-</a>	1681
1631	Xinya Du. 2025. <a href="#">Llm4sr: A survey on large lan-</a>	<a href="#">tives on the potential for generative AI in their fields.</a>	1682
1632	<a href="#">guage models for scientific research</a> . <i>arXiv preprint</i>	<i>CoRR, abs/2304.01420.</i>	1683
1633	<i>arXiv:2501.04306.</i>		
1634		Xinyi Mou, Xuanwen Ding, Qi He, Liang Wang, Jing-	1684
1635		cong Liang, Xinnong Zhang, Libo Sun, Jiayu Lin,	1685
1636		Jie Zhou, Xuanjing Huang, and Zhongyu Wei. 2024.	1686



1687 From individual to society: A survey on social simulation driven by large language model-based agents. *CoRR*, abs/2412.03563. 1740

1688 1741

1689 1742

1690 Panitan Muangkammuen, Fumiyo Fukumoto, Jiyi Li, and Yoshimi Suzuki. 2022. Exploiting labeled and unlabeled data via transformer fine-tuning for peer-review score prediction. In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 2233–2240. Association for Computational Linguistics. 1744

1691 1745

1692 1746

1693 1747

1694 1748

1695 1749

1696 1750

1697 1751

1698 Arief Purnama Muharram and Ayu Purwarianti. 2024. Enhancing natural language inference performance with knowledge graph for COVID-19 automated fact-checking in Indonesian language. *CoRR*, abs/2409.00061. 1752

1699 1753

1700 1754

1701 1755

1702 1756

1703 Bo Ni and Markus J. Buehler. 2023. Mechagents: Large language model multi-agent collaborations can solve mechanics problems, generate new data, and integrate knowledge. *CoRR*, abs/2311.08166. 1757

1704 1758

1705 1759

1706 1760

1707 Ziqi Ni, Yahao Li, Kaijia Hu, Kunyuan Han, Ming Xu, Xingyu Chen, Fengqi Liu, Yicong Ye, and Shuxin Bai. 2024. Matpilot: an llm-enabled AI materials scientist under the framework of human-machine collaboration. *CoRR*, abs/2411.08063. 1761

1708 1762

1709 1763

1710 1764

1711 1765

1712 Harshit Nigam, Manasi Patwardhan, Lovekesh Vig, and Gautam Shroff. 2024. Acceleron: A tool to accelerate research ideation. *CoRR*, abs/2403.04382. 1766

1713 1767

1714 1768

1715 Kazuya Nishimura, Kuniaki Saito, Tosho Hirasawa, and Yoshitaka Ushiku. 2024. Toward related work generation with structure and novelty statement. In *Proceedings of the Fourth Workshop on Scholarly Document Processing (SDP 2024)*, pages 38–57. 1770

1716 1771

1717 1772

1718 1773

1719 1774

1720 Liang Niu, Nian Xue, and Christina Pöpper. 2023. Unveiling the sentinels: Assessing AI performance in cybersecurity peer review. *CoRR*, abs/2309.05457. 1775

1721 1776

1722 1777

1723 1778

1724 1779

1725 OpenAI. 2023. GPT-4 technical report. *CoRR*, abs/2303.08774. 1780

1726 1781

1727 1782

1728 1783

1729 1784

1730 1785

1731 1786

1732 1787

1733 1788

1734 1789

1735 1790

1736 1791

1737 1792

1738 1793

1739 1794

Liangming Pan, Yunxiang Zhang, and Min-Yen Kan. 2023b. Investigating zero- and few-shot generalization in fact verification. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, IJCNLP 2023 - Volume 1: Long Papers, Nusa Dua, Bali, November 1 - 4, 2023*, pages 511–524. Association for Computational Linguistics. 1795

Pat Pataranutaporn, Nattavudh Powdthavee, and Pattie Maes. 2025. Can AI solve the peer review crisis? A large scale experiment on llm’s performance and biases in evaluating economics papers. *CoRR*, abs/2502.00070. 1796

Alison Pease, Simon Colton, Chris Warburton, Athanasios Nathanail, Irina Preda, Daniel Arnold, Daniel Winterstein, and Mike Cook. 2019. The importance of applying computational creativity to scientific and mathematical domains. In *10th International Conference on Computational Creativity, ICC3 2019*, pages 250–257. Association for Computational Creativity.

Iratxe Pinedo, Mikel Larrañaga, and Ana Arruarte. 2024. Arzigo: A recommendation system for scientific articles. *Inf. Syst.*, 122:102367.

Milton Pividori and Casey S. Greene. 2024. A publishing infrastructure for artificial intelligence (ai)-assisted academic authoring. *J. Am. Medical Informatics Assoc.*, 31(9):2103–2113.

Andrea Polonioli. 2021. The ethics of scientific recommender systems. *Scientometrics*, 126(2):1841–1848.

Stanislas Polu and Ilya Sutskever. 2020. Generative language modeling for automated theorem proving. *CoRR*, abs/2009.03393.

Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2024. Ideasynt: Iterative research idea development through evolving and composing idea facets with literature-grounded feedback. *CoRR*, abs/2410.04025.

Biqing Qi, Kaiyan Zhang, Haoxiang Li, Kai Tian, Sihang Zeng, Zhang-Ren Chen, and Bowen Zhou. 2023. Large language models are zero shot hypothesis proposers. *CoRR*, abs/2311.05965.

Hanhao Qu, Yu Cao, Jun Gao, Liang Ding, and Ruifeng Xu. 2022. Interpretable proof generation via iterative backward reasoning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2968–2981. Association for Computational Linguistics.

Gollam Rabby, Diyana Muhammed, Prasenjit Mitra, and Sören Auer. 2025. Iterative hypothesis generation for scientific discovery with monte carlo nash equilibrium self-refining trees. *arXiv preprint arXiv:2503.19309*.

Dragomir R. Radev, Pradeep Muthukrishnan, Vahed Qazvinian, and Amjad Abu-Jbara. 2013. [The ACL anthology network corpus](#). *Lang. Resour. Evaluation*, 47(4):919–944.

Zachary Robertson. 2023. [GPT4 is slightly helpful for peer-review assistance: A pilot study](#). *CoRR*, abs/2307.05492.

Ambrose Robinson, William Thorne, Ben P. Wu, Abdullah Pandor, Munira Essat, Mark Stevenson, and Xingyi Song. 2023. [Bio-sieve: Exploring instruction tuning large language models for systematic review automation](#). *CoRR*, abs/2308.06610.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, and 1 others. 2024. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475.

Kai Ruan, Xuan Wang, Jixiang Hong, and Hao Sun. 2024a. [Liveideabench: Evaluating llms’ scientific creativity and idea generation with minimal context](#). *CoRR*, abs/2412.17596.

Yixiang Ruan, Chenyin Lu, Ning Xu, Yuchen He, Yixin Chen, Jian Zhang, Jun Xuan, Jianzhang Pan, Qun Fang, Hanyu Gao, and 1 others. 2024b. [An automatic end-to-end chemical synthesis development platform powered by large language models](#). *Nature communications*, 15(1):10160.

Michele Salvagno, Fabio Silvio Taccone, and Alberto Giovanni Gerli. 2023. Can artificial intelligence help for scientific writing? *Critical care*, 27(1):75.

Shubhra Kanti Karmaker Santu, Sanjeev Kumar Sinha, Naman Bansal, Alex Knipper, Souvika Sarkar, John Salvador, Yash Mahajan, Sri Guttikonda, Mousumi Akter, Matthew Freestone, and Matthew C. Williams Jr. 2024. [Prompting llms to compose meta-review drafts from peer-review narratives of scholarly manuscripts](#). *CoRR*, abs/2402.15589.

Mourad Sarrouiti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. [Evidence-based fact-checking of health-related claims](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 3499–3512. Association for Computational Linguistics.

Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2023. [Are emergent abilities of large language models a mirage?](#) In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Laurie A. Schintler, Connie L. McNeely, and James Witte. 2023. [A critical examination of the ethics of ai-mediated peer review](#). *CoRR*, abs/2309.12356.

Samuel Schmidgall and Michael Moor. 2025. [Agent-trxiv: Towards collaborative autonomous research](#). *arXiv preprint arXiv:2503.18102*.

Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Zicheng Liu, and Emad Barsoum. 2025. [Agent laboratory: Using llm agents as research assistants](#). *arXiv preprint arXiv:2501.04227*.

Yakub Sebastian, Eu-Gen Siew, and Sylvester O. Ori-maye. 2017. [Emerging approaches in literature-based discovery: techniques and performance review](#). *The Knowledge Engineering Review*, 32:e12.

Basuki Setio and Masatoshi Tsuchiya. 2022. [The quality assist: A technology-assisted peer review based on citation functions to predict the paper quality](#). *IEEE Access*, 10:126815–126831.

Abdul Shahid, Muhammad Tanvir Afzal, Moloud Abdar, Mohammad Ehsan Basiri, Xujuan Zhou, Neil Y Yen, and Jia-Wei Chang. 2020. Insights into relevant knowledge extraction techniques: a comprehensive review. *The Journal of Supercomputing*, 76:1695–1733.

Yijia Shao, Yucheng Jiang, Theodore A. Kanell, Peter Xu, Omar Khattab, and Monica S. Lam. 2024. [Assisting in writing wikipedia-like articles from scratch with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 6252–6278. Association for Computational Linguistics.

Ritu Sharma, Dinesh Gopalani, and Yogesh Kumar Meena. 2023. [An anatomization of research paper recommender system: Overview, approaches and challenges](#). *Eng. Appl. Artif. Intell.*, 118:105641.

Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. 2022. [Mred: A meta-review dataset for structure-controllable text generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2521–2535. Association for Computational Linguistics.

Yongliang Shen, Kaitao Song, Xu Tan, Wenqi Zhang, Kan Ren, Siyu Yuan, Weiming Lu, Dongsheng Li, and Yueting Zhuang. 2024. [Taskbench: Benchmarking large language models for task automation](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.

Zhengliang Shi, Shen Gao, Zhen Zhang, Xiuying Chen, Zhumin Chen, Pengjie Ren, and Zhaochun Ren. 2023. [Towards a unified framework for reference retrieval and related work generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 5785–5799. Association for Computational Linguistics.

1912	Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2024.	Cheng Tan, Dongxin Lyu, Siyuan Li, Zhangyang Gao,	1967
1913	<a href="#">Can llms generate novel research ideas? A large-scale</a>	Jingxuan Wei, Siqu Ma, Zicheng Liu, and Stan Z.	1968
1914	<a href="#">human study with 100+ NLP researchers.</a>	Li. 2024. <a href="#">Peer review as A multi-turn and long-</a>	1969
1915	<i>CoRR</i> , abs/2409.04109.	<a href="#">context dialogue with role-based interactions.</a>	1970
		<i>CoRR</i> , abs/2406.05688.	1971
1916	Zachary S. Siegel, Sayash Kapoor, Nitya Nagdir,	Xiangru Tang, Xingyao Zhang, Yanjun Shao, Jie Wu,	1972
1917	Benedikt Stroebl, and Arvind Narayanan. 2024.	Yilun Zhao, Arman Cohan, Ming Gong, Dongmei	1973
1918	<a href="#">Core-bench: Fostering the credibility of published re-</a>	Zhang, and Mark Gerstein. 2024a. <a href="#">Step-back profil-</a>	1974
1919	<a href="#">search through a computational reproducibility agent</a>	<a href="#">ing: Distilling user history for personalized scientific</a>	1975
1920	<a href="#">benchmark.</a>	<i>CoRR</i> , abs/2406.14275.	1976
	<i>CoRR</i> , abs/2409.11363.		
1921	Michael D. Skarlinski, Sam Cox, Jon M. Laurent,	Xuemei Tang, Xufeng Duan, and Zhenguang G Cai.	1977
1922	James D. Braza, Michaela M. Hinks, Michael J.	2024b. Are llms good literature review writers? eval-	1978
1923	Hammerling, Manvitha Ponnampati, Samuel G. Ro-	uating the literature review writing ability of large	1979
1924	driques, and Andrew D. White. 2024. <a href="#">Language</a>	language models. <i>arXiv preprint arXiv:2412.13612</i> .	1980
1925	<a href="#">agents achieve superhuman synthesis of scientific</a>		
1926	<a href="#">knowledge.</a>	Min Tao, Xinmin Yang, Gao Gu, and Bohan Li. 2020.	1981
	<i>CoRR</i> , abs/2409.13740.	Paper recommend based on lda and pagerank. In	1982
1927	Peiyang Song, Kaiyu Yang, and Anima Anandkumar.	<i>Artificial Intelligence and Security: 6th International</i>	1983
1928	2024. <a href="#">Towards large language models as copilots for</a>	<i>Conference, ICAIS 2020, Hohhot, China, July 17–</i>	1984
1929	<a href="#">theorem proving in lean.</a>	<i>20, 2020, Proceedings, Part III</i> 6, pages 571–584.	1985
	<i>CoRR</i> , abs/2404.12534.	Springer.	1986
1930	Giulio Starace, Oliver Jaffe, Dane Sherburn, James	Nitya Thakkar, Mert Yuksekgonul, Jake Silberg, Ani-	1987
1931	Aung, Jun Shern Chan, Leon Maksin, Rachel Dias,	mesh Garg, Nanyun Peng, Fei Sha, Rose Yu,	1988
1932	Evan Mays, Benjamin Kinsella, Wyatt Thompson,	Carl Vondrick, and James Zou. 2025. <a href="#">Can llm</a>	1989
1933	and 1 others. 2025. <a href="#">Paperbench: Evaluating ai’s</a>	<a href="#">feedback enhance review quality? a randomized</a>	1990
1934	<a href="#">ability to replicate ai research.</a>	<i>study of 20k reviews at iclr 2025.</i>	1991
1935	<i>arXiv preprint arXiv:2504.01848</i> .	<i>arXiv preprint arXiv:2504.09737</i> .	1992
1936	Vaios Stergiopoulos, Michael Vassilakopoulos, Eleni	Amitayush Thakur, George Tsoukalas, Yeming Wen,	1993
1937	Tousidou, and Antonio Corral. 2024. <a href="#">An academic</a>	Jimmy Xin, and Swarat Chaudhuri. 2024. <a href="#">An in-</a>	1994
1938	<a href="#">recommender system on large citation data based on</a>	<a href="#">context learning agent for formal theorem-proving.</a>	1995
1939	<a href="#">clustering, graph modeling and deep learning.</a>	In <i>First Conference on Language Modeling</i> .	1996
1940	<i>Knowl. Inf. Syst.</i> , 66(8):4463–4496.	Mike Thelwall and Abdullah Yaghi. 2024. <a href="#">Evaluating</a>	1997
1941	Haoyang Su, Renqi Chen, Shixiang Tang, Xinzhe	<a href="#">the predictive capacity of chatgpt for academic peer</a>	1998
1942	Zheng, Jingzhe Li, Zhenfei Yin, Wanli Ouyang, and	<a href="#">review outcomes across multiple platforms.</a>	1999
1943	Nanqing Dong. 2024. <a href="#">Two heads are better than one:</a>	<i>CoRR</i> , abs/2411.09763.	2000
1944	<a href="#">A multi-agent system has the potential to improve</a>		
1945	<a href="#">scientific idea generation.</a>	James Thorne, Andreas Vlachos, Christos	2001
	<i>CoRR</i> , abs/2410.09403.	Christodoulopoulos, and Arpit Mittal. 2018.	2002
1946	Purin Sukpanichnant, Anna Rapberger, and Francesca	<a href="#">FEVER: a large-scale dataset for fact extraction</a>	2003
1947	Toni. 2024. <a href="#">Peerarg: Argumentative peer review with</a>	<a href="#">and verification.</a>	2004
1948	<a href="#">llms.</a>	In <i>Proceedings of the 2018</i>	
	<i>CoRR</i> , abs/2409.16813.	<i>Conference of the North American Chapter of the</i>	2005
1949	Lu Sun, Stone Tao, Junjie Hu, and Steven P. Dow. 2024.	<i>Association for Computational Linguistics: Human</i>	2006
1950	<a href="#">Metawriter: Exploring the potential and perils of AI</a>	<i>Language Technologies, NAACL-HLT 2018, New</i>	2007
1951	<a href="#">writing support in scientific peer review.</a>	<i>Orleans, Louisiana, USA, June 1-6, 2018, Volume</i>	2008
1952	<i>Proc. ACM Hum. Comput. Interact.</i> , 8(CSCW1):1–32.	<i>1 (Long Papers)</i> , pages 809–819. Association for	2009
1953	Teo Susnjak, Peter Hwang, Napoleon H. Reyes, Andre	Computational Linguistics.	2010
1954	L. C. Barczak, Timothy R. McIntosh, and Surangika	Yangjie Tian, Xungang Gu, Aijia Li, He Zhang, Ruohua	2011
1955	Ranathunga. 2024. <a href="#">Automating research synthe-</a>	Xu, Yunfeng Li, and Ming Liu. 2024. <a href="#">Overview of</a>	2012
1956	<a href="#">sis with domain-specific large language model fine-</a>	<a href="#">the NLPCC2024 shared task 6: Scientific literature</a>	2013
1957	<a href="#">tuning.</a>	<a href="#">survey generation.</a>	2014
	<i>CoRR</i> , abs/2404.08680.	In <i>Natural Language Processing</i>	
1958	Don R. Swanson. 1986. <a href="#">Undiscovered public knowl-</a>	<i>and Chinese Computing - 13th National CCF Con-</i>	2015
1959	<a href="#">edge.</a>	<i>ference, NLPCC 2024, Hangzhou, China, November</i>	2016
1960	<i>The Library Quarterly: Information, Commu-</i>	<i>1-3, 2024, Proceedings, Part V</i> , volume 15363 of	2017
	<i>nity, Policy</i> , 56(2):103–118.	<i>Lecture Notes in Computer Science</i> , pages 400–408.	2018
1961	Nathan J Szymanski, Bernardus Rendy, Yuxing Fei,	Springer.	2019
1962	Rishi E Kumar, Tanjin He, David Milsted, Matthew J	Venktesh V, Abhijit Anand, Avishek Anand, and Vinay	2020
1963	McDermott, Max Gallant, Ekin Dogus Cubuk, Amil	Setty. 2024. <a href="#">Quantemp: A real-world open-domain</a>	2021
1964	Merchant, and 1 others. 2023. <a href="#">An autonomous labo-</a>	<a href="#">benchmark for fact-checking numerical claims.</a>	2022
1965	<a href="#">ratory for the accelerated synthesis of novel materials.</a>	In <i>Proceedings of the 47th International ACM SIGIR</i>	2023
1966	<i>Nature</i> , 624(7990):86–91.		



2024	Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024, pages 650–660. ACM.	2081
2025		2082
2026		2083
2027	Juraj Vladika and Florian Matthes. 2023. <a href="#">Scientific fact-checking: A survey of resources and approaches</a> .	2084
2028	In <i>Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023</i> , pages 6215–6230. Association for Computational Linguistics.	2085
2029		2086
2030		2087
2031		2088
2032		2089
2033	Juraj Vladika and Florian Matthes. 2024a. <a href="#">Comparing knowledge sources for open-domain scientific claim verification</a> . In <i>Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024</i> , pages 2103–2114. Association for Computational Linguistics.	2090
2034		2091
2035		2092
2036		2093
2037		2094
2038		2095
2039		2096
2040		2097
2041	Juraj Vladika and Florian Matthes. 2024b. <a href="#">Improving health question answering with reliable and time-aware evidence retrieval</a> . In <i>Findings of the Association for Computational Linguistics: NAACL 2024, Mexico City, Mexico, June 16-21, 2024</i> , pages 4752–4763. Association for Computational Linguistics.	2098
2042		2099
2043		2100
2044		2101
2045		2102
2046		2103
2047	Juraj Vladika, Phillip Schneider, and Florian Matthes. 2024. <a href="#">Healthfc: Verifying health claims with evidence-based medical fact-checking</a> . In <i>Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy</i> , pages 8095–8107. ELRA and ICCL.	2104
2048		2105
2049		2106
2050		2107
2051		2108
2052		2109
2053		2110
2054	David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. <a href="#">Fact or fiction: Verifying scientific claims</a> . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020</i> , pages 7534–7550. Association for Computational Linguistics.	2111
2055		2112
2056		2113
2057		2114
2058		2115
2059		2116
2060		2117
2061		2118
2062	David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hannaneh Hajishirzi. 2022a. <a href="#">Scifact-open: Towards open-domain scientific claim verification</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022</i> , pages 4719–4734. Association for Computational Linguistics.	2119
2063		2120
2064		2121
2065		2122
2066		2123
2067		2124
2068		2125
2069		2126
2070	David Wadden, Kyle Lo, Lucy Lu Wang, Arman Cohan, Iz Beltagy, and Hannaneh Hajishirzi. 2022b. <a href="#">Multivers: Improving scientific claim verification with weak supervision and full-document context</a> . In <i>Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022</i> , pages 61–76. Association for Computational Linguistics.	2127
2071		2128
2072		2129
2073		2130
2074		2131
2075		2132
2076		2133
2077		2134
2078	David Wadden, Kejian Shi, Jacob Morrison, Aakanksha Naik, Shruti Singh, Nitzan Barzilay, Kyle Lo, Tom Hope, Luca Soldaini, Shannon Zejiang Shen, Doug Downey, Hannaneh Hajishirzi, and Arman Cohan. 2024. <a href="#">Sciriff: A resource to enhance language model instruction-following over scientific literature</a> . <i>CoRR</i> , abs/2406.07835.	2135
2079		2136
2080		2137
		2138
	Gengyu Wang, Kate Harwood, Lawrence Chillrud, Amith Ananthram, Melanie Subbiah, and Kathleen R. McKeown. 2023a. <a href="#">Check-covid: Fact-checking COVID-19 news claims with scientific evidence</a> . In <i>Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023</i> , pages 14114–14127. Association for Computational Linguistics.	2139
		2140
	Haiming Wang, Huajian Xin, Zhengying Liu, Wenda Li, Yinya Huang, Jianqiao Lu, Zhicheng Yang, Jing Tang, Jian Yin, Zhenguo Li, and Xiaodan Liang. 2024a. <a href="#">Proving theorems recursively</a> . <i>CoRR</i> , abs/2405.14414.	2141
		2142
	Haiming Wang, Huajian Xin, Chuanyang Zheng, Zhengying Liu, Qingxing Cao, Yinya Huang, Jing Xiong, Han Shi, Enze Xie, Jian Yin, Zhenguo Li, and Xiaodan Liang. 2024b. <a href="#">Lego-prover: Neural theorem proving with growing libraries</a> . In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	2143
		2144
	Haiming Wang, Ye Yuan, Zhengying Liu, Jianhao Shen, Yichun Yin, Jing Xiong, Enze Xie, Han Shi, Yujun Li, Lin Li, Jian Yin, Zhenguo Li, and Xiaodan Liang. 2023b. <a href="#">Dt-solver: Automated theorem proving with dynamic-tree sampling guided by proof-level value function</a> . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023</i> , pages 12632–12646. Association for Computational Linguistics.	2145
		2146
	Linghe Wang, Minhwa Lee, Ross Volkov, Luan Tuyen Chau, and Dongyeop Kang. 2025. <a href="#">Scholawrite: A dataset of end-to-end scholarly writing process</a> . <i>CoRR</i> , abs/2502.02904.	2147
		2148
	Pancheng Wang, Shasha Li, Kunyuan Pang, Liangliang He, Dong Li, Jintao Tang, and Ting Wang. 2022a. <a href="#">Multi-document scientific summarization from a knowledge graph-centric view</a> . In <i>Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022</i> , pages 6222–6233. International Committee on Computational Linguistics.	2149
		2150
	Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2024c. <a href="#">Scimon: Scientific inspiration machines optimized for novelty</a> . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024</i> , pages 279–299. Association for Computational Linguistics.	2151
		2152
	Wenxiao Wang, Lihui Gu, Liye Zhang, Yunxiang Luo, Yi Dai, Chen Shen, Liang Xie, Binbin Lin, Xiaofei He, and Jieping Ye. 2024d. <a href="#">Scipip: An llm-based scientific paper idea proposer</a> . <i>CoRR</i> , abs/2410.23166.	2153

2139	Yidong Wang, Qi Guo, Wenjin Yao, Hongbo Zhang,	Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren,	2195
2140	Xin Zhang, Zhen Wu, Meishan Zhang, Xinyu Dai,	Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li, and	2196
2141	Min Zhang, Qingsong Wen, Wei Ye, Shikun Zhang,	Xiaodan Liang. 2024. <a href="#">Deepseek-prover: Advancing</a>	2197
2142	and Yue Zhang. 2024e. <a href="#">Autosurvey: Large language</a>	<a href="#">theorem proving in llms through large-scale synthetic</a>	2198
2143	<a href="#">models can automatically write surveys.</a> <i>CoRR</i> ,	<a href="#">data.</a> <i>CoRR</i> , abs/2405.14333.	2199
2144	abs/2406.10252.		
2145	Yifan Wang, Yiping Song, Shuai Li, Chaoran Cheng,	Jing Xiong, Jianhao Shen, Ye Yuan, Haiming Wang,	2200
2146	Wei Ju, Ming Zhang, and Sheng Wang. 2022b.	Yichun Yin, Zhengying Liu, Lin Li, Zhijiang Guo,	2201
2147	<a href="#">Disencite: Graph-based disentangled representation</a>	Qingxing Cao, Yinya Huang, Chuanyang Zheng,	2202
2148	<a href="#">learning for context-specific citation generation.</a> In	Xiaodan Liang, Ming Zhang, and Qun Liu. 2023.	2203
2149	<i>Thirty-Sixth AAAI Conference on Artificial Intelli-</i>	<a href="#">TRIGO: benchmarking formal mathematical proof re-</a>	2204
2150	<i>gence, AAAI 2022, Thirty-Fourth Conference on In-</i>	<a href="#">duction for generative language models.</a> In <i>Proceed-</i>	2205
2151	<i>novative Applications of Artificial Intelligence, IAAI</i>	<i>ings of the 2023 Conference on Empirical Methods</i>	2206
2152	<i>2022, The Twelveth Symposium on Educational Ad-</i>	<i>in Natural Language Processing, EMNLP 2023, Sin-</i>	2207
2153	<i>vances in Artificial Intelligence, EAAI 2022 Virtual</i>	<i>gapore, December 6-10, 2023</i> , pages 11594–11632.	2208
2154	<i>Event, February 22 - March 1, 2022</i> , pages 11449–	Association for Computational Linguistics.	2209
2155	11458. AAAI Press.		
2156	Jane Webster and Richard T. Watson. 2002. Analyzing	Ziyang Xu. 2025. <a href="#">Patterns and purposes: A cross-</a>	2210
2157	the past to prepare for the future: Writing a literature	<a href="#">journal analysis of ai tool usage in academic writing.</a>	2211
2158	review. <i>MIS Q.</i> , 26(2).	<i>Preprint</i> , arXiv:2502.00632.	2212
2159	Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo	An Yang, Baosong Yang, Beichen Zhang, Binyuan	2213
2160	Zhang, Jindong Wang, Yue Zhang, and Linyi	Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayi-	2214
2161	Yang. 2024. <a href="#">Cycleresearcher: Improving auto-</a>	heng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian	2215
2162	<a href="#">mated research via automated review.</a> <i>CoRR</i> ,	Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Ji-	2216
2163	abs/2411.00816.	axi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and	2217
2164		22 others. 2024a. <a href="#">Qwen2.5 technical report.</a> <i>CoRR</i> ,	2218
2165		abs/2412.15115.	2219
2166			
2167	Nigel L. Williams, Stanislav Ivanov, and Dimitrios	Kaiyu Yang and Jia Deng. 2019. <a href="#">Learning to prove the-</a>	2220
2168	Buhalis. 2023. <a href="#">Algorithmic ghost in the research</a>	<a href="#">orems via interacting with proof assistants.</a> In <i>Pro-</i>	2221
2169	<a href="#">shell: Large language models and academic knowl-</a>	<i>ceedings of the 36th International Conference on Ma-</i>	2222
2170	<a href="#">edge creation in management research.</a> <i>CoRR</i> ,	<i>chine Learning, ICML 2019, 9-15 June 2019, Long</i>	2223
2171	abs/2303.07304.	<i>Beach, California, USA</i> , volume 97 of <i>Proceedings</i>	2224
2172		<i>of Machine Learning Research</i> , pages 6984–6994.	2225
2173	Jinxuan Wu, Wenhan Chao, Xian Zhou, and Zhunchen	PMLR.	2226
2174	Luo. 2023. <a href="#">Characterizing and verifying scientific</a>		
2175	<a href="#">claims: Qualitative causal structure is all you need.</a>	Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chala-	2227
2176	In <i>Proceedings of the 2023 Conference on Empirical</i>	mala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J.	2228
2177	<i>Methods in Natural Language Processing, EMNLP</i>	Prenger, and Animashree Anandkumar. 2023a. <a href="#">Le-</a>	2229
2178	<i>2023, Singapore, December 6-10, 2023</i> , pages 13428–	<a href="#">andojo: Theorem proving with retrieval-augmented</a>	2230
2179	13439. Association for Computational Linguistics.	<a href="#">language models.</a> In <i>Advances in Neural Information</i>	2231
2180		<i>Processing Systems 36: Annual Conference on Neu-</i>	2232
2181	Zijian Wu, Jiayu Wang, Dahua Lin, and Kai Chen. 2024.	<i>ral Information Processing Systems 2023, NeurIPS</i>	2233
2182	<a href="#">Lean-github: Compiling github LEAN repositories</a>	<i>2023, New Orleans, LA, USA, December 10 - 16,</i>	2234
2183	<a href="#">for a versatile LEAN prover.</a> <i>CoRR</i> , abs/2407.17227.	2023.	2235
2184			
2185	Amelie Wüthrl, Yarik Menchaca Resendiz, Lara Grim-	Zhishen Yang, Raj Dabre, Hideki Tanaka, and Naoaki	2236
2186	minger, and Roman Klinger. 2024a. <a href="#">What makes</a>	Okazaki. 2023b. <a href="#">Scicap+: A knowledge augmented</a>	2237
2187	<a href="#">medical claims (un)verifiable? analyzing entity and</a>	<a href="#">dataset to study the challenges of scientific figure cap-</a>	2238
2188	<a href="#">relation properties for fact verification.</a> In <i>Proceed-</i>	<a href="#">tioning.</a> In <i>Proceedings of the Workshop on Scientific</i>	2239
2189	<i>ings of the 18th Conference of the European Chap-</i>	<i>Document Understanding co-located with 37th AAAI</i>	2240
2190	<i>ter of the Association for Computational Linguistics,</i>	<i>Conference on Artificial Intelligence (AAAI 2023), Re-</i>	2241
2191	<i>EACL 2024 - Volume 1: Long Papers, St. Julian's,</i>	<i>mote, February 14, 2023</i> , volume 3656 of <i>CEUR</i>	2242
2192	<i>Malta, March 17-22, 2024</i> , pages 2046–2058. Asso-	<i>Workshop Proceedings.</i> CEUR-WS.org.	2243
2193	ciation for Computational Linguistics.		
2194			
2188	Amelie Wüthrl, Dustin Wright, Roman Klinger, and	Zonglin Yang, Xinya Du, Junxian Li, Jie Zheng, Sou-	2244
2189	Isabelle Augenstein. 2024b. <a href="#">Understanding fine-</a>	janya Poria, and Erik Cambria. 2024b. <a href="#">Large lan-</a>	2245
2190	<a href="#">grained distortions in reports of scientific findings.</a>	<a href="#">guage models for automated open-domain scientific</a>	2246
2191	In <i>Findings of the Association for Computational</i>	<a href="#">hypotheses discovery.</a> In <i>Findings of the Association</i>	2247
2192	<i>Linguistics, ACL 2024, Bangkok, Thailand and vir-</i>	<i>for Computational Linguistics, ACL 2024, Bangkok,</i>	2248
2193	<i>tual meeting, August 11-16, 2024</i> , pages 6175–6191.	<i>Thailand and virtual meeting, August 11-16, 2024,</i>	2249
2194	Association for Computational Linguistics.	pages 13545–13565. Association for Computational	2250
		Linguistics.	2251

2252	Zonglin Yang, Xinya Du, Rui Mao, Jinjie Ni, and Erik Cambria. 2023c. <a href="#">Logical reasoning over natural language as knowledge representation: A survey</a> . <i>CoRR</i> , abs/2303.12023.	2309
2253		2310
2254		2311
2255		2312
2256	Zonglin Yang, Wanhao Liu, Ben Gao, Tong Xie, Yuqiang Li, Wanli Ouyang, Soujanya Poria, Erik Cambria, and Dongzhan Zhou. 2024c. <a href="#">Moosechem: Large language models for rediscovering unseen chemistry scientific hypotheses</a> . <i>CoRR</i> , abs/2410.07076.	2313
2257		2314
2258		2315
2259		
2260		2316
2261		2317
2262	Michihiro Yasunaga, Jungo Kasai, Rui Zhang, Alexander R. Fabbri, Irene Li, Dan Friedman, and Dragomir R. Radev. 2019. <a href="#">Scisummnet: A large annotated corpus and content-impact models for scientific paper summarization with citation networks</a> . In <i>The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019</i> , pages 7386–7393. AAAI Press.	2318
2263		2319
2264		2320
2265		2321
2266		2322
2267		2323
2268		2324
2269		
2270		2325
2271		2326
2272		2327
2273		2328
2274	Geyan Ye, Xibao Cai, Houtim Lai, Xing Wang, Junhong Huang, Longyue Wang, Wei Liu, and Xiangxiang Zeng. 2024. <a href="#">Drugassist: A large language model for molecule optimization</a> . <i>CoRR</i> , abs/2401.10334.	
2275		2329
2276		2330
2277		2331
2278	Haofei Yu, Zhaochen Hong, Zirui Cheng, Kunlun Zhu, Keyang Xuan, Jinwei Yao, Tao Feng, and Jiaxuan You. 2024a. <a href="#">Researchtown: Simulator of human research community</a> . <i>CoRR</i> , abs/2412.17767.	2332
2279		2333
2280		2334
2281		2335
2282	Luyao Yu, Qi Zhang, Chongyang Shi, An Lao, and Liang Xiao. 2024b. <a href="#">Reinforced subject-aware graph neural network for related work generation</a> . In <i>Knowledge Science, Engineering and Management - 17th International Conference, KSEM 2024, Birmingham, UK, August 16-18, 2024, Proceedings, Part I</i> , volume 14884 of <i>Lecture Notes in Computer Science</i> , pages 201–213. Springer.	2336
2283		2337
2284		
2285		2338
2286		2339
2287		2340
2288		2341
2289		
2290	Mengxia Yu, Wenhao Yu, Lingbo Tong, and Meng Jiang. 2022. Scientific comparative argument generation.	2342
2291		2343
2292		2344
2293	Jiakang Yuan, Xiangchao Yan, Botian Shi, Tao Chen, Wanli Ouyang, Bo Zhang, Lei Bai, Yu Qiao, and Bowen Zhou. 2025. Dolphin: Closed-loop open-ended auto-research through thinking, practice, and feedback. <i>arXiv preprint arXiv:2501.03916</i> .	2345
2294		2346
2295		
2296		2347
2297	Weizhe Yuan and Pengfei Liu. 2022. <a href="#">Kid-review: Knowledge-guided scientific review generation with oracle pre-training</a> . In <i>Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022</i> , pages 11639–11647. AAAI Press.	2348
2298		2349
2299		2350
2300		2351
2301		2352
2302		
2303		2353
2304		2354
2305		2355
2306	Weizhe Yuan, Pengfei Liu, and Graham Neubig. 2022. <a href="#">Can we automate scientific reviewing?</a> <i>J. Artif. Intell. Res.</i> , 75:171–212.	2356
2307		2357
2308		2358
		2359
		2360
		2361
		2362
		2363
		2364
		2365
	Fengzhu Zeng and Wei Gao. 2023. <a href="#">Prompt to be consistent is better than self-consistent? few-shot and zero-shot fact verification with pre-trained language models</a> . In <i>Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023</i> , pages 4555–4569. Association for Computational Linguistics.	
	Qi Zeng, Mankeerat Sidhu, Ansel Blume, Hou Pong Chan, Lu Wang, and Heng Ji. 2024. Scientific opinion summarization: Paper meta-review generation dataset, methods, and evaluation. In <i>Artificial Intelligence for Research and Democracy: First International Workshop, AI4Research 2024, and 4th International Workshop, DemocrAI 2024, Held in Conjunction with IJCAI 2024, Jeju, South Korea, August 5, 2024, Proceedings</i> , page 20. Springer Nature.	
	Qi Zeng, Mankeerat Sidhu, Hou Pong Chan, Lu Wang, and Heng Ji. 2023. <a href="#">Meta-review generation with checklist-guided iterative introspection</a> . <i>CoRR</i> , abs/2305.14647.	
	Lei Zhang, Yuge Zhang, Kan Ren, Dongsheng Li, and Yuqing Yang. 2024a. <a href="#">Mlcopilot: Unleashing the power of large language models in solving machine learning tasks</a> . In <i>Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024</i> , pages 2931–2959. Association for Computational Linguistics.	
	Shujian Zhang, Chengyue Gong, Lemeng Wu, Xingchao Liu, and Mingyuan Zhou. 2023. <a href="#">Automl-gpt: Automatic machine learning with GPT</a> . <i>CoRR</i> , abs/2305.02499.	
	Xiaocheng Zhang, Xi Wang, Yifei Lu, Zhuangzhuang Ye, Jianing Wang, Mengjiao Bao, Peng Yan, and Xiaohong Su. 2024b. <a href="#">Augmenting the veracity and explanations of complex fact checking via iterative self-revision with llms</a> . <i>CoRR</i> , abs/2410.15135.	
	Xingjian Zhang, Yutong Xie, Jin Huang, Jinge Ma, Zhaoying Pan, Qijia Liu, Ziyang Xiong, Tolga Ergen, Dongsub Shim, Honglak Lee, and Qiaozhu Mei. 2024c. <a href="#">MASSW: A new dataset and benchmark tasks for ai-assisted scientific workflows</a> . <i>CoRR</i> , abs/2406.06357.	
	Xuan Zhang and Wei Gao. 2023. <a href="#">Towards llm-based fact verification on news claims with a hierarchical step-by-step prompting method</a> . In <i>Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, IJCNLP 2023 -Volume 1: Long Papers, Nusa Dua, Bali, November 1 - 4, 2023</i> , pages 996–1011. Association for Computational Linguistics.	
	Yu Zhang, Xiusi Chen, Bowen Jin, Sheng Wang, Shuiwang Ji, Wei Wang, and Jiawei Han. 2024d. <a href="#">A comprehensive survey of scientific large language models</a>	



and their applications in scientific discovery. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 8783–8817. Association for Computational Linguistics.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023a. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Xueliang Zhao, Wenda Li, and Lingpeng Kong. 2023b. *Decomposing the enigma: Subgoal-based demonstration learning for formal theorem proving*. *CoRR*, abs/2305.16366.

Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. 2022. *minif2f: a cross-system benchmark for formal olympiad-level mathematics*. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.

Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Lei Shen, Zihan Wang, Andi Wang, Yang Li, and 1 others. 2023a. Codegeex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5673–5684.

Yizhen Zheng, Huan Yee Koh, Jiaxin Ju, Anh T. N. Nguyen, Lauren T. May, Geoffrey I. Webb, and Shirui Pan. 2023b. *Large language models for scientific synthesis, inference and explanation*. *CoRR*, abs/2310.07984.

Yangqiaoyu Zhou, Haokun Liu, Tejes Srivastava, Hongyuan Mei, and Chenhao Tan. 2024. *Hypothesis generation with large language models*. *CoRR*, abs/2404.04326.

Kun Zhu, Xiaocheng Feng, Xiachong Feng, Yingsheng Wu, and Bing Qin. 2023. *Hierarchical catalogue generation for literature review: A benchmark*. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 6790–6804. Association for Computational Linguistics.

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

## A Further Discussion

**Open Question: What is the difference between AI for science and AI for research?** We posit that AI for research constitutes a subset of AI for science. While AI for research primarily focuses on supporting or automating the research process, it is not domain-specific and places greater emphasis

on methodological advancements. In contrast, AI for science extends beyond the research process to include result-oriented discovery processes within specific domains, such as materials design, drug discovery, biology, and the solution of partial differential equations (Zheng et al., 2023b; AI4Science and Quantum, 2023; Zhang et al., 2024d).

### Open Question: What is the difference between hypothesis generation and scientific discovery?

Hypothesis generation, which is primarily based on literature-based review (LBD) (Swanson, 1986; Sebastian et al., 2017), emphasizing the process by which researchers generate new concepts, solutions, or approaches through existing research and their own reasoning. Scientific discovery encompasses not only hypothesis generation, but also innovation in fields like molecular optimization and drug development (Ye et al., 2024; Liu et al., 2024b), driven by outcome-oriented results.

### Open Question: What is the difference between systematic literature review and related work generation?

Existing research frequently addresses the systematic literature survey, which constitutes a component of the knowledge synthesis process during hypothesis formulation, alongside the related work generation phase in manuscript writing (Luo et al., 2025). However, we argue that these two tasks are distinct in nature. The systematic literature survey primarily focuses on summarizing knowledge extracted from diverse scientific documents, thereby assisting researchers in acquiring an initial understanding of a specific field (Altamimi and Menai, 2022). In contrast, related work generation focuses on the writing process, emphasizing selection of pertinent literature and effective content structuring (Nishimura et al., 2024).

### Discussion: Potential links between artificial intelligence systems and human research practice

- In research paper recommendation, PaperWeaver (Lee et al., 2024) offers an interactive page that allows users to modify the topics they are interested in.
- In systematic literature review, Block and Kuckertz (2024) highlights the significant role of humans, including setting correct questions and individualized problem-solving and theorizing. Meanwhile, Hsu et al. (2024) emphasizes manual correction during the outline generation process.

- In hypothesis generation, AI engages more closely with human researchers, ranging from scenarios where humans provide the core ideas and AI contributes by iteratively refining them (Pu et al., 2024), to more collaborative settings where humans and AI engage in dialogue to facilitate new scientific discoveries (Ni et al., 2024; Liu et al., 2024b; Ye et al., 2024).
- In scientific claim verification, Altuncu et al. (2023); Das et al. (2023) highlight the critical role of experts in countering fake scientific news and advocate for the incorporation of expert opinions as a form of evidence.
- In theorem proving, Song et al. (2024) proposes leveraging LLMs as assistants to human researchers by generating suggested proof steps throughout the proving process.
- In experiment Validation, Ni et al. (2024) enhances the experimental setup through human-AI dialogue, whereas Li et al. (2024d) incorporates human input and real-time adjustments during the execution phase to optimize experimental design.
- In manuscript writing, Ifargan et al. (2024); Feng et al. (2024); Du et al. (2022a) require human intervention to suggest improvements to AI-generated paragraphs and enhance their quality through interactive methods.
- In peer review, Kumar et al. (2023); Darrin et al. (2024) advocate for assigning the responsibility of generating meta-reviews to human researchers. The role of AI is to assist by identifying conflicts among reviewers' opinions and supporting the chair in the scoring process, rather than independently assigning scores.

**Discussion: The involvement of AI in manuscript writing** The application of AI in manuscript writing has been accompanied by significant controversy. As LLMs demonstrated advanced capabilities, an increasing number of researchers began adopting these systems for scholarly composition (Liang et al., 2024b; Gao et al., 2023a). This trend raised concerns within the academic community (Salvagno et al., 2023), with scholars explicitly opposing the attribution of authorship to AI systems (Lee,

2023). Despite these reservations, the substantial time efficiencies offered by this technology led researchers to gradually accept AI-assisted writing practices (Gruda, 2024; Huang and Tan, 2023; Chen, 2023). This shift ultimately led to formal guidelines issued by leading academic journals (Ganjavi et al., 2024; Xu, 2025).

**Discussion: Some areas that have not been discussed** In addition to the eight areas discussed above, there are other lines of work that also aim to support scientific research, such as reading assistance (Kang et al., 2020; Head et al., 2021; Lo et al., 2023), which helps researchers read academic papers; literature processing<sup>1</sup>, which handles documents in various formats to provide effective data for subsequent tasks; as well as code and data generation (Bauer et al., 2024; Zheng et al., 2023a), which serve as a foundation for experimental validation. However, as our focus is on the core process of scientific research, we have chosen to omit these aspects from the main text.

## B Challenges

### B.1 Hypothesis Formulation

**Knowledge Synthesize** Existing paper recommendation tools predominantly rely on the meta-data of existing publications to suggest related articles, which often results in a lack of user-specific targeting and insufficiently detailed presentation that hampers comprehension. Leveraging LLMs can facilitate the construction of dynamic user profiles, enabling personalized literature recommendations and enhancing the richness of the contextual information provided for each recommended article, ultimately improving the user experience. In the process of generating systematic literature reviews, our practical experience reveals that the outline generation tools often produces redundant results with insufficient hierarchical structure. Moreover, the full-text generation process is prone to hallucinations—for instance, statements may not correspond to the cited articles—a pervasive issue in large language models (Huang et al., 2023; Bolaños et al., 2024; Susnjak et al., 2024). This problem can be ameliorated by enhancing the foundational model capabilities or by incorporating citation tracing.

**Hypothesis Generation** Most existing tools generate hypotheses by designing prompts or construct-

<sup>1</sup><https://sdproc.org>

ing systematic frameworks, which heavily rely on the capabilities of pre-trained models. However, these methods struggle to balance the novelty, feasibility, and validity of the hypotheses (Li et al., 2024c). Furthermore, our investigation reveals that many current approaches adopt novelty and feasibility as evaluation metrics; these metrics are either difficult to quantify or require manual scoring, which can introduce bias. To date, there is no unified benchmark to compare the various methods, and we believe that future research should prioritize establishing a unified metric that objectively reflects the strengths and weaknesses of different approaches.

## B.2 Hypothesis Validation

Most existing scientific claim verification tools are largely confined to specific domains, exhibiting poor generalizability, which limits their practical applicability (Vladika and Matthes, 2023). Theorem proving, the scarcity of relevant data adversely affects performance improvements through training, results across different proof assistants are not directly comparable, and the lack of standardized evaluation benchmarks presents numerous challenges. Moreover, current approaches remain predominantly in the research stage and lack practical tools that facilitate interaction with researchers (Li et al., 2024e). Experiment Validation, as automatically generated experiments often suffer from a lack of methodological rigor, practical feasibility, and alignment with the original research objectives (Lou et al., 2024). All these fields require rigorous logical reasoning, and I believe that the recent surge in advanced reasoning technologies could potentially address these issues.

## B.3 Manuscript Publication

Similar to systematic literature surveys, manuscript writing is also adversely affected by hallucination issues (Athaluri et al., 2023; Huang et al., 2023). Even when forced citation generation is employed, incorrect references may still be introduced (Aljamaan et al., 2024). Furthermore, the text generated by models requires meticulous examination by researchers to avoid ethical concerns, such as plagiarism risks (Salvagno et al., 2023). AI-generated manuscript reviews frequently provide vague suggestions and are susceptible to biases (Chamoun et al., 2024; Drori and Te’eni, 2024). Additionally, during meta-review generation, models can be misled by erroneous information arising from

the manuscript review process (Kumar et al., 2023). To address these issues, it may be necessary for the industry to establish appropriate regulations or to employ AI-based methods for detecting AI-generated papers and reviews (Lin et al., 2023a).

## C Ethical Considerations

AI has demonstrated significant potential in enhancing productivity by mitigating human limitations, thereby motivating increased investigation into its capacity to accelerate the research process (Messerli and Crockett, 2024). Nevertheless, its integration into scientific workflows introduces a range of ethical concerns (Fecher et al., 2023; Morris, 2023), including algorithmic biases, data privacy issues, risks of plagiarism, and the broader implications of AI-generated content for research communities. In this work, we examine these ethical challenges across the key stages of the research lifecycle: hypothesis formulation, validation, and publication.

During the hypothesis formulation stage, research paper recommendation systems and literature reviews are commonly employed; however, they often suffer from limitations that can lead to the formation of information bubbles and restrict exposure to diverse viewpoints. Furthermore, these systems tend to reinforce recognition disparities between prominent and lesser-known researchers and may inadvertently contribute to the dissemination of misinformation (Polonioli, 2021; Bolaños et al., 2024). To address these biases, recommendation algorithms can be enhanced by emphasizing content-based rather than author-based recommendations and by incorporating robust evaluation mechanisms to strengthen the credibility of suggested materials.

In contrast, AI-driven hypothesis generation presents more pronounced ethical challenges. First, the attribution of intellectual property rights and authorship for AI-generated hypotheses remains ambiguous (Majumder et al., 2024a). Additionally, the widespread generation of low-quality content poses a risk of diluting the integrity of the academic landscape (Hu et al., 2024a), while the potential misuse of such technologies for illicit purposes cannot be overlooked (Si et al., 2024). Addressing these concerns necessitates the development of robust accountability frameworks, the assignment of clear responsibility for AI-generated outputs to researchers, and the establishment of appropriate legal and regulatory mechanisms.



During the hypothesis validation phase, automated systems for scientific fact-checking remain underdeveloped. This limitation may be exploited by malicious actors to create advanced misinformation generators capable of circumventing existing fact-checking tools (Wadden et al., 2022b). Likewise, in the context of experimental validation, there is a risk of unethical or legally questionable experiments being designed (Eger et al., 2025). These concerns underscore the need for continued research into model safety.

During the manuscript publication stage, several challenges remain. Text generated by AI models may carry a risk of plagiarism (Salvagno et al., 2023; Gupta and Pruthi, 2025), while AI-assisted peer reviews often offer vague feedback and exhibit inherent biases (Schintler et al., 2023; Drori and Te’eni, 2024; Pataranutaporn et al., 2025). To address these issues, the development of robust detection methods is essential. However, current detection tools are still in the early stages of maturity (Gupta and Pruthi, 2025).

## D Ability Comparison

An effective survey should not only summarize existing methods within a field but also provide comparative analyses of different approaches. However, the domain of AI for Research remains in its early stages, with many areas lacking standardized benchmarks and even established evaluation metrics. To facilitate a clearer understanding of the distinctions among various methods, we draw on existing literature (Kang et al., 2023; Bolaños et al., 2024; Luo et al., 2025; Vladika and Matthes, 2023; Yang et al., 2023c; Li and Ouyang, 2022, 2024; Lin et al., 2023a) and adopt attribute graphs to compare representative approaches within each subfield, as illustrated in table §1 to table §8.

Method	Human-Computer Interaction	LLM	Required Information	Return Information	Relevance Source
ComLittee (Kang et al., 2023)	✓	-	Authorship Graphs	Meta data with relevant authors	R, Co, Ci
ArZiGo (Pinedo et al., 2024)	✓	-	User Interest	Meta data	R
PaperWeaver (Lee et al., 2024)	✓	✓	Collected Papers	Meta data with description	R
Kang et al. (2022)	-	-	Author’s social network relationships +Reference relationship	Meta data with relevant authors	R

Table 1: Research Paper Recommendation, we referred to Kang et al. (2023) for comparing different methods, where R represents Paper recommender score, Co represents Co-author relationship, and Ci represents Cited author relationship.

Method	Research Field	Across Stages	Human Interaction	Task	Input	Output	Evaluation Method
AutoSurvey (Wang et al., 2024e)	Any	✓	-	Outline Generation, +Full-text Generation	Title & Full Content	Literature Survey	LLM & Human
CHIME (Hsu et al., 2024)	Biomedicine	-	✓	Outline Generation	Title & Full Content	Hierarchical Outline	Automatic Metrics
Knowledge Navigator (Katz et al., 2024)	Any	-	-	Outline Generation	Title & Full Content	Hierarchical Outline	LLM & Human
Relatedly (Palani et al., 2023)	Any	-	-	Full-text Generation	Title & Related Work	Literature Survey	Human
STORM (Shao et al., 2024)	Any	-	-	Outline Generation, +Full-text Generation	Title & Full Content	Literature Survey	LLM & Automatic Metrics

Table 2: Scientific Literature Review, we referred to Bolaños et al. (2024) and made modifications, thereby comparing different methods.

Method	Research Field	Across Stages	Human Interaction	Multi-agent	Trained Model	Online RAG	Novelty	Feasibility	Validity
COI (Li et al., 2024a)	Any	✓	-	-	-	✓	✓	✓	✓
Learn2Gen (Li et al., 2024c)	Artification Intelligence	✓	-	-	✓	-	✓	✓	✓
MatPilot (Ni et al., 2024)	Materials Science	✓	✓	✓	-	-	✓	✓	-
SciAgents (Ghafarirollahi and Buehler, 2024)	Any	-	✓	✓	-	-	✓	✓	-
SciMON (Wang et al., 2024c)	Any	-	-	-	✓	-	✓	-	-

Table 3: Hypothesis Generation, we referred to Luo et al. (2025) and made modifications, thereby comparing different methods.

Method	Input	Document Retrieval	Human Interaction	Rationale Selection	Evidence Format	Output
MULTIVERS (Wadden et al., 2022b)	Claim & scientific abstract	Provided	-	Longformer	Document	Label & sentence-level rationales
SFAVEL (Bazaga et al., 2024)	Claim	Pre-trained Language Model	-	-	knowledge graph	Top-K Facts & Corresponding Relevance Scores
ProToCo (Zeng and Gao, 2023)	Claim-Evidence Pair	Provided	-	-	Sentence	Label
MAGIC (Kao and Yen, 2024)	Claim	Provided	-	Dense Passage Retriever	Sentence	Label
aedFaCT (Altuncu et al., 2023)	News Article	Google Search	✓	Human	Document	Evidence

Table 4: Scientific Claim Verification, we referred to Vladika and Matthes (2023) and made modifications, thereby comparing different methods.

Method	Generation Based	Stepwise	Heuristic Search	Informal or Formal	Human-authored Realistic Proof
IBR (Qu et al., 2022)	-	✓	✓	Informal	-
GPT-f (Polu and Sutskever, 2020)	✓	✓	-	Formal	✓
DT-Solver (Wang et al., 2023b)	✓	✓	✓	Formal	✓
POETRY (Wang et al., 2024a)	✓	-	-	Formal	✓

Table 5: Theorem proving, we referred to Yang et al. (2023c) and made modifications, thereby comparing different methods.

Method	Research Field	Across Stages	Human Interaction	Multi-agent	Task	Input	External tools
AutoML-GPT (Zhang et al., 2023)	Artification Intelligence	-	-	-	Automated Machine Learning	Task-oriented Prompts	-
Chemcrow (Bran et al., 2024)	Chemistry	-	✓	-	Chemical Task	Task Description	✓
DOLPHIN (Yuan et al., 2025)	Any	✓	-	✓	Automated Scientific Research	Idea	✓
MechAgents (Ni and Buehler, 2023)	Physics	-	-	✓	Mechanical Problem	-	-
Manning et al. (2024)	Social Science	✓	-	✓	Simulating Human	-	-

Table 6: Experiment Validation: we use attribute diagrams to compare different schemes, and the table design refers to Hypothesis Generation.

Method	Across Stages	Human Interaction	Task	Input	Evaluation Method
AI Scientist (Lu et al., 2024)	✓	-	Full-text Generation	Manuscript Template & Experimental Results & Hypothesis	LLM
data-to-paper (Ifargan et al., 2024)	✓	✓	Full-text Generation	Experimental Results & Hypothesis	-
ScholaCite (Martin-Boyle et al., 2024)	-	-	Related Work Generation	Title & Abstract & Citation	Citation Graph Metrics
SciLit (Gu and Hahnloser, 2023)	✓	-	Citation Generation	Keywords	Automatic Metrics
Gu and Hahnloser (2024)	-	-	Citation Generation	Citation Intent & Keywords	Human

Table 7: Manuscript Writing, we referred to Li and Ouyang (2022, 2024) and made modifications, thereby comparing different methods.

Method	Across Stages	Human Interaction	Paper Review	Meta Review	Multi-agent	Trained Model	Output
Gamma-Trans (Muangkammuen et al., 2022)	-	-	✓	-	-	✓	Peer-review Score
MARG (D’Arcy et al., 2024a)	-	-	✓	-	✓	-	Peer-review Comments
CycleResearcher (Weng et al., 2024)	✓	-	✓	-	-	✓	Peer-review Comments & Score
PeerArg (Sukpanichnant et al., 2024)	-	-	-	✓	-	-	Final Decision
GLIMPSE (Darrin et al., 2024)	-	✓	-	✓	-	-	Summary of Peer-review

Table 8: Peer Review, we referred to Lin et al. (2023a) and made modifications, thereby comparing different methods.

Task	Benchmark	Domain	Size	Input	Output	Metric
Hypothesis Formulation	SCHOLAT (Li et al., 2020)	Research Paper Recommendation	34,518	-	-	Recall, Precision, F1-score
	ACL selection network (Tao et al., 2020)	Research Paper Recommendation	18,718	Topics	Related Papers	Accuracy
	CiteSeer (Kang et al., 2021)	Research Paper Recommendation	1,100	Paper	Related Papers	Correlation Coefficient
	SciReviewGen (Kasanishi et al., 2023)	Systematic Literature Review	10,000+	Abstracts	literature review	ROUGE
	FacetSum (Meng et al., 2021)	Systematic Literature Review	60,024	Source Text+Facet	Summary of Facet	ROUGE
	BigSurvey (Liu et al., 2022)	Systematic Literature Review	7,000+	Abstracts	Survey Paragraph	ROUGE, F1-score
	SCHOLARQABENCH (Asai et al., 2024)	Systematic Literature Review	2,200	Question	Answer with Citations	Accuracy, Coverage, Citations
	HiCaD (Zhu et al., 2023)	Systematic Literature Review	7,600	Reference Papers	Catalogues	+ Relevance, Usefulness
	CLUSTREC-COVID (Katz et al., 2024)	Systematic Literature Review	2,284	Titles, Abstracts	Topic	Catalogue Edit Distance Similarity (CEDS)
	CHIME (Hou et al., 2024)	Systematic Literature Review	2,174	Topic	Hierarchies	+ Catalogue Quality Estimate (CQE)
	Tian et al. (2024)	Systematic Literature Review	700	Subject, Reference	Title, Content	Clusters per Topic
	MASSW (Zhang et al., 2024c)	Hypothesis Generation	152,000	Context of Literature	Hypothesis	F1-score
	IdeaBench (Guo et al., 2024)	Hypothesis Generation	2,374	Instruction, Background Information	Hypothesis	BLEU, ROUGE, BERTScore, + Cosine Similarity, BLEURT
	SCIMON (Wang et al., 2024c)	Hypothesis Generation	-	Background Context	Idea	Insight Score, BERTScore, Novelty, + LLM Similarity Rating, Feasibility
	MOOSEYang et al. (2024b)	Hypothesis Generation	50	Background, Inspiration	Hypothesis	ROUGE, BERTScore
	DISCOVERYBENCH (Majumder et al., 2024b)	Hypothesis Generation	1,167	Data	Discovery	+BARTScore, Novelty
	LiveIdeaBench (Ruan et al., 2024a)	Hypothesis Generation	-	Scientific Keywords	Idea	Validity, Novelty
	Kumar et al. (2024)	Hypothesis Generation	100	Paper without Future Work	Idea	+ Helpfulness
Hypothesis Validation	SciRIF (Wadden et al., 2024)	Scientific Claim Verification	137,000	Evidence, Task prompt	Structured Paragraph	Idea Alignment Score, Idea Distinctness Index
	SCIFACT (Wadden et al., 2020)	Scientific Claim Verification	1,409	Claim, Evidence	Rationale Sentences, Label	F1, BLEU
	SCIFACT-OPEN (Wadden et al., 2022a)	Scientific Claim Verification	279	Claim, Evidence	Rationale Sentences, Label	Precision, Recall, Micro-F1
	MISSCI (Glockner et al., 2024b)	Scientific Claim Verification	435	Claim, Premise, Context	Verification	Precision, Recall, Micro-F1
	FEVER (Thorne et al., 2018)	Scientific Claim Verification	185,445	Claim, Evidence	Label, Necessary Evidence	Micro F1-score, P@L Arg@1
	XClaimCheck (Kao and Yen, 2024)	Scientific Claim Verification	16,177	Claim, Evidence	Label, Argument	+ METEOR Score, BERTScore
	HEALTHYVER (Sarouti et al., 2021)	Scientific Claim Verification	14,330	Claim, Evidence	Label	+ NLI-A, NLI-S, Matches@1
	QuanTemp (V et al., 2024)	Scientific Claim Verification	15,514	Claim, Evidence	Label	F1-Score, Oracle Accuracy
	SCITAB (Lu et al., 2023)	Scientific Claim Verification	1,225	Claim, Evidence	Label	+ Accuracy, Recall
	Check-COVID (Wang et al., 2023a)	Scientific Claim Verification	1,504	Claim, Evidence	Label	Macro-F1, Accuracy
	HealthFC (Vladika et al., 2024)	Scientific Claim Verification	750	Claim, Evidence	Label	Macro Precision, Macro Recall
	FACTKG (Kim et al., 2023)	Scientific Claim Verification	108,000	Claim, Evidence	Label	+ Macro F1-score, Accuracy
	BEAR-FACT (Würl et al., 2024a)	Scientific Claim Verification	1,448	Claim, Evidence	Label	Weighted-F1 Score, Macro-F1, BLEU, + BERTScore, Cohen’s Kappa Score
	MINIFZ (Zheng et al., 2022)	Theorem Proving	488	Problem, Theorem	Proof	+ Human Evaluation
	FIMO (Liu et al., 2023a)	Theorem Proving	149	Problem, Theorem, statements	Proof	Macro-F1
	LeanDojo (Yang et al., 2023a)	Theorem Proving	98,734	Problem, Theorem	Proof	Accuracy, Precision, Recall, Macro-F1
	Lean-github (Wu et al., 2024)	Theorem Proving	28,597	Problem, Theorem	Proof	Precision, Recall, F1-Macro
	TRIGO-real (Xiong et al., 2023)	Theorem Proving	427	Problem, Theorem	Proof	Accuracy
Manuscript Publication	TRIGO-web (Xiong et al., 2023)	Theorem Proving	453	Problem, Theorem	Proof	F1-Score
	TRIGO-gen (Xiong et al., 2023)	Theorem Proving	-	Problem, Theorem	Proof	Pass Rate
	CoqGym (Yang and Deng, 2019)	Theorem Proving	71,000	Problem, Theorem	Proof	Pass Rate, Accuracy, EM@n
	MLAgentBench (Huang et al., 2024b)	Experiment Validation	13	-	-	Pass Rate, Accuracy, EM@n
	AAAR-1.0 (Lou et al., 2024)	Experiment Validation	-	Instance, Papers	Design, Explanation	Success Rate
	TASKBENCH (Shen et al., 2024)	Experiment Validation	17,331	-	-	Competence, Efficiency
	Spider2-V (Cao et al., 2024a)	Experiment Validation	494	Task	Experiment Execution	S-F1, S-Precision, S-Recall
	CORE-Bench (Siegel et al., 2024)	Experiment Validation	270	Task Requirements	Experiment Result	+ S-Match, ROUGE
	LAB-Bench (Laurent et al., 2024)	Experiment Validation	2400	Multiple-choice Question	Answer	ROUGE, t-F1, v-F1
	PaperBench (Sturace et al., 2025)	Experiment Validation	20	Paper, Additional Information	Code	+Normalized Edit Distance
	SUPER (Bogin et al., 2024)	Experiment Validation	801	Task Requirements	-	Success Rate
	ScienceAgentBench (Chen et al., 2025)	Experiment Validation	102	Task Instruction, Dataset Information	Program	Accuracy
	SciCap+ (Yang et al., 2023b)	Manuscript Writing	414,000	Figure, OCR tokens	Caption	Accuracy, Precision, Coverage
	AAN Corpus (Radev et al., 2013)	Manuscript Writing	-	-	-	Replication Score
	SciSummNet (Yasunaga et al., 2019)	Manuscript Writing	1,000	Paper, Citation Sentence	Summary	Accuracy, Landmark-Based Evaluation
	CiteBench (Funkquist et al., 2023)	Manuscript Writing	358,765	Cited Papers, Context	Citation Text	Valid Execution Rate, Success Rate, CodeBERTScore, API Cost
	ALCE (Gao et al., 2023b)	Manuscript Writing	3,000	Question	Answer with Citations	BLEU, ROUGE, METEOR
	GCite (Wang et al., 2022b)	Manuscript Writing	2,500	Citing/Cited Paper	Citation Text	+ CIDEr, SPICE
	ARXIVEDITS (Jiang et al., 2022b)	Manuscript Writing	1,000	Sentence Pairs	Sentence, Intent	ROUGE
	CASIMIR (Jourdan et al., 2024)	Manuscript Writing	15,646	Original Sentence	Revised Sentence	ROUGE, BERTScore
Manuscript Publication	ParaRev (Jourdan et al., 2025)	Manuscript Writing	48,203	Original Paragraph	Revised Paragraph	Recall, Precision
	SCHOLAWRITE (Wang et al., 2025)	Manuscript Writing	62,000	Before-text	Writing Intention, After-text	BLEU, ROUGE
	MReD (Shen et al., 2023)	Peer Review	7,089	Reviews	Meta-Review	Precision, Recall, F1-score
	ORSUM/Zeng et al., 2024)	Peer Review	15,062	Reviews	Meta-Review	ROUGE-L, BERTScore, FACTCC
	PeerRead v1 (Kang et al., 2018)	Peer Review	107,000	Reviews	Accept/Reject	+ SummaC, DiscoScore
	NLPeer (Dycke et al., 2023)	Peer Review	5,000	Reviews, Paper	Review Score, Connection, + Review Category	Accuracy
	AMPERE (Hua et al., 2019)	Peer Review	400	Review	Review with Type	MRSE, F1-macro
	MOPRD (Lin et al., 2023b)	Peer Review	6,578	Reviews, Paper	Editorial Decision, Review, + Meta-Review, Author Rebuttal	+ Precision, Recall
	ARIES (D’Arcy et al., 2024b)	Peer Review	1,720	Review Comment, Edits	Comment-Edit Pairs	Precision, Recall, F1-score
	ASAP-Review (Yuan et al., 2022)	Peer Review	-	Paper	Review	Aspect Coverage, Aspect Recall, + Semantic Equivalence
	ReviewMT (Tan et al., 2024)	Peer Review	26,841	Paper	Review Dialogue	+ Human: Recommendation Accuracy (RAcc), + Informativeness (Info), Aspect-level, + Constructiveness (Acon) and Summary accuracy
	ReAct (Choudhary et al., 2021)	Peer Review	6,250	Paper	Classification of Review	ROUGE, BLEU, METEOR
	PEERSUM (Li et al., 2023)	Peer Review	-	Review	Meta-Review	Accuracy
						ROUGE, BERTScore, UniEval, ACC

Table 9: An overview of benchmarks on AI for research. In the Input, Output, and Metric columns, the ‘+’ symbol indicates that the row is a continuation of the previous row.



Tool	Research Recommendation	Paper	Systematic Literature Review	Hypothesis Generation	Scientific Claim Verification	Theorem Proving	Experiment Verification	Manuscript Writing	Peer Review	Reading Assistance
Connected Paper	✓									
Inciteful	✓									
Litmaps	✓									
Pasa	✓									
Research Rabbit	✓									
Semantic Scholar	✓									✓
GenGO	✓									✓
Jenni AI	✓							✓		✓
Elicit	✓		✓							
Undermind	✓		✓							
OpenScholar	✓		✓							
ResearchBuddies	✓		✓							
Hyperwrite	✓		✓					✓		
Concensus	✓		✓		✓					
Iris.ai	✓		✓		✓					✓
MirrorThink	✓		✓		✓		✓			✓
SciSpace	✓		✓					✓	✓	✓
Ask YourPDF	✓		✓		✓			✓	✓	✓
Iflytek	✓		✓		✓	✓	✓	✓		✓
FutureHouse	✓		✓	✓			✓			
Enago Read	✓		✓	✓	✓	✓				✓
Aminer	✓		✓	✓	✓	✓	✓	✓		✓
OpenResearcher	✓		✓	✓		✓	✓	✓	✓	✓
ResearchFlow	✓		✓		✓	✓	✓	✓	✓	✓
You.com	✓		✓	✓	✓	✓	✓	✓	✓	✓
GPT Researcher			✓							
PICO Portal			✓							
SurveyX			✓							
Scinence42:Dora			✓					✓		
STORM			✓					✓		
ChatDOC			✓							✓
Scite			✓							✓
Silatus			✓							✓
Agent Laboratory			✓				✓	✓		
Sider			✓					✓		✓
Quillbot			✓					✓	✓	✓
Scholar AI			✓		✓		✓	✓	✓	✓
AI-Researcher			✓	✓			✓	✓	✓	
AI Scientist				✓			✓	✓	✓	
Isabelle						✓				
LeanCopilot						✓				
Llmstep						✓				
Proverbot9001						✓				
chatgpt_academic								✓		
gpt_academic								✓		
HeadlineAnalyzer								✓		
Langsmith Editor								✓		
Textero.ai								✓		
Wordvice.AI								✓		
Writesonic								✓		
Writefull								✓	✓	
Covidence									✓	
Penelope.ai									✓	
Byte-science										✓
Cool Papers										✓
Explainpaper										✓
Uni-finder										✓

Table 10: Tools for Research Paper Assistance

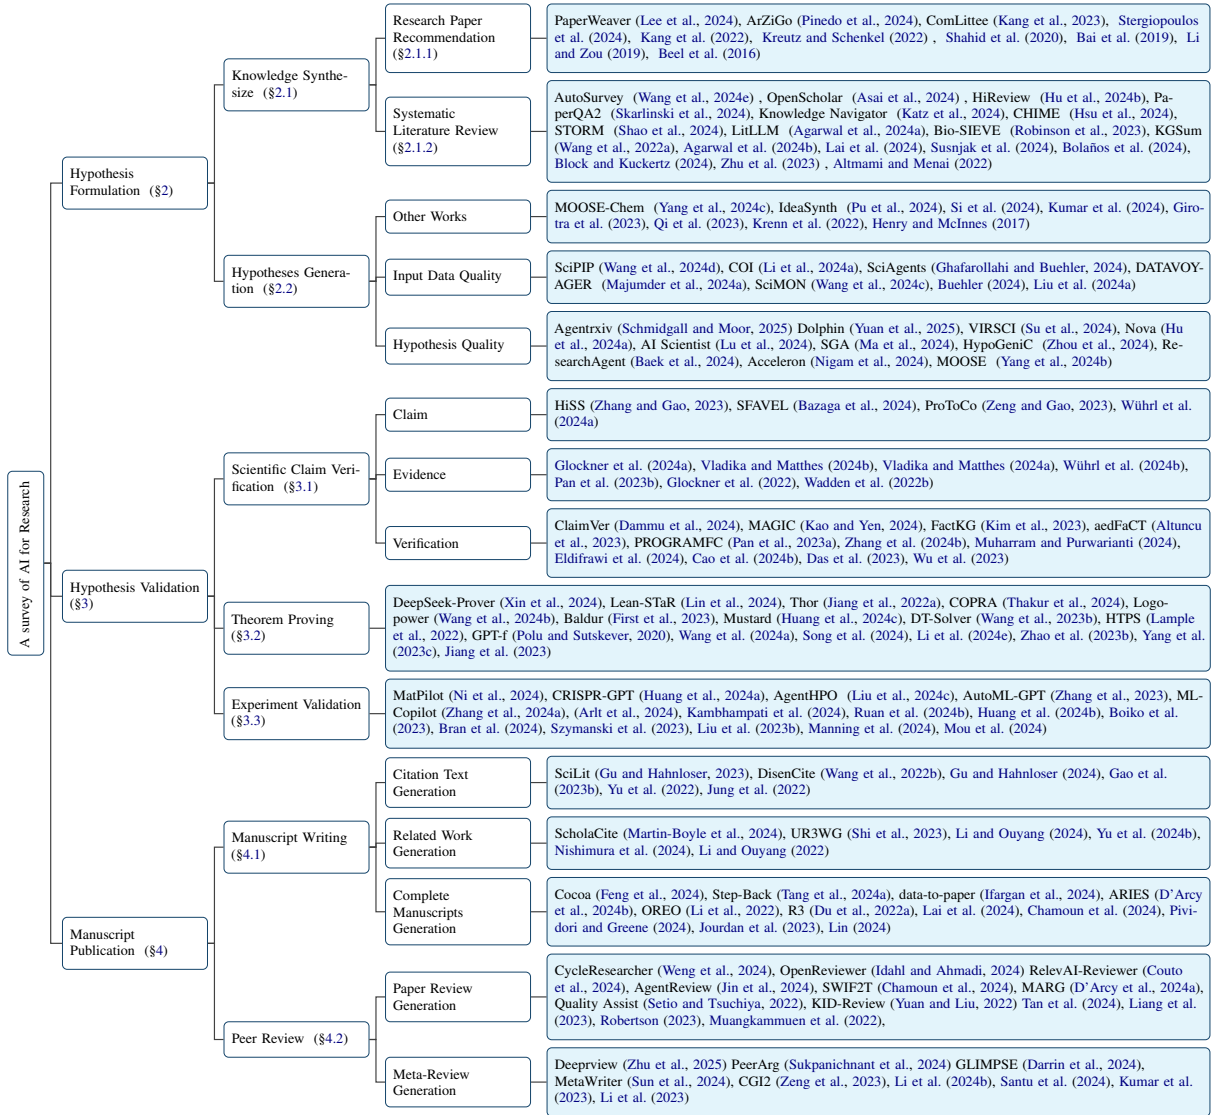


Figure 6: Taxonomy of Hypothesis Formulation, Hypothesis Validation and Manuscript Publication (Full Edition).