

Puzzles to Paradigms: A Survey of Human–Agent Collaboration in Scientific Innovation

Zhenni Liang¹ Haoxuan Zhang¹ Haihua Chen^{1,*}

Ting Xiao¹ Junhua Ding¹ Yang Zhang^{1,*}

¹University of North Texas, Denton, TX, USA

{zhenniliang, haoxuanzhang}@my.unt.edu

{haihua.chen, ting.xiao, junhua.ding, yang.zhang}@unt.edu

*Corresponding authors

Abstract

Agentic AI systems are increasingly embedded in scientific workflows, including literature synthesis, hypothesis generation, tool use, writing, and peer review. However, existing surveys and evaluations often emphasize task performance while underexamining how humans and agents coordinate under different epistemic conditions of scientific work. This survey introduces **CSCW×Kuhn**, a coordination-aware framework that connects computer-supported cooperative work dimensions with distinct scientific contexts. Using this framework, we synthesize recent human–agent systems and benchmarks to examine how coordination demands and failure risks may vary across contexts. Our synthesis identifies five recurring failure modes: procedural drift, grounding substitution, evaluative flattening, proxy optimization, and epistemic black-boxing, and argues for evaluations that move beyond task accuracy and toward grounding, provenance, disagreement preservation, and long-horizon scientific reliability. Our corpus and coding artifact are available at <https://github.com/jenlia024/CSCWxKuhn-artifact>.

1 Introduction

Agentic AI systems are increasingly entering scientific research workflows. These recent systems can retrieve and synthesize literature, generate hypotheses, operate scientific tools, write and revise manuscripts, simulate peer review, coordinate multiple agents, and execute data-analysis or experimental procedures (Wang et al., 2023; Chen et al., 2025a; Wei et al., 2025; Zheng et al., 2025; Wang et al., 2024a; Guo et al., 2024; Zhang et al., 2026b). Apart from the automation of scientific tasks, these emerging developments are also reshaping how scientific work is organized, monitored, and evaluated. Scientific innovation is rarely a single-step problem (Kuhn, 2012)—it requires researchers to frame problems, compare evidence, maintain provenance,

interpret uncertainty, revise assumptions, and decide whether a result is worth pursuing. As agentic systems become increasingly involved in these workflows, task completion is no longer a sufficient measure of success; effective and reliable human–agent coordination under various epistemic conditions matters too.

This question of human–agent coordination is central, as scientific workflows vary widely in uncertainty, evidence standards, and evaluation criteria. Some tasks resemble puzzle-solving within established methods (Kuhn, 2012), such as screening papers, executing analysis code, or following a protocol. Others involve evaluation of anomalies, where humans and agents must judge uncertain findings, contested claims, novelty, or reviewer disagreement. Yet others support paradigm-shift exploration, where the goal is to search for alternative concepts, distant analogies, or new research directions. These contexts impose different coordination demands. A system built for consistent, bounded execution may falter in contexts where evidence is contested; another built for exploration or discovery—generating plausible hypotheses—may be unreliable where provenance or validation paths remain opaque.

Computer-supported cooperative work (CSCW) offers a fruitful theoretical foundation for analyzing these situations. CSCW treats collaboration as a sociotechnical achievement that depends on articulation work, coordination mechanisms, shared understanding, and awareness of ongoing activity (Schmidt and Bannon, 1992; Schmidt and Simone, 1996; Clark and Brennan, 1991; Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002). These concerns are central to scientific workflows, where researchers must maintain alignment in goals, evidence, methods, assumptions, intermediate artifacts, and responsibility over multiple stages (Olson and Olson, 2000; Neang et al., 2021).

When agents enter such workflows, they do not

merely automate isolated tasks; they become part of distributed coordination arrangements. A CSCW perspective thus helps explain why task-level performance is insufficient: a system may produce a seemingly correct output while failing to preserve workflow state, evidential grounding, disagreement, provenance, or human accountability.

Existing surveys have made substantial progress in mapping AI systems for science, agentic architectures, and human–AI collaboration (Table 3). *AI4Science surveys* organize systems along research pipelines, including literature review, hypothesis generation, experimentation, analysis, and writing (Chen et al., 2025a; Zheng et al., 2025; Tie et al., 2026). *Agentic-system surveys* analyze planning, tool use, memory, retrieval, and multi-agent coordination (Wang et al., 2024a; Guo et al., 2024; Tran et al., 2025). *Human–agent collaboration surveys* examine role allocation, agency distribution, interaction patterns, and control with human involvement (Gomez et al., 2024; Vats et al., 2025; Zhang et al., 2025b; Zou et al., 2026). Although several surveys address the scientific epistemic context (Table 3), none explicitly link it to coordination demand and failure-mode analysis within a unified framework. In parallel, current evaluation methods often emphasize bounded task success, output correctness, or benchmark performance (Lou et al., 2025; Chen et al., 2025d; Majumder et al., 2025; Wijk et al., 2025), while giving less attention to workflow integrity, evidential grounding, disagreement preservation, provenance, and accountability (Xiong et al., 2025; Jin et al., 2024; Russo Latona et al., 2025).

In this survey, we address the analytical coverage gap by introducing **CSCW×Kuhn**, a framework for analyzing human–agent collaboration in scientific innovation. We use the term “human–agent collaboration” for systems in which agents contribute to scientific workflows (e.g., planning, judgment, execution, revision) and where human decisions remain relevant for final scientific output. Drawing on CSCW, we characterize coordination demand through articulation work, common ground, and awareness. Using Kuhn’s account of scientific change (Kuhn, 2012), we distinguish three epistemic modes: puzzle-solving in normal science, anomaly evaluation, and paradigm-shift exploration. Together, these perspectives ground our analysis.

CSCW×Kuhn proposes that coordination failure risks may emerge when the structure of human–

agent collaboration is misaligned with the epistemic demands of the scientific task. Specifically, our synthesis suggests that *procedural drift* is salient in bounded execution workflows, *grounding substitution* is critical in hypothesis generation and literature-based reasoning, and *evaluative flattening* becomes consequential when agents participate in peer review or novelty assessment. This framing helps explain why task-level success can still mask failures in scientific reliability. Figure 1 summarizes our framework. Our survey makes the following three contributions:

- **A coordination-aware taxonomy.** We introduce **CSCW×Kuhn**, a coordination-aware analytical framework and taxonomy that connects CSCW coordination dimensions with epistemic modes of scientific work.
- **A synthesis of recent systems.** We organize representative human–agent scientific systems across puzzle-solving in normal science, anomaly evaluation, and paradigm-shift exploration to illustrate how collaboration patterns and system risks vary by epistemic context.
- **An evaluation agenda.** We identify five recurring coordination failure modes: procedural drift, grounding substitution, evaluative flattening, proxy optimization, and epistemic black-boxing, and translate them into evaluation targets beyond task accuracy.

We use CSCW×Kuhn to synthesize recent systems and benchmarks over three scientific contexts. Section 2 defines the framework and formalizes coordination demand. Section 3 applies the framework to representative human–agent systems across puzzle-solving, anomaly evaluation, and paradigm-shift exploration. Section 4 analyzes current benchmarks and evaluation methods, showing where task-centric evaluation under-measures coordination quality. Section 5 discusses implications for system design, evaluation, and governance, and Section 6 concludes. This survey is not an exhaustive systematic review; it is intended as a focused, representative conceptual survey. Appendix C details the search, screening, and coding procedures.

2 A Coordination Framework for Human–Agent Scientific Collaboration

Human–agent collaboration in scientific innovation is more than assigning tasks to agents or evaluat-




Figure 1: CSCW×Kuhn framework for human-agent scientific collaboration. It maps CSCW coordination dimensions to epistemic modes of scientific innovation, and proposes that higher epistemic uncertainty is associated with greater coordination, grounding, transparency, and oversight demands.

ing final outputs. Scientific work unfolds through multi-step workflows (e.g., problem framing, literature review, hypothesis generation, experimentation, analysis, review). In these workflows, humans and agents must coordinate goals, assumptions, evidence, intermediate states, and responsibilities. The required coordination varies across scientific contexts. Some tasks involve stable goals and accepted criteria; others involve uncertain claims and contested evidence; yet others require exploring entirely new conceptual directions.

We present CSCW×Kuhn, a framework for analyzing these differences. Concepts from CSCW ground our framework in prior work on cooperative work and coordination. In CSCW, collaboration requires the organization of interdependent activities, the repair of workflow breakdowns, the maintenance of common understandings, and the visibility of ongoing work (Schmidt and Bannon, 1992; Schmidt and Simone, 1996; Clark and Brennan, 1991; Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002). This perspective is useful because agent capabilities are deeply embedded in interdependent scientific workflows. Agents may retrieve evidence, decompose tasks, execute tools, critique claims, or revise outputs, while humans oversee, interpret scientific meaning, and decide if the resulting work is credible. Moreover,

Kuhn’s account of scientific change helps distinguish routine puzzle-solving, anomaly evaluation, and paradigm-shift exploration (Kuhn, 2012). We combine CSCW and Kuhn to analyze how coordination demands and evaluation targets change in context, across human-agent scientific workflows.

2.1 Coordination Dimensions

We characterize coordination demand with three dimensions informed by CSCW:

- **Articulation work** (α_t): The work required to organize, assign, sequence, integrate, and repair interdependent tasks. This follows CSCW accounts of articulation work as the often-invisible organizing work that makes cooperative activity possible (Strauss, 1985; Schmidt and Bannon, 1992).
- **Common ground** (γ_t): The shared understanding of goals, assumptions, constraints, evidence, methods, and task state. This dimension draws on grounding theory, which emphasizes that collaboration requires participants to establish and maintain sufficient mutual understanding for the activity at hand (Clark and Brennan, 1991).
- **Awareness** (ω_t): The visibility of actions, intermediate outputs, decisions, errors, and consequences. This dimension follows CSCW work

on awareness in shared workspaces, where participants need to understand what others are doing, what has changed, and how their own actions should adapt (Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002).

At workflow step t , we define coordination demand as:

$$\phi_t = (\alpha_t, \gamma_t, \omega_t), \quad (1)$$

where α_t , γ_t , and ω_t correspond to articulation work, common ground, and awareness. This tuple is a qualitative, ordinal coding representation of what a human-agent system must support. Appendix A.2 provides further discussion of this framing.

2.2 Scientific Contexts

Coordination demands depend on the scientific context. We use Kuhn’s theory (Kuhn, 2012) as an analytical heuristic and distinguish three contexts:

- **Puzzle-solving in normal science:** Work conducted within accepted goals, methods, and evaluation criteria.
- **Anomaly evaluation:** Work that assesses unexpected findings, uncertain claims, or contested contributions.
- **Paradigm-shift exploration:** Exploratory work that searches for alternative concepts, cross-domain connections, or new research directions.

Let K denote the set of scientific contexts:

$$K = \left\{ \begin{array}{l} \text{Puzzle-solving in normal science,} \\ \text{Anomaly evaluation,} \\ \text{Paradigm-shift exploration} \end{array} \right\}, \quad (2)$$

where $k_t \in K$ denotes the dominant context at step t . These categories are not strict boundaries; a scientific workflow may move from routine execution to anomaly evaluation or exploratory reasoning. Appendix A.3 provides a more in-depth discussion of this framing.

2.3 Framework Formulation

We model human-agent scientific collaboration as an iterative workflow that transforms an initial research intent into evidence-backed outputs. In systems with substantial agent-agent coordination, we treat these as an architectural feature of the broader

sociotechnical workflow when humans retain control over objectives, evaluation, validation, or outputs, even if execution is autonomous.

Let $S = \{s_1, s_2, \dots, s_n\}$ denote the set of workflow stages (e.g., problem framing, experimentation, review). The coordination demand at step t depends on both the scientific context and the workflow stage:

$$\phi_t = f(k_t, s_t), \quad (3)$$

where $k_t \in K$ and $s_t \in S$.

This formulation captures the main idea of CSCW×Kuhn—different scientific contexts require different forms of human-agent coordination. In normal-science puzzle-solving, coordination mainly concerns reliable execution and tracking of workflow. In anomaly evaluation, it shifts toward evidential comparison, disagreement preservation, and evaluative alignment. In paradigm-shift exploration, coordination becomes more demanding because goals, assumptions, and validation criteria may be unstable.

Under this framework, workflow stage and epistemic mode are separate, distinct dimensions. The stage indicates where in the research workflow an activity occurs; epistemic mode characterizes the scientific conditions under which it occurs. We illustrate this independence with the following two examples.

Within puzzle-solving, literature review and protocol execution occupy different workflow stages but share a common epistemic mode, with stable goals and evaluation criteria. Their differing coordination demands thus stem from workflow stage, not from a shift in epistemic conditions. Conversely, literature-related work can occur under different epistemic modes at the same workflow stage. For instance, under puzzle-solving, routine evidence retrieval demands mainly reliable search and coverage, whereas under anomaly evaluation, conflicting literature assessment demands mainly evidential comparison and disagreement preservation.

Additionally, this framework clarifies why task-level evaluation is insufficient: a system may complete a benchmark task while failing to preserve workflow integrity or human accountability. Thus, we use CSCW×Kuhn as a coding framework that links scientific context, coordination demand, failure mode, and evaluation target.

We operationalize CSCW×Kuhn as a qualitative coding scheme. For each system or workflow segment, we code for the dominant epistemic context,

the level of coordination support provided by the system along articulation work, common ground, and awareness, and the corresponding failure risks and evaluation targets. Appendix B provides the full rubric and decision rules.

3 Collaboration Patterns Across Scientific Contexts

Corpus and Coding. To apply CSCW×Kuhn, recent human-agent systems and benchmarks were coded along four dimensions: epistemic mode, system cluster, coordination demand, and potential failure risk. The corpus was constructed through a focused search—rather than an exhaustive systematic review—of recent work on scientific agents, AI for science, human-agent collaboration, tool-augmented workflows, peer review, and multi-agent scientific discovery. We use this coding to organize representative systems across puzzle-solving in normal science, anomaly evaluation, and paradigm-shift exploration. Full details on search sources, inclusion/exclusion criteria, screening, and coding decisions are provided in Appendix C. Table 1 summarizes, for each scientific context, the dominant coordination pattern that CSCW×Kuhn identifies across these systems. The system-level coding is detailed in Appendix D (Table 2).

3.1 Puzzle-Solving in Normal Science

Stable task frames. Puzzle-solving in normal science refers to work conducted within accepted goals, methods, and evaluation criteria (Kuhn, 2012). In this context, the main coordination challenge is reliable execution. Articulation work concerns task decomposition, tool use, protocol following, and workflow integration. Common ground depends on shared assumptions about domain constraints, data formats, methods, and success criteria. Awareness requires visible intermediate outputs, status updates, errors, and checkpoints.

Observed collaboration pattern. Systems in this context mainly support bounded scientific workflows, including literature retrieval, evidence screening, molecular design, protocol execution, and shared planning (Wang et al., 2025b; Ye et al., 2025; Kon et al., 2025; Feng et al., 2026). In evidence-oriented workflows, citation interpretation can depend on broader document context and source-level information surrounding citations (Zhang et al., 2021, 2022). Citation-function

modeling has also been advanced through hybrid data augmentation and broader deep-learning approaches (Zhang et al., 2023a, 2025c), while recent survey work examines the growing intersection between LLMs and citation analysis (Zhang et al., 2023b). Agents typically assist by decomposing tasks, expanding search, generating candidates, executing procedures, or maintaining workflow artifacts. Human involvement remains critical for checking constraints, validating intermediate results, and correcting deviations.

Evaluation gap. Current evaluations often emphasize task completion, output correctness, or benchmark performance. These measures are useful but incomplete. They may show that an agent completes a predefined task, while missing whether the system preserved the intended workflow, handled intermediate states correctly, and supported human oversight in long-horizon processes. Recent work on long-horizon agent memory likewise shows that downstream task success may obscure whether an agent retrieves the correct historical evidence, particularly under constrained context budgets (Li et al., 2026). These limitations create a risk of *procedural drift*, in which an agent’s actions deviate from the intended workflow, protocol, constraints, or task state, even when the final output appears successful.

Worked example: experiment design. AAAR-1.0 (Lou et al., 2025) evaluates scientific tasks with well-defined input–output specifications (see Table 5 for its coverage mapping). In its experiment design task, an LLM receives pre-experiment paper context and generates an experiment plan with motivations. It is reported that while some generated experiments can be innovative and diverse, many others are trivial, infeasible, or divergent from the original objective. The findings concern output quality rather than the intermediate generation process. As an illustrative scenario grounded in this benchmark design, a system could produce an experiment plan and receive a favorable score for its output. Yet violations of intermediate constraints or workflow deviations could remain unobserved.

3.2 Anomaly Evaluation

Contested scientific judgment. Anomaly evaluation refers to the assessment of uncertain findings, contested claims, or contributions whose significance is not yet settled. This includes novelty assessment, peer review, rebuttal analysis, critique

Scientific Context	Representative System Clusters	Coordination Focus	Main Failure Risks	Evaluation Implications
Puzzle-solving in normal science	Literature/evidence workflows; design and optimization; protocol execution; shared planning	task decomposition, protocol execution, workflow tracking; domain constraints, data formats, methods, success criteria; intermediate outputs, checkpoints, errors, deviations	Procedural drift; grounding substitution; hidden execution errors; weak long-horizon oversight	Audit traces; intermediate state checks; tool-use inspection
Anomaly evaluation	Hypothesis and novelty evaluation; LLM-as-reviewer; rebuttal and meta-review; institutional review	review criteria, novelty checks, rebuttal and decision workflows; paper–literature–benchmark–venue alignment; uncertainty, disagreement, unsupported claims, confidence signals	Evaluative flattening; shallow fluency; majority-opinion bias; weak adjudication of contested claims	Claim–evidence alignment checks; disagreement and judgment preservation
Paradigm-shift exploration	Cross-domain ideation; falsification and feedback loops; recursive AI scientist workflows; verification/search; infrastructure	evolving goals, search strategies, validation plans; problem–concept–domain–evidence alignment; idea provenance, retrieved evidence, rankings, experimental outcomes	Grounding substitution; proxy optimization; epistemic black-boxing; unsupported or non-testable ideas	Assumption tracing; provenance checks; uncertainty reports; expert critiques; downstream validation

Table 1: Summary of coordination patterns with corresponding evaluation implications across scientific contexts. We note that workflow stages are not exclusive to a single epistemic mode. The same stage may recur across scientific contexts with different coordination demands, and a single mode may span multiple workflow stages.

generation, and meta-review. In this context, the coordination challenge shifts from reliable execution to evaluative alignment. Here, humans and agents must compare evidence, surface disagreement, and make judgment criteria explicit.

Observed collaboration pattern. Systems in this context often convert scientific judgment into structured workflows (Xiong et al., 2025; Lin et al., 2025; Zhu et al., 2025; Chang et al., 2025; Wu et al., 2026). They retrieve comparison evidence, decompose review criteria, identify unsupported claims, expose reviewer disagreement, or support rebuttal reasoning. We therefore view these agents as most appropriate for structuring judgment rather than replacing human evaluation.

Coordination demands. Articulation work involves decomposing scientific merit into criteria such as novelty, soundness, relevance, actionability, and evidential support. Common ground requires alignment among the paper, prior literature, benchmarks, venue expectations, reviewer concerns, and rebuttal evidence. Awareness necessitates visibility into uncertainty, disagreements, unsupported claims, confidence signals, and revisions before humans make decisions.

Evaluation gap. Existing evaluations often rely on prediction accuracy, agreement with human la-

bels, review-quality rankings, groundedness, specificity, or actionability. These metrics do not fully capture whether a system helps humans preserve minority views, distinguish substantive novelty from superficial difference, or adjudicate evidential conflict. This limitation is consequential because AI-assisted review systems may reshape reviewer confidence, rebuttal behavior, and decision outcomes (Russo Latona et al., 2025; Liang et al., 2024; Jin et al., 2024; Zhao et al., 2026). We characterize this risk as *evaluative flattening*, in which contested scientific judgment is reduced to simplified scores, rankings, or majority opinions that obscure multidimensional criteria and disagreement.

Worked example: novelty assessment.

SchNovel (Lin et al., 2025) evaluates whether LLMs can identify the more novel paper in a pair using titles, abstracts, and optional metadata, with performance measured by accuracy (Table 5). The findings show that novelty judgments are sensitive to factors unrelated to scholarly contribution. Models exhibit position bias, and adding institutional affiliation can produce a preference for papers associated with a top research university. This illustrates how a structured novelty judgment can be affected by contextual cues, even when the evaluation reduces the decision to a novelty label.

3.3 Paradigm-Shift Exploration

Unstable goals and criteria. Paradigm-shift exploration refers to upstream work that expands the conceptual search space, connects distant bodies of knowledge, and tests whether speculative directions deserve further investigation (Kuhn, 2012). Unlike normal-science puzzle-solving, this context does not assume that goals, comparison spaces, or validation criteria are stable already.

Observed collaboration pattern. Systems in this context support cross-domain ideation, falsification loops, recursive research workflows, verification search, and scientific infrastructure (Wang et al., 2024b; Baek et al., 2025; Radensky et al., 2026; Gottweis et al., 2026; Ghareeb et al., 2026; Zhang et al., 2025a; Lu et al., 2026; Romera-Paredes et al., 2024; Shojaee et al., 2025). In our framework, these systems are most credible when speculative generation is paired with evidence retrieval, critique, executable checks, provenance tracking, or expert revision.

Coordination demands. Within our framework, paradigm-shift exploration can impose particularly strong coordination demands among the tasks analyzed. Articulation work must stabilize evolving goals, search strategies, and validation plans. Common ground is fragile, as agents may introduce unfamiliar concepts, distant analogies, or weakly related literature. Awareness is essential because researchers must inspect the provenance of the idea, retrieved evidence, rankings, experimental results, and verification signals before deciding whether a direction is scientifically meaningful.

Evaluation gap. Current evaluation remains limited for paradigm-shift exploration. Criteria such as novelty, plausibility, or benchmark success do not establish whether a system supports durable conceptual reframing. The relevant question is not necessarily whether an agent can generate surprising ideas, but whether those ideas remain grounded, testable, interpretable, and useful under scrutiny from experts and follow-up validation.

Worked example: scientific ideation. Si et al. (2025) compare LLM-generated research ideas with ideas from expert NLP researchers using blinded expert evaluation. LLM-generated ideas were rated significantly more novel than human ideas, but slightly weaker on feasibility. The authors identify failures in LLM self-evaluation along

with limited diversity in generation. They further note that idea-level judgments alone do not establish whether these advantages translate into meaningful research outcomes, and propose executing the ideas as full projects for end-to-end evaluation. This illustrates how favorable novelty judgments can coexist with unresolved questions about feasibility and downstream scientific validity (Table 5).

Recurring failure modes. Over the three scientific contexts, we identify patterns of five coordination failure modes. The term *failure risk* is used to denote the potential for a given failure mode to arise in a system or context. *Procedural drift* occurs when collaboration deviates from the intended workflow, protocol, constraint set, or task state. *Grounding substitution* occurs when plausibility, fluency, or apparent novelty replaces evidential support or domain validity. *Evaluative flattening* occurs when contested scientific judgment is reduced to simplified scores, rankings, or majority opinions. *Proxy optimization* occurs when systems optimize benchmark metrics while losing alignment with scientific goals. *Epistemic black-boxing* occurs when provenance, assumptions, rationale, intermediate states, tool use, or responsibility become difficult to inspect or contest.

4 Benchmarks and Evaluation Methods

Current benchmarks provide useful measures of task performance but only partially evaluate human-agent collaboration in scientific innovation. As summarized in Table 5, most benchmarks focus on task- or output-level metrics (e.g., retrieval quality, answer correctness, hypothesis plausibility, review quality, code execution, end-to-end task completion). These metrics improve comparability but often under-measure whether humans and agents preserve workflow integrity, maintain evidential grounding, surface disagreement, and support accountable scientific judgment.

4.1 Limits of Task-Centric Evaluation

Workflow integrity. Many benchmarks decompose scientific work into standalone tasks with clear input-output formats. This design supports reproducible evaluation but can obscure whether an agent follows the intended workflow across intermediate steps. As illustrated in Section 3.1, such benchmarks are suitable for measuring local task success but are less able to detect *procedural drift*, where an agent deviates from protocols or con-

straints. Considering this, evaluation should include trace audits, intermediate state checks, tool-use inspection, and procedure-fidelity measures in addition to final accuracy.

Evidential grounding. Output-level verification does not necessarily establish scientific grounding. Generated hypotheses or reviews may appear plausible while remaining weakly supported or insufficiently tied to the underlying scientific claims (Xiong et al., 2025; Ou et al., 2025). CLAIM-CHECK (Ou et al., 2025) further shows that LLMs struggle to associate reviewer-identified weaknesses with the specific claims they dispute, and to generate specific, claim-grounded critiques from scratch. As an illustrative scenario, a generated claim may seem well supported due to relevant material being retrieved; however, the cited evidence may not directly support the claim or cannot be traced. This reflects the risk of *grounding substitution*, where plausibility, fluency, or apparent novelty replaces evidence. Evaluation should test whether claims are supported by relevant sources, whether evidence is traceable, and whether generated conclusions remain consistent with the retrieved literature and data.

Scientific judgment. Some benchmarks isolate criteria that are easier to score than scientific merit as a whole. As discussed in Section 3.2, SchNovel (Lin et al., 2025) advances the evaluation of scholarly novelty; however, novelty is one dimension of scientific judgment, not a proxy for it in full. In anomaly evaluation, this creates the risk of *evaluative flattening*, where contested judgments are reduced to a single score. Evaluation should instead preserve tradeoffs among multiple dimensions of scientific value: novelty, validity, feasibility, significance, evidence quality, and disagreement.

Exploratory validity. Paradigm-shift exploration is especially difficult to evaluate because goals, assumptions, and validation criteria are often unstable. As illustrated in Section 3.3, favorable novelty judgments do not necessarily establish feasibility or downstream scientific validity. Existing ideation benchmarks assess novelty, feasibility, diversity, and quality of ideas, yet generated ideas may require expert assessment and empirical validation before becoming actionable (Si et al., 2025; Ruan et al., 2026). Evaluation should therefore include assumption tracing, provenance checks, uncertainty reports, critiques from experts, and vali-

dation from long-horizon follow-ups.

4.2 Toward Coordination-Aware Evaluation

Evaluation targets. An evaluation that accounts for coordination should complement task metrics with four additional targets. First, it should assess *process fidelity*: whether agents preserve protocols, constraints, intermediate states, and intended workflows. Second, it should assess *claim–evidence alignment*: whether generated claims are supported by relevant, traceable, and domain-appropriate evidence. Third, it should assess *judgment preservation*: whether systems maintain multiple evaluation criteria and make disagreement visible. Fourth, it should assess *epistemic transparency*, addressing the risk of epistemic black-boxing, by testing whether users can inspect assumptions, provenance, uncertainty, tool use, and responsibility (Zhang et al., 2026a).

Human oversight. Human oversight should vary with epistemic context. In puzzle-solving settings, humans should primarily engage in verification, such as checking workflow adherence, tool use, constraint satisfaction, and intermediate outputs. In anomaly evaluation, humans should focus on overseeing comparisons of evidence, reviewer disagreement, minority perspectives, and distinctions between substantive novelty and superficial difference. In paradigm-shift exploration, oversight should involve examining how candidate ideas were generated, what assumptions shaped these ideas, what evidence supports them, and what types of follow-up would make them testable.

Implications for benchmark design. Future benchmarks should move beyond final-output scoring to evaluate collaboration as a process. This requires in-depth traces, interpretable intermediate artifacts, explicit evidence links, human checkpoints for intervention, and longitudinal validation. Such evaluation does not replace task accuracy, but instead situates accuracy within the broader coordination demands of scientific work. Under CSCW×Kuhn, benchmark completion is not sufficient evidence for success; the operative criterion is whether the system can support reliable coordination under the epistemic conditions of the scientific workflow.

5 Discussion

Implications for evaluation. CSCW×Kuhn re-frames evaluation from a question of task com-

pletion by agents to whether a collaborative system—consisting of humans and agents—supports the coordination demands of scientific work. We emphasize this distinction because the same benchmark score may reflect different levels of scientific reliability across various epistemic contexts. Evaluation should not stop at output correctness in puzzle-solving settings; it should also test whether the system preserves procedure fidelity, constraint satisfaction, tool-use validity, and consistency in intermediate states. With anomaly evaluation, the issue goes beyond whether an agent predicts a human label or generates a plausible review; evaluations in this context should assess whether agents help preserve disagreement, compare alternatives, expose evidential gaps, and distinguish substantive novelty from superficial difference. Evaluation requires even greater caution in paradigm-shift exploration. Novelty and plausibility are insufficient unless the system also supports provenance, feasibility assessment, falsifiability, and follow-up validation. Thus, coordination-aware evaluation complements task metrics by asking what kinds of scientific accountability remain visible to humans.

Implications for system design. This perspective cautions against interpreting current benchmark progress as support for fully autonomous scientific discovery, and suggests limits to strong claims about end-to-end automation. For instance, ScienceAgentBench (Chen et al., 2025d) evaluates bounded workflow tasks and frames agents as assisting scientists with data processing and visualization, not substitutes for scientific judgment. In a related vein, benchmarks for hypothesis generation and novelty assessment show that plausible or original outputs still require grounding, expert interpretation, and downstream validation (Si et al., 2025; Xiong et al., 2025). System design should therefore prioritize modular assistance, explicit human checkpoints, interpretable intermediate artifacts, and domain-specific validation interfaces. Agent outputs should be treated as candidate analyses, hypotheses, reviews, or drafts rather than finalized scientific conclusions. This is crucial in long-horizon workflows where small errors can accumulate and amplify into misleading scientific claims.

Implications for peer review and institutional governance. Human–agent collaboration also raises institutional questions when agents participate in evaluative settings such as peer review. Review and meta-review help determine which work

receives credibility, attention, and publication opportunities. AgentReview (Jin et al., 2024) shows that review outcomes are shaped by reviewer heterogeneity, latent variables, and privacy constraints, and that biased reviewers, in particular, can materially affect decisions. Related evidence suggests that AI-assisted reviewing can influence review scores and acceptance outcomes, notably for borderline submissions (Russo Latona et al., 2025). Findings imply that the main governance issue is under what disclosure, monitoring, audit, and accountability regimes AI should be used. Without such safeguards, agentic judgments can shape scientific decisions without a clear record of how or by whom. Institutions therefore need policies that specify when AI assistance is permitted, how its use should be disclosed, how review evidence should be preserved, and how responsibility is assigned (Zhang et al., 2026a).

Open challenges. Several challenges remain. First, coordination quality is difficult to benchmark: it is processual and often visible only through intermediate traces, so future benchmarks need richer logs, evidence links, and human-in-the-loop protocols. Second, long-horizon scientific value is hard to measure. Ideas that appear novel at the time of evaluation may later prove to be unproductive; thus, longitudinal studies are needed. Third, CSCW×Kuhn should be empirically validated, testing whether researchers can reliably annotate its dimensions across real workflows. Finally, governance should be integrated into system development rather than treated as post-hoc.

6 Conclusion

This survey frames Human–Agent Collaboration in Scientific Innovation as an epistemically conditioned coordination process. We introduce CSCW×Kuhn to connect scientific context, coordination demand, failure mode, and evaluation target across puzzle-solving, anomaly evaluation, and paradigm-shift exploration. Our synthesis suggests that current evaluations often under-measure workflow integrity, evidential grounding, disagreement preservation, transparency, and human oversight. We argue that assessing scientific agents requires looking past task performance to whether agents support accountable coordination across various epistemic conditions.

Limitations

This survey has several limitations. First, the literature on agentic scientific systems is expanding rapidly, so our corpus cannot exhaustively cover all emerging systems, benchmarks, and deployment practices. Second, our coding relies on information reported in published papers; when system designs or evaluation procedures are underspecified, coordination demands and failure risks are interpreted from available evidence. Third, CSCW×Kuhn codes the dominant epistemic condition of each system or workflow segment, while real scientific workflows may transition across puzzle-solving, anomaly evaluation, and paradigm-shift exploration. Future work can extend CSCW×Kuhn through larger corpora, empirical annotation studies, and deployment-based validation.

Ethics Statement

This paper is a survey of human–agent collaboration in scientific innovation and does not introduce new human–subject experiments or collect personal or sensitive data. We release a bibliographic corpus and analytic coding artifact derived from publicly available research papers and benchmarks. We acknowledge that human–agent scientific systems may raise ethical concerns related to accountability, transparency, overreliance on automated judgments, and the potential misuse of agentic systems in scientific evaluation. Accordingly, we emphasize that such systems should be designed and evaluated with human oversight, provenance tracking, uncertainty disclosure, and clear responsibility assignment.

Information About Use of AI Assistants

AI assistants were used only for polishing language, correcting grammar, and improving the clarity and readability of the manuscript. They were not used to generate original scientific claims, conduct literature screening, make coding decisions, perform analysis, or produce conclusions without author verification. All content, references, interpretations, and final writing decisions were reviewed and approved by the authors, who take full responsibility for the manuscript.

Acknowledgments

We would like to thank the anonymous reviewers and meta-reviewer for their constructive feedback,

which helped us improve our presentation and clarity of this survey.

References

- Samuel Alber, Bowen Chen, Eric Sun, Alina Isakova, Aaron J. Wilk, and James Zou. 2026. [CellVoyager: AI CompBio agent generates new insights by autonomously analyzing biological data](#). *Nature Methods*, 23:749–759.
- Akari Asai, Jacqueline He, Rulin Shao, Weijia Shi, Amanpreet Singh, Joseph Chee Chang, Kyle Lo, Luca Soldaini, Sergey Feldman, et al. 2026. [Synthesizing scientific literature with retrieval-augmented language models](#). *Nature*, 650:857–863.
- Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2025. [ResearchAgent: Iterative research idea generation over scientific literature with large language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6709–6738, Albuquerque, New Mexico. Association for Computational Linguistics.
- Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, et al. 2025. [MLE-bench: Evaluating machine learning agents on machine learning engineering](#). In *Proceedings of the International Conference on Learning Representations*.
- Yuan Chang, Ziyue Li, Hengyuan Zhang, Yuanbo Kong, Yanru Wu, Hayden Kwok-Hay So, Zhijiang Guo, Liya Zhu, and Ngai Wong. 2025. [TreeReview: A dynamic tree of questions framework for deep and efficient LLM-based scientific peer review](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 15651–15682, Suzhou, China. Association for Computational Linguistics.
- Qiguang Chen, Mingda Yang, Libo Qin, Jinhao Liu, Zheng Yan, Jiannan Guan, Dengyun Peng, Yiyang Ji, Hanjing Li, et al. 2025a. [AI4Research: A survey of artificial intelligence for scientific research](#). *arXiv preprint arXiv:2507.01903v2*.
- Tingting Chen, Srinivas Anumasa, Beibei Lin, Vedant Shah, Anirudh Goyal, and Dianbo Liu. 2025b. [Auto-Bench: An automated benchmark for scientific discovery in LLMs](#). *arXiv preprint arXiv:2502.15224*.
- Wei Chen, Han Ding, Meng Yuan, Zhao Zhang, Deqing Wang, and Fuzhen Zhuang. 2025c. [Bridging social psychology and LLM reasoning: Conflict-aware meta-review generation via cognitive alignment](#). *arXiv preprint arXiv:2503.13879*.
- Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao, Chen

- Wei, et al. 2025d. [ScienceAgentBench: Toward rigorous assessment of language agents for data-driven scientific discovery](#). In *Proceedings of the International Conference on Learning Representations*.
- Herbert H. Clark and Susan E. Brennan. 1991. Grounding in communication. In *Perspectives on Socially Shared Cognition*, pages 127–149. American Psychological Association.
- Le Cong, David Smerkous, Xiaotong Wang, Di Yin, Zaxi Zhang, Ruofan Jin, Yinkai Wang, Michal Gerasimiuk, Ravi K. Dinesh, et al. 2025. [LabOS: The AI-XR co-scientist that sees and works with humans](#). *arXiv preprint arXiv:2510.14861v2*.
- Paul Dourish and Victoria Bellotti. 1992. Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, pages 107–114.
- Jiangshu Du, Yibo Wang, Wenting Zhao, Zhongfen Deng, Shuaiqi Liu, Renze Lou, Henry Peng Zou, Pranav Narayanan Venkit, Nan Zhang, et al. 2024. [LLMs assist NLP researchers: Critique paper \(meta-reviewing\)](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5081–5099, Miami, Florida, USA. Association for Computational Linguistics.
- Haoyang Fang, Boran Han, Nick Erickson, Xiyuan Zhang, Su Zhou, Anirudh Dagar, Jiani Zhang, Ali Caner Turkmen, Cuixiong Hu, et al. 2025. [MLZero: A multi-agent system for end-to-end machine learning automation](#). In *Advances in Neural Information Processing Systems*.
- K. J. Kevin Feng, Kevin Pu, Matt Latzke, Tal August, Pao Siangliulue, Jonathan Bragg, Daniel S. Weld, Amy X. Zhang, and Joseph Chee Chang. 2026. [Co-coa: Co-planning and co-execution with AI agents](#). In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery.
- Shanghai Gao, Richard Zhu, Pengwei Sui, Zhenglun Kong, Sufian Aldogom, Yepeng Huang, Ayush Noori, Reza Shamji, Krishna Parvataneni, et al. 2025. [Democratizing AI scientists using ToolUniverse](#). *arXiv preprint arXiv:2509.23426v2*.
- Ali Essam Ghareeb, Benjamin Chang, Ludovico Mitchener, Angela Yiu, Caralyn J. Szostkiewicz, Dmytro Shved, Gavin J. Gyimesi, Jon M. Laurent, Samantha M. Wright, et al. 2026. [A multi-agent system for automating scientific discovery](#). *Nature*.
- Catalina Gomez, Sue Min Cho, Shichang Ke, Chien-Ming Huang, and Mathias Unberath. 2024. [Human-AI collaboration is not very collaborative yet: A taxonomy of interaction patterns in AI-assisted decision making from a systematic review](#). *Frontiers in Computer Science*, 6.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Petar Sirkovic, Artiom Myaskovsky, Grzegorz Glowaty, Felix Weissenberger, Alessio Orlandi, et al. 2026. [Accelerating scientific discovery with Co-Scientist](#). *Nature*.
- Palash Goyal, Mihir Parmar, Yiwen Song, Hamid Palangi, Tomas Pfister, and Jinsung Yoon. 2026. [ScholarPeer: A context-aware multi-agent framework for automated peer review](#). *arXiv preprint arXiv:2601.22638*.
- Chenyang Gu, Jiahao Cheng, Meicong Zhang, Pujun Zheng, Jinqian Zheng, and Guoxiu He. 2026. [MoRI: Learning motivation-grounded reasoning for scientific ideation in large language models](#). *arXiv preprint arXiv:2603.19044v3*.
- Xuemei Gu and Mario Krenn. 2025. [Interesting scientific idea generation using knowledge graphs and LLMs: Evaluations with 100 research group leaders](#). *arXiv preprint arXiv:2405.17044v3*.
- Sikun Guo, Amir Hassan Shariatmadari, Guangzhi Xiong, Albert Huang, Myles Kim, Corey M. Williams, Stefan Bekiranov, and Aidong Zhang. 2025. [IdeaBench: Benchmarking large language models for research idea generation](#). In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, pages 5888–5899, New York, NY, USA. Association for Computing Machinery.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xi-angliang Zhang. 2024. [Large language model based multi-agents: A survey of progress and challenges](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 8048–8057.
- Carl Gutwin and Saul Greenberg. 2002. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW)*, 11(3):411–446.
- Tal Ifargan, Lukas Hafner, Maor Kern, Ori Alcalay, and Roy Kishony. 2025. [Autonomous LLM-driven research — from data to human-verifiable research papers](#). *NEJM AI*, 2(1).
- Peter Jansen, Oyvind Tafjord, Marissa Radensky, Pao Siangliulue, Tom Hope, Bhavana Dalvi Mishra, Bodhisattwa Prasad Majumder, Daniel S. Weld, and Peter Clark. 2025. [CodeScientist: End-to-end semi-automated scientific discovery with code-based experimentation](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 13370–13467, Vienna, Austria. Association for Computational Linguistics.
- Pengcheng Jiang, Jiacheng Lin, Zhiyi Shi, Zifeng Wang, Luxi He, Yichen Wu, Ming Zhong, Peiyang Song, Qizheng Zhang, et al. 2026. [Adaptation of agentic AI: A survey of post-training, memory, and skills](#). *arXiv preprint arXiv:2512.16301v3*.

- Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, Yijia Xiao, and Jindong Wang. 2024. [AgentReview: Exploring peer review dynamics with LLM agents](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1208–1226, Miami, Florida, USA. Association for Computational Linguistics.
- Liqiang Jing, Zhehui Huang, Xiaoyang Wang, Wenlin Yao, Wenhao Yu, Kaixin Ma, Hongming Zhang, Xinya Du, and Dong Yu. 2025. [DSBench: How far are data science agents from becoming data science experts?](#) In *Proceedings of the International Conference on Learning Representations*.
- Nolan Koblishke, Hyunseok Jang, Kristen Menou, and Mohamad Ali-Dib. 2025. [Gravity-Bench-v1: A benchmark on gravitational physics discovery for agents](#). In *Proceedings of the 42nd International Conference on Machine Learning*.
- Patrick Tser Jern Kon, Jiachen Liu, Qiuyi Ding, Yiming Qiu, Zhenning Yang, Yibo Huang, Jayanth Srinivasa, Myungjin Lee, Mosharaf Chowdhury, and Ang Chen. 2025. [Curie: Toward rigorous and automated scientific experimentation with AI agents](#). *arXiv preprint arXiv:2502.16069*.
- Thomas S. Kuhn. 2012. *The Structure of Scientific Revolutions*, 4 edition. University of Chicago Press. 50th Anniversary Edition; with an introductory essay by Ian Hacking.
- Ed Li, Junyu Ren, Xintian Pan, Cat Yan, Chuanhao Li, Dirk Bergemann, and Zhuoran Yang. 2025a. [Build your personalized research group: A multi-agent framework for continual and interactive science automation](#). *arXiv preprint arXiv:2510.15624*.
- Jinze Li, Yang Zhang, Xin Yang, Jiayi Qu, Jinfeng Xu, Shuo Yang, Junhua Ding, and Edith Cheuk-Han Ngai. 2026. [OCR-Memory: Optical context retrieval for long-horizon agent memory](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10409–10420, San Diego, California, United States. Association for Computational Linguistics.
- Ruochen Li, Teerth Patel, Qingyun Wang, and Xinya Du. 2025b. [MLR-Copilot: Autonomous machine learning research based on large language models agents](#). *arXiv preprint arXiv:2408.14033*.
- Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, et al. 2024. [Monitoring AI-modified content at scale: A case study on the impact of ChatGPT on AI conference peer reviews](#). In *Forty-first International Conference on Machine Learning*.
- Ethan Lin, Zhiyuan Peng, and Yi Fang. 2025. [Evaluating and enhancing large language models for novelty assessment in scholarly publications](#). In *Proceedings of the 1st Workshop on AI and Scientific Discovery: Directions and Opportunities*, pages 46–57, Albuquerque, New Mexico, USA. Association for Computational Linguistics.
- Shengchao Liu, Jiong Xiao Wang, Yijin Yang, Chengpeng Wang, Ling Liu, Hongyu Guo, and Chaowei Xiao. 2024a. [Conversational drug editing using retrieval and domain feedback](#). In *Proceedings of the International Conference on Learning Representations*.
- Shengchao Liu, Hannan Xu, Yan Ai, Huanxin Li, Yoshua Bengio, and Harry Guo. 2025a. [Expert-guided LLM reasoning for battery discovery: From AI-driven hypothesis to synthesis and characterization](#). *arXiv preprint arXiv:2507.16110*.
- Sizhe Liu, Yizhou Lu, Siyu Chen, Xiyang Hu, Jieyu Zhao, Yingzhou Lu, and Yue Zhao. 2025b. [DrugAgent: Automating AI-aided drug discovery programming through LLM multi-agent collaboration](#). *arXiv preprint arXiv:2411.15692v2*.
- Yujie Liu, Zonglin Yang, Tong Xie, Jinjie Ni, Ben Gao, Yuqiang Li, Shixiang Tang, Wanli Ouyang, Erik Cambria, and Dongzhan Zhou. 2026. [ResearchBench: Benchmarking LLMs in scientific discovery via inspiration-based task decomposition](#). *arXiv preprint arXiv:2503.21248v3*.
- Zijun Liu, Kaiming Liu, Yiqi Zhu, Xuanyu Lei, Zonghan Yang, Zhenhe Zhang, Peng Li, and Yang Liu. 2024b. [AIGS: Generating science from AI-powered automated falsification](#). OpenReview Archive.
- Renze Lou, Hanzi Xu, Sijia Wang, Jiangshu Du, Ryo Kamoi, Xiaoxin Lu, Jian Xie, Yuxuan Sun, Yusen Zhang, et al. 2025. [AAAR-1.0: Assessing AI’s potential to assist research](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 40361–40383. PMLR.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. 2026. [Towards end-to-end automation of AI research](#). *Nature*, 651:914–919.
- Yi Luo, Linghang Shi, Yihao Li, Aobo Zhuang, Yeyun Gong, Ling Liu, and Chen Lin. 2025a. [From intention to implementation: Automating biomedical research via LLMs](#). *Science China Information Sciences*, 68(170105).
- Ziming Luo, Zonglin Yang, Zexin Xu, Wei Yang, and Xinya Du. 2025b. [LLM4SR: A survey on large language models for scientific research](#). *arXiv preprint arXiv:2501.04306*.
- Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh Meena, Aryan Prakhar, Tirth Vora, Tushar Khot, Ashish Sabharwal, and Peter Clark. 2025. [DiscoveryBench: Towards data-driven discovery with large language models](#). In *Proceedings of the International Conference on Learning Representations*.
- Ludovico Mitchener, Angela Yiu, Benjamin Chang, Mathieu Bourdenx, Tyler Nadolski, Arvis Sulovari, Eric C. Landsness, Daniel L. Barabasi, Sidharth Narayanan, et al. 2025. [Kosmos: An AI](#)

- scientist for autonomous discovery. *arXiv preprint arXiv:2511.02824v2*.
- David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G. Altman, and The PRISMA Group. 2009. Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *BMJ*, 339:b2535.
- Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vincent Moens, Mikhail Plekhanov, Amar Budhiraja, Despoina Magka, et al. 2025. *MLGym: A new framework and benchmark for advancing AI research agents*. In *Proceedings of the Conference on Language Modeling*.
- Andrew B. Neang, Will Sutherland, Michael W. Beach, and Charlotte P. Lee. 2021. Data integration as coordination: The articulation of data work in an ocean science collaboration. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–25.
- Alexander Novikov, Ng  n V  , Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, et al. 2025. *AlphaEvolve: A coding agent for scientific and algorithmic discovery*. *arXiv preprint arXiv:2506.13131*.
- Gary M. Olson and Judith S. Olson. 2000. Distance matters. *Human-Computer Interaction*, 15(2-3):139–178.
- Jiefu Ou, William Walden, Kate Sanders, Zhengping Jiang, Kaiser Sun, Jeffrey Cheng, William Jurayj, Miriam Wanner, Shaobo Liang, et al. 2025. *CLAIM-CHECK: How grounded are LLM critiques of scientific papers?* In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 21712–21735, Suzhou, China. Association for Computational Linguistics.
- Matthew J. Page, Joanne E. McKenzie, Patrick M. Bossuyt, Isabelle Boutron, Tammy C. Hoffmann, Cynthia D. Mulrow, Larissa Shamseer, Jennifer M. Tetzlaff, Elie A. Akl, et al. 2021. *The PRISMA 2020 statement: an updated guideline for reporting systematic reviews*. *BMJ*, 372:n71.
- Jingyuan Qi, Zian Jia, Minqian Liu, Wangzhi Zhan, Junkai Zhang, Xiaofei Wen, Jingru Gan, Jianpeng Chen, Qin Liu, et al. 2025. *MetaScientist: A human-AI synergistic framework for automated mechanical metamaterial design*. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 404–436, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yansheng Qiu, Haoquan Zhang, Zhaopan Xu, Ming Li, Diping Song, Zheng Wang, and Kaipeng Zhang. 2025. *AI Idea Bench 2025: AI research idea generation benchmark*. *arXiv preprint arXiv:2504.14191v3*.
- Marissa Radensky, Simra Shahid, Raymond Fok, Pao Siangliulue, Tom Hope, and Daniel S. Weld. 2026. *Scideator: Human-LLM compound system for scientific ideation through facet recombination and novelty evaluation*. In *Proceedings of the ACM Conference on AI and Agentic Systems*, pages 348–374. Association for Computing Machinery.
- Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, et al. 2024. *Mathematical discoveries from program search with large language models*. *Nature*, 625:468–475.
- Kai Ruan, Xuan Wang, Jixiang Hong, Peng Wang, Yang Liu, and Hao Sun. 2026. *Evaluating LLMs’ divergent thinking capabilities for scientific idea generation with minimal context*. *Nature Communications*, 17:3625.
- Giuseppe Russo Latona, Manoel Horta Ribeiro, Tim R. Davidson, Veniamin Veselovsky, and Robert West. 2025. *The AI review lottery: Widespread AI-assisted peer reviews boost paper scores and acceptance rates*. *Proceedings of the ACM on Human-Computer Interaction*, 9(7):1–28.
- Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu, and Emad Barsoum. 2025. *Agent Laboratory: Using LLM agents as research assistants*. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5977–6043, Suzhou, China. Association for Computational Linguistics.
- Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW seriously: Supporting articulation work. *Computer supported cooperative work (CSCW)*, 1(1):7–40.
- Kjeld Schmidt and Carla Simone. 1996. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Computer Supported Cooperative Work (CSCW)*, 5(2):155–200.
- Chenyang Shao, Dehao Huang, Yu Li, Keyu Zhao, Weiquan Lin, Yining Zhang, Qingbin Zeng, Zhiyu Chen, Tianxing Li, et al. 2025. *OmniScientist: Toward a co-evolving ecosystem of human and AI scientists*. *arXiv preprint arXiv:2511.16931v2*.
- Erzhuo Shao, Yifang Wang, Yifan Qian, Zhenyu Pan, Han Liu, and Dashun Wang. 2026. *SciSciGPT: Advancing human-AI collaboration in the science of science*. *Nature Computational Science*, 6:301–315.
- Chuhan Shi, Xiaoquan Ren, Yifang Wang, Junze Li, Yushi Sun, Yawen Luo, and Rui Sheng. 2026. *A survey of human-AI collaboration for scientific discovery*. *TechRxiv*. Preprint, version 2.
- Parshin Shojaee, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, and Chandan K. Reddy. 2025.

- LLM-SR: Scientific equation discovery via programming with large language models. In *Proceedings of the International Conference on Learning Representations*.
- Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2025. Can LLMs generate novel research ideas? A large-scale human study with 100+ NLP researchers. In *The Thirteenth International Conference on Learning Representations*.
- Yiwen Song, Yale Song, Tomas Pfister, and Jinsung Yoon. 2026a. PaperOrchestra: A multi-agent framework for automated AI research paper writing. *arXiv preprint arXiv:2604.05018v1*.
- Zhangde Song, Jieyu Lu, Yuanqi Du, Botao Yu, Thomas M. Pruyn, Yue Huang, Kehan Guo, Xiuzhe Luo, Yuanhao Qu, et al. 2026b. Evaluating large language models in scientific discovery. *arXiv preprint arXiv:2512.15567v2*.
- Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, et al. 2025. Paper-Bench: Evaluating AI’s ability to replicate AI research. In *Proceedings of the 42nd International Conference on Machine Learning*.
- Anselm Strauss. 1985. Work and the division of labor. *Sociological quarterly*, 26(1):1–19.
- Haoyang Su, Renqi Chen, Shixiang Tang, Zhenfei Yin, Xinzhe Zheng, Jinzhe Li, Biqing Qi, Qi Wu, Hui Li, et al. 2025. Many heads are better than one: Improved scientific idea generation by a LLM-based multi-agent system. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 28201–28240, Vienna, Austria. Association for Computational Linguistics.
- Jiabin Tang, Lianghao Xia, Zhonghang Li, and Chao Huang. 2025. AI-Researcher: Autonomous scientific innovation. In *Advances in Neural Information Processing Systems*.
- Guiyao Tie, Pan Zhou, and Lichao Sun. 2026. A survey of AI scientists. *arXiv preprint arXiv:2510.23045v5*.
- Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D. Nguyen. 2025. Multi-agent collaboration mechanisms: A survey of LLMs. *arXiv preprint arXiv:2501.06322*.
- Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12):2293–2303.
- Rosni Vasu, Chandrayee Basu, Bhavana Dalvi Mishra, Cristina Sarasua, Peter Clark, and Abraham Bernstein. 2025. HyperER: Literature-grounded hypothesis generation and distillation with provenance. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25413–25438, Suzhou, China. Association for Computational Linguistics.
- Vanshika Vats, Marzia Binta Nizam, Minghao Liu, Ziyuan Wang, Richard Ho, Mohnish Sai Prasad, Vincent Titterton, Sai Venkat Malreddy, Riya Aggarwal, et al. 2025. A survey on human-AI collaboration with large foundation models. *arXiv preprint arXiv:2403.04931v3*.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, et al. 2023. Scientific discovery in the age of artificial intelligence. *Nature*, 620:47–60.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, et al. 2024a. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(186345).
- Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2024b. SciMON: Scientific inspiration machines optimized for novelty. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 279–299, Bangkok, Thailand. Association for Computational Linguistics.
- Yunfei Wang, Yeqin Zhang, Yuyang Wu, Liang Lu, Phi Le Nguyen, Xiaoliang Wang, and Cam-Tu Nguyen. 2025a. MLAlgo-Bench: Can machines implement machine learning algorithms? In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 14298–14329, Suzhou, China. Association for Computational Linguistics.
- Zifeng Wang, Lang Cao, Qiao Jin, Joey Chan, Nicholas Wan, Behdad Afzali, Hyun-Jin Cho, Chang-In Choi, Mehdi Emamverdi, et al. 2025b. A foundation model for human-AI collaboration in medical literature mining. *Nature Communications*, 16:8361.
- Jiaqi Wei, Yuejin Yang, Xiang Zhang, Yuhao Chen, Xiang Zhuang, Zhangyang Gao, Dongzhan Zhou, Guangshuai Wang, Zhiqiang Gao, et al. 2025. From AI for science to agentic science: A survey on autonomous scientific discovery. *arXiv preprint arXiv:2508.14111v2*.
- Yixuan Weng, Minjun Zhu, Guangsheng Bao, Hongbo Zhang, Jindong Wang, Yue Zhang, and Linyi Yang. 2025. CycleResearcher: Improving automated research via automated review. In *The Thirteenth International Conference on Learning Representations*.
- Yixuan Weng, Minjun Zhu, Qiujie Xie, QiYao Sun, Zhen Lin, Sifan Liu, and Yue Zhang. 2026. DeepScientist: Advancing frontier-pushing scientific findings progressively. In *Proceedings of the International Conference on Learning Representations*.

- Hjalmar Wijk, Tao Roa Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence Chan, Michael Chen, Joshua M. Clymer, et al. 2025. **RE-Bench: Evaluating frontier AI R&D capabilities of language model agents against human experts**. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 66772–66832. PMLR.
- Sihong Wu, Yiling Ma, Yilun Zhao, Tiansheng Hu, Owen Jiang, Manasi Patwardhan, and Arman Cohen. 2026. **RbtAct: Rebuttal as supervision for actionable review feedback generation**. *arXiv preprint arXiv:2603.09723v2*.
- Wenqing Wu, Chengzhi Zhang, and Yi Zhao. 2025. **Automated novelty evaluation of academic paper: A collaborative approach integrating human and large language model knowledge**. *Journal of the Association for Information Science and Technology*, 76(11):1452–1469.
- Yanzheng Xiang, Hanqi Yan, Shuyin Ouyang, Lin Gui, and Yulan He. 2025. **SciReplicate-Bench: Benchmarking LLMs in agent-driven algorithmic reproduction from research papers**. In *Proceedings of the Conference on Language Modeling*.
- Guangzhi Xiong, Eric Xie, Corey Williams, Myles Kim, Amir Hassan Shariatmadari, Sikun Guo, Stefan Bekiranov, and Aidong Zhang. 2025. **Toward reliable scientific hypothesis generation: Evaluating truthfulness and hallucination in large language models**. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pages 7849–7857.
- Weikai Xu, Chengrui Huang, Shen Gao, and Shuo Shang. 2025. **LLM-based agents for tool learning: A survey**. *Data Science and Engineering*, 10:533–563.
- Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune, and David Ha. 2025. **The AI scientist-v2: Workshop-level automated scientific discovery via agentic tree search**. *arXiv preprint arXiv:2504.08066*.
- Zonglin Yang, Xinya Du, Junxian Li, Jie Zheng, Soujanya Poria, and Erik Cambria. 2024. **Large language models for automated open-domain scientific hypotheses discovery**. In *ICML 2024 Workshop on AI for Science*.
- Zonglin Yang, Wanhao Liu, Ben Gao, Tong Xie, Yuqiang Li, Wanli Ouyang, Soujanya Poria, Erik Cambria, and Dongzhan Zhou. 2025. **MOOSE-Chem: Large language models for rediscovering unseen chemistry scientific hypotheses**. In *Proceedings of the International Conference on Learning Representations*.
- Geyan Ye, Xibao Cai, Houtim Lai, Xing Wang, Junhong Huang, Longyue Wang, Wei Liu, and Xiangxiang Zeng. 2025. **DrugAssist: A large language model for molecule optimization**. *Briefings in Bioinformatics*, 26(1):bbae693.
- Zahra Zahedi and Subbarao Kambhampati. 2021. **Human-AI symbiosis: A survey of current approaches**. *arXiv preprint arXiv:2103.09990*.
- Bo Zhang, Shiyang Feng, Xiangchao Yan, Jiakang Yuan, Runmin Ma, Yusong Hu, Zhiyin Yu, Xiaohan He, Songtao Huang, et al. 2025a. **InternAgent: When agent becomes the scientist – building closed-loop system from hypothesis to verification**. *arXiv preprint arXiv:2505.16938*.
- Haoxuan Zhang, Ruochi Li, Zhenni Liang, Mehri Sattari, Phat Vo, Collin Qu, Ting Xiao, Junhua Ding, Yang Zhang, and Haihua Chen. 2026a. **AdaQE-CG: Adaptive query expansion for web-scale generative AI model and data card generation**. In *Proceedings of the ACM Web Conference 2026*, WWW ’26, pages 4360–4371, New York, NY, USA. Association for Computing Machinery.
- Haoxuan Zhang, Ruochi Li, Yang Zhang, Ting Xiao, Jiangping Chen, Junhua Ding, and Haihua Chen. 2026b. **Evolving roles of LLMs in scientific innovation: Assistant, collaborator, scientist, and evaluator**. *Preprint*, arXiv:2507.11810.
- Shuning Zhang, Hui Wang, and Xin Yi. 2025b. **Exploring collaboration patterns and strategies in human-AI co-creation through the lens of agency: A scoping review of the top-tier HCI literature**. *Proceedings of the ACM on Human-Computer Interaction*, 9(7):1–43.
- Yang Zhang, Yufei Wang, Quan Z. Sheng, Adnan Mahmood, Wei Emma Zhang, and Rongying Zhao. 2021. **TDM-CFC: Towards document-level multi-label citation function classification**. In *Web Information Systems Engineering – WISE 2021*, volume 13081 of *Lecture Notes in Computer Science*, pages 363–376, Cham. Springer.
- Yang Zhang, Yufei Wang, Quan Z. Sheng, Adnan Mahmood, Wei Emma Zhang, and Rongying Zhao. 2023a. **Hybrid data augmentation for citation function classification**. In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Yang Zhang, Yufei Wang, Quan Z. Sheng, Lina Yao, Haihua Chen, Kai Wang, Adnan Mahmood, Wei Emma Zhang, Munazza Zaib, et al. 2025c. **Deep learning meets bibliometrics: A survey of citation function classification**. *Journal of Informetrics*, 19(1):101608.
- Yang Zhang, Yufei Wang, Kai Wang, Quan Z. Sheng, Lina Yao, Adnan Mahmood, Wei Emma Zhang, and Rongying Zhao. 2023b. **When large language models meet citation: A survey**. *arXiv preprint arXiv:2309.09727*.
- Yang Zhang, Rongying Zhao, Yufei Wang, Haihua Chen, Adnan Mahmood, Munazza Zaib, Wei Emma Zhang, and Quan Z. Sheng. 2022. **Towards employing native information in citation function classification**. *Scientometrics*, 127(11):6557–6577.

Yunxiang Zhang, Muhammad Khalifa, Shitanshu Bhushan, Grant D. Murphy, Lajanugen Logeswaran, Jaekyeom Kim, Moontae Lee, Honglak Lee, and Lu Wang. 2025d. [MLRC-Bench: Can language agents solve machine learning research challenges?](#) In *Advances in Neural Information Processing Systems*.

Bingchen Zhao, Jenny Zhang, Chenxi Whitehouse, Mingqi Jiang, Michael Shvartsman, Abhishek Charnalia, Despoina Magka, Tatiana Shavrina, Derek Dunfield, et al. 2026. [APRES: An agentic paper revision and evaluation system](#). *arXiv preprint arXiv:2603.03142*.

Tianshi Zheng, Zheyue Deng, Hong Ting Tsang, Weiqi Wang, Jiaxin Bai, Zihao Wang, and Yangqiu Song. 2025. [From automation to autonomy: A survey on large language models in scientific discovery](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 17733–17750, Suzhou, China. Association for Computational Linguistics.

Xianrui Zhong, Bowen Jin, Siru Ouyang, Yanzen Shen, Qiao Jin, Yin Fang, Zhiyong Lu, and Jiawei Han. 2025. [Benchmarking retrieval-augmented generation for chemistry](#). In *Proceedings of the Conference on Language Modeling*.

Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang. 2025. [DeepReview: Improving LLM-based paper review with human-like deep thinking process](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29330–29355, Vienna, Austria. Association for Computational Linguistics.

Xiang Zhuang, Keyan Ding, Tianwen Lyu, Yinuo Jiang, Xiaotong Li, Zhuoyi Xiang, Zeyuan Wang, Ming Qin, Kehua Feng, et al. 2025. [Advancing biomolecular understanding and design following human instructions](#). *Nature Machine Intelligence*, 7:1154–1167.

Henry Peng Zou, Wei-Chieh Huang, Yaozu Wu, Jizhou Guo, Yankai Chen, Chunyu Miao, Hoang Nguyen, Yue Zhou, Weizhi Zhang, et al. 2026. [LLM-based human-agent collaboration and interaction systems: A survey](#). *arXiv preprint arXiv:2505.00753v5*.

A Survey Positioning

This appendix provides supplementary details for the survey scope, CSCW and Kuhnian conceptual grounding, operational coding rubric, corpus construction, analytic coding, and extended tables. Appendix A.1 positions our work relative to existing surveys. Appendix A.2 details the CSCW foundations of our coordination coding. Appendix A.3 clarifies how Kuhn’s account of scientific change is used as an analytical lens. Appendix B provides the operational rubric for coding epistemic context, coordination demand, failure modes, and evaluation

targets. Appendix C describes corpus construction and coding procedures. Appendix D reports details of system-level coding used to support the synthesis in the main text. Appendix E provides additional tables for operational definitions, survey comparison, and benchmark mapping.

A.1 Comparison with Existing Surveys

To position CSCW×Kuhn, we compare our survey with prior work on AI for science, agentic systems, multi-agent systems, and human–agent collaboration. These survey traditions provide important foundations, but differ in primary analytical focus. AI4Science surveys typically organize systems by scientific tasks or research pipelines. Agentic-system surveys focus on model architecture, planning, tool use, memory, and multi-agent coordination. Human–agent collaboration surveys provide stronger accounts of human roles, interaction, and agency distribution.

Our survey pursues a different analytical objective. Whereas prior surveys mainly organize systems by task, architecture, or application domain, we analyze human–agent scientific collaboration as an epistemically conditioned coordination process—connecting scientific context, coordination demand, human role, evaluation targets, and failure modes in a unified coding framework. We do not intend to reclassify prior surveys’ systems, nor suggest that their conclusions are incorrect. We instead reinterpret representative systems through a coordination-aware lens to address the analytical coverage gap. Table 3 summarizes our positioning.

A.2 CSCW Foundations for Coordination Coding

CSCW provides the conceptual foundation for the coordination dimensions used in CSCW×Kuhn. We use CSCW as a theory of cooperative work that emphasizes the practical work required to make collaboration possible. Prior research in CSCW has shown that cooperative systems must support task execution, and also the organization of interdependent activities, the maintenance of shared understandings, and the visibility of ongoing work (Schmidt and Bannon, 1992; Schmidt and Simone, 1996; Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002).

This appendix clarifies how we translate CSCW concepts into coding dimensions. **Articulation work** refers to the organizing and repair work required to coordinate interdependent tasks, includ-

ing sequencing, role assignment, handoffs, integration, and recovery from breakdowns (Strauss, 1985; Schmidt and Bannon, 1992). In human-agent scientific systems, this includes agent planning, tool selection, protocol execution, workflow decomposition, and human intervention points. **Common ground** refers to the extent to which humans and agents share sufficient understanding of goals, assumptions, constraints, evidence, methods, and task state for coordinated action (Clark and Brennan, 1991). In scientific workflows, this includes alignment between research questions, retrieved literature, data, claims, benchmarks, and evaluation criteria. **Awareness** refers to the visibility of actions, intermediate states, errors, decisions, and consequences during collaborative work (Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002). In human-agent scientific systems, this includes logs, provenance, uncertainty signals, rationales, disagreement traces, and checkpoints that allow humans to monitor and contest agent behavior.

These three dimensions were selected because they are theoretically grounded in CSCW and operationally useful for analyzing agentic scientific workflows. The dimensions allow us to code for whether a system supports the coordination work needed for reliable scientific collaboration, rather than only if it produces a correct output. They also connect directly to the failure modes identified in the main text: procedural drift reflects failures of articulation work; grounding substitution reflects failures of common ground; evaluative flattening and epistemic black-boxing reflect failures of awareness, inspectability, and contestability.

A.3 Kuhnian Lens for Scientific Contexts

We use Kuhn’s account of scientific change as an analytical heuristic for distinguishing different epistemic conditions in scientific work (Kuhn, 2012). In Kuhn’s framework, scientific communities operate within paradigms that define legitimate problems, accepted methods, evidential standards, and criteria for successful explanation. Under stable paradigms, most work proceeds as normal science, where researchers solve puzzles within established assumptions. Scientific work may also involve anomalies (i.e., findings or claims that do not fit prevailing expectations). Some anomalies are resolved within the existing paradigm, whereas others motivate deeper conceptual reconsideration.

This distinction is useful for analyzing human-agent collaboration since different epistemic con-

ditions impose different coordination demands. Puzzle-solving in normal science requires reliable execution, workflow tracking, and constraint management. Anomaly evaluation requires comparison among competing interpretations, careful treatment of evidence, and preservation of disagreement. Paradigm-shift exploration requires broader conceptual search, provenance over speculative ideas, and mechanisms for feasibility, falsifiability, and follow-up validation.

Accordingly, we operationalize three epistemic modes: *puzzle-solving in normal science*, *anomaly evaluation*, and *paradigm-shift exploration*. These modes are not intended as exhaustive categories of all scientific activity, nor do they imply that real workflows remain fixed within a single mode. Instead, they provide a tractable coding scheme for analyzing how coordination demands change as scientific work moves from stable execution to contested judgment and exploratory reframing.

B Operational Coding Rubric

We operationalize CSCW $\times$ Kuhn as a qualitative coding scheme that links epistemic context, coordination demand, failure mode, and evaluation target. The primary coding unit is the named system or artifact. When a paper presents clearly separable workflow components, those components inform the system-level coding, but each system receives one dominant epistemic-mode assignment for comparative analysis. For each system, we first identify the dominant epistemic context, then assess coordination support along three CSCW dimensions: articulation work (α), common ground (γ), and awareness (ω). We then identify potential failure risks, corresponding to the failure modes defined in the main text, and map them to evaluation targets.

Epistemic context. **Puzzle-solving in normal science** is assigned when the task has stable goals, established methods, and predefined success criteria, such as protocol execution, benchmarked prediction, evidence extraction, or constrained optimization. **Anomaly evaluation** is assigned when the task involves judging uncertain, contested, novel, or potentially problematic claims, such as peer review, novelty assessment, rebuttal analysis, critique, or claim verification. **Paradigm-shift exploration** is assigned when the task involves open-ended search for alternative hypotheses, cross-domain analogies, conceptual recombination, or long-horizon research directions where

goals and evaluation criteria are not fully stable.

Coordination demand. We code each CSCW coordination dimension on an ordinal 0–3 scale. These scores are not treated as quantitative performance measures; rather, they provide a transparent way to compare what forms of coordination support are made available by different systems.

- **Articulation work** (α). 0 = no task decomposition or coordination structure; 1 = fixed pipeline with predefined steps; 2 = dynamic planning with explicit human checkpoints; 3 = adaptive replanning across agents and humans, including hand-offs, repair, and revision.
- **Common ground** (γ). 0 = no shared state or explicit task context; 1 = prompt-level task instruction only; 2 = explicit shared memory, workspace, or task state; 3 = shared assumptions, evidence, uncertainty, constraints, and role model are visible and maintained.
- **Awareness** (ω). 0 = final output only; 1 = logs or execution traces are available; 2 = intermediate states, provenance, and decision paths are visible; 3 = humans can inspect, contest, intervene, and revise system actions or outputs.

Worked coding examples. To illustrate how the ordinal rubric applies in practice, let us walk through two examples of the coding process and the rationale behind the coordination demand scores. Consider the systems: LEADS (Wang et al., 2025b) and ResearchAgent (Baek et al., 2025).

LEADS (Wang et al., 2025b) is coded at the system level as puzzle-solving in normal science with $(\alpha, \gamma, \omega) = (1, 2, 3)$. Its literature-mining workflow is decomposed into six predefined sub-tasks; thus **articulation work** is coded as $\alpha = 1$ (Wang et al., 2025b). The reason why it is not coded as $\alpha = 2$ or 3 is because task order and structure are set in advance, and the paper does not describe any mechanism for dynamic task replanning. **Common ground** is coded as $\gamma = 2$, since screening and extraction operate in a structured task context, including PICO criteria, target-field definitions, eligibility assessments, and shared persistent artifacts (ranked spreadsheet) (Wang et al., 2025b). That exceeds bare prompt-level instruction ($\gamma = 1$), though no explicit role model or shared assumptions are maintained ($\gamma = 3$). **Awareness** is coded as $\omega = 3$ because human experts can inspect criterion-level predictions and rationales,

select studies independently of the AI ranking, and directly edit AI-extracted field values before finalizing them (Wang et al., 2025b). This would not be a simple decision log ($\omega = 1$) and exceeds $\omega = 2$.

ResearchAgent (Baek et al., 2025) is coded at the system level as paradigm-shift exploration with $(\alpha, \gamma, \omega) = (1, 2, 2)$. Although its problem, method, and experiment-design outputs are iteratively revised by ReviewingAgents, the workflow remains predefined, and the paper does not report any inclusion of human checkpoints during generation, yielding $\alpha = 1$ for **articulation work** rather than 2 (replanning is agent-only, not human-in-the-loop) or 0 (a decomposed structure does exist) (Baek et al., 2025). **Common ground** is coded as $\gamma = 2$ because the workflow maintains structured context across stages through prior idea components and rationales, retrieved literature, an entity-centric knowledge store, and shared evaluation criteria (Baek et al., 2025). The workflow exceeds prompt-level instruction ($\gamma = 1$) but lacks any persistent, explicit layer for uncertainty or assumptions required for $\gamma = 3$. **Awareness** is coded as $\omega = 2$ because rationales, reviews, and ratings expose intermediate evaluative state (Baek et al., 2025). This goes beyond simple logs or execution traces ($\omega = 1$). However, in this system, human researchers participate primarily in post-hoc evaluation rather than in-loop intervention or revision required for $\omega = 3$ (Baek et al., 2025).

These ordinal scores are qualitative coding judgments about the coordination support exposed by the system design, not performance measures or claims about system quality.

Failure risk and evaluation target. Potential failure risks are identified by comparing the epistemic demands of a task with the coordination support provided by the system. **Procedural drift** is coded when a system performs multi-step work with limited task-state tracking, checkpointing, or replanning support. **Grounding substitution** is coded when outputs depend on evidence but provide weak provenance, unsupported claims, or insufficient source alignment. **Evaluative flattening** is coded when contested scientific judgment is reduced to narrow scores, rankings, or majority opinions. **Proxy optimization** is coded when a system is primarily optimized for benchmark metrics without evidence of downstream scientific validity. **Epistemic black-boxing** is coded when assumptions, intermediate reasoning, provenance, tool use,

or responsibility boundaries are difficult to inspect or contest.

The corresponding evaluation targets are: workflow coordination for procedural drift, claim-evidence grounding for grounding substitution, multidimensional scientific judgment for evaluative flattening, downstream validation for proxy optimization, and inspectability or contestability for epistemic black-boxing.

C Corpus Construction and Analytic Coding

C.1 Survey Protocol

Our goal was to construct a representative corpus of agentic systems, benchmarks, and evaluation studies relevant to human-agent collaboration in scientific innovation. The corpus was designed to support conceptual synthesis and framework development rather than to serve as an exhaustive systematic review. We therefore prioritized breadth across scientific workflows, agentic system designs, and evaluation settings.

Search sources. We searched ACL Anthology, arXiv, OpenReview, ACM Digital Library, IEEE Xplore, Google Scholar, and publisher websites for recent work on scientific agents, AI for science, LLM-based scientific workflows, multi-agent systems, and human-agent collaboration. We prioritized work published between 2024 and 2026 because this period captures the rapid emergence of LLM-based agents, tool-augmented systems, and multi-agent scientific workflows. Earlier work was included when it provided important conceptual grounding or representative system designs.

Search keywords. Search queries combined terms related to scientific workflows, language agents, collaboration, and evaluation. Representative queries included: *LLM scientific discovery*, *AI scientist*, *scientific agents*, *AI for science*, *AI4Science*, *human-agent collaboration*, *human-AI collaboration*, *human-AI collaborative hypothesis generation*, *RAG scientific literature*, *LLM-as-reviewer*, *automated peer review*, *scientific ideation LLM*, *hypothesis generation LLM*, *tool-use agents science*, *tool-augmented LLM scientific workflow*, and *multi-agent scientific discovery*.

Screening process. Candidate papers were screened in three stages. First, titles and abstracts were screened for relevance to scientific

research, discovery, evaluation, or scholarly workflows. Second, full papers were assessed to determine whether the system involved an agentic or LLM-based component relevant to planning, retrieval, critique, judgment, tool use, experimentation, writing, or evaluation, with continuous human input during execution not required for inclusion. Systems with substantial internal agent-agent coordination were included when scientific goals, constraints, validation standards, or accountability remained human-governed; these were treated as human-agent scientific collaboration at the broader sociotechnical level, with agent-agent coordination analyzed as an architectural mechanism within that workflow. Third, papers were selected for synthesis if they contributed analytical coverage for at least one epistemic mode or system cluster in our framework.

Screening results. We identified a total of 130 records. After removing 3 duplicate records, 127 unique candidate papers were screened. During title and abstract screening, 43 papers were excluded, leaving 84 papers for full-text assessment. After full-text assessment, 54 papers were included in the final synthesis, representing 59 named systems and related artifacts. Figure 2 presents a flowchart of the survey protocol, inspired and adapted from the PRISMA framework (Moher et al., 2009; Page et al., 2021) for transparent reporting. As this work is a focused conceptual survey and not a systematic review, we adopt its reporting structure selectively rather than strictly. The accompanying corpus and system-level coding artifact is maintained at <https://github.com/jenlia024/CSCWxKuhn-artifact>.

C.2 Analytic Coding

Inclusion criteria. We included papers that met at least two of the following criteria: (1) the system supports scientific research, scientific discovery, scholarly evaluation, or research-workflow automation; (2) the system uses LLMs, agents, retrieval-augmented generation, tool use, multi-agent coordination, review, critique, or ideation; (3) the system includes or implies a human role in task planning, judgment, oversight, revision, interpretation, or validation; and (4) the paper reports a concrete system, benchmark, method, or evaluation study relevant to scientific work.

Exclusion criteria. We excluded papers that focused only on general-purpose NLP tasks with-




Figure 2: Flow diagram of our corpus selection.

out a scientific workflow, domain-specific prediction without an agentic or collaborative component, purely theoretical discussions without a system or evaluation target, and applications where human involvement was absent, incidental, or not relevant to the scientific output. We also excluded papers whose relevance was limited to background AI capability rather than scientific collaboration, evaluation, or workflow coordination. Papers that met the inclusion criteria were removed if they also met any exclusion criterion.

Coding dimensions. Each included work was coded along four dimensions: *epistemic mode*, *system cluster*, *coordination demand*, and *potential failure risk*. The coding was designed to support the comparative synthesis in the main text rather than to produce mutually exclusive labels for all possible system properties.

- **Epistemic mode** identifies the dominant scientific context supported by the system: puzzle-solving in normal science, anomaly evaluation, or paradigm-shift exploration. When a system could fit multiple modes, we assigned the mode corresponding to its primary workflow role and evaluation target.
- **System cluster** captures the dominant workflow function of the system, such as literature and evidence workflows, design and optimization, protocol execution, shared orchestration, nov-

elty evaluation, peer review, rebuttal support, cross-domain ideation, falsification, recursive AI-scientist workflows, verification/search, or scientific infrastructure.

- **Coordination demand** records how the system supports the three CSCW dimensions used in our framework: articulation work (α), common ground (γ), and awareness (ω). Articulation work concerns task organization and workflow integration; common ground concerns shared assumptions, goals, evidence, and constraints; awareness concerns the visibility of actions, intermediate states, rationales, errors, and outcomes.
- **Potential failure risk** identifies the main coordination vulnerability implied by the system design, task setting, interface, or evaluation setup. These risks were used to synthesize recurring failure modes across systems.

Failure-mode synthesis. After coding individual systems, we compared failure risks across epistemic modes and system clusters. We synthesized five recurring coordination failure modes. **Procedural drift** refers to deviations from intended workflows, protocols, constraints, or task states. **Grounding substitution** occurs when fluency, plausibility, or apparent novelty substitutes for evidential support or domain validity. **Evaluative flattening** occurs when contested scientific judgment is reduced to simplified scores, rankings, criteria, or majority opinions. **Proxy optimization** occurs when systems optimize benchmark metrics or local evaluation signals while losing alignment with broader scientific goals. **Epistemic black-boxing** occurs when provenance, assumptions, intermediate reasoning, tool use, or responsibility become difficult for humans to inspect, contest, or revise.

Use of coding in the survey. The coding was used in two ways. First, it supported the system-level synthesis in Table 2, where representative systems are grouped by epistemic mode, system cluster, coordination demand, and failure risk. Second, it supported the benchmark analysis in Table 5, where evaluation studies are compared according to their coverage of workflow coordination, human oversight, scientific judgment, and downstream validation. This procedure allows the survey to move beyond a list of systems and instead analyze how coordination requirements change across scientific contexts.

Reliability and limitations. The current coding should be interpreted as analytic coding for conceptual synthesis, not as a finalized systematic-review annotation. Coding decisions were checked against the reported system design, workflow role, and evaluation setup of each paper. Because many systems span multiple scientific stages, epistemic modes are not treated as mutually exclusive categories. We assigned the dominant mode to support comparison, while acknowledging that real scientific workflows may move from routine puzzle-solving to anomaly evaluation or exploratory reframing. Future work can extend this codebook through independent annotators, inter-annotator agreement, larger corpora, and deployment-based validation.

D System-Level Coding

Table 2 summarizes the system-level coding used to construct the synthesis in the main text. For each epistemic mode, we group representative systems into workflow-oriented clusters and identify their primary coordination demands and potential failure risks. Icons indicate artifact type, while the coordination-demand column reports articulation work (α), common ground (γ), and awareness (ω). Ordinal coding and metadata are provided in the machine-readable artifact linked in Appendix C.

E Supplementary Tables

This section provides supplementary tables that support the main analysis. Table 2 provides the full system-level coding of representative human-agent scientific collaboration systems, grouped by epistemic mode and workflow cluster. Table 3 compares existing survey traditions with our survey. Table 4 defines the main terms used throughout the paper. Table 5 maps benchmark and evaluation studies by scientific domain, target task, metric, and coverage gap. The coding procedure behind these tables is described in Appendix B.

System Cluster	Representative Systems	Coordination Demand	Potential Failure Risks
Puzzle-Solving in Normal Science			
Literature & evidence workflows	  LEADS (Wang et al., 2025b)	α : subtask decomposition γ : PICO, eligibility, field alignment ω : ranked candidates, rationales, extraction outputs for expert oversight	Missed studies; borderline-case errors; extraction errors
Design & optimization agents	 ChatDrug (Liu et al., 2024a)   DrugAssist (Ye et al., 2025)   InstructBioMol (Zhuang et al., 2025)   ChatBattery (Liu et al., 2025a)   MetaScientist (Qi et al., 2025)	α : generation, editing, optimization subtasks γ : property-format-constraint alignment ω : candidates, feedback, explanations, validation scores	Unsafe, uninterpretable, or infeasible candidate molecules
Execution & protocol systems	  Curie (Kon et al., 2025)   CycleResearcher (Weng et al., 2025)   BioResearcher (Luo et al., 2025a)   LabOS (Cong et al., 2025)	α : protocol planning, code execution, documentation γ : method-data-variable, reproducibility alignment ω : step state, quality checks, deviations, errors, results	Protocol failures; variable, dataset, or deviation mishandling
Shared planning & orchestration systems	  DrugAgent (Liu et al., 2025b)   Agent Laboratory (Schmidgall et al., 2025)   MLZero (Fang et al., 2025)   freephdlabor (Li et al., 2025a)   data-to-paper (Ifargan et al., 2025)   SciSciGPT (Shao et al., 2026)   PaperOrchestra (Song et al., 2026a)   Cocoa (Feng et al., 2026)	α : plan-subtask-role-tool orchestration, review checkpoints γ : goal-artifact-constraint-workspace alignment ω : plans, reports, memory, feedback, intervention points	Context loss; role misalignment; degraded intermediate information
Anomaly Evaluation			
Hypothesis generation & novelty evaluation	  SchNovel / RAG-Novelly (Lin et al., 2025)   TruthHypo / KnowHD (Xiong et al., 2025)   Knowledge Interaction Model / Fusion module (Wu et al., 2025)   HypER (Vasu et al., 2025)   MoRI (Gu et al., 2026)	α : novelty, truthfulness, grounding checks γ : review-metadata-literature alignment, knowledge graphs ω : scores, rationales, retrieved evidence, unsupported claims	Conflation of perceived novelty and plausibility with grounded substance
LLM-as-reviewer & critique systems	  DeepReview /  DeepReviewer-14B (Zhu et al., 2025)   CycleReviewer (Weng et al., 2025)   TreeReview (Chang et al., 2025)   MLR-Copilot (Li et al., 2025b)   ScholarPeer (Goyal et al., 2026)	α : novelty, soundness, baseline checks γ : venue-literature-benchmark alignment ω : evidence-cited answers, retrieved literature, question trees, verification logs, confidence signals	Shallow fluency; missed flaws or baselines; hallucinated or unusable feedback
Rebuttal & meta-review systems	  ReviewCritique (Du et al., 2024)   CAF (Chen et al., 2025c)   RbtAct (Wu et al., 2026)	α : deficiency, conflict, actionable feedback checks γ : paper-review-rebuttal-decision alignment ω : deficient review segments, conflict maps, rebuttal uptake, evidence signals	Majority-opinion bias; overreliance on reviews; weak adjudication
Institutional review systems	  Distributional GPT quantification (Liang et al., 2024)   AgentReview (Jin et al., 2024)   APRES (Zhao et al., 2026)	α : revision, review simulation, decision modeling, AI-use monitoring γ : rubric-reviewer role-rebuttal-decision alignment ω : revision comparisons, simulated review traces, decision shifts, AI-use prevalence	Black-box review or revision outputs; simplified evaluative judgments
Paradigm-Shift Exploration			
Cross-domain ideation systems	  SciMON (Wang et al., 2024b)   SciMuse (Gu and Krenn, 2025)   VIRSCI (Su et al., 2025)   ResearchAgent (Baek et al., 2025)   CodeScientist (Jansen et al., 2025)   Scideator (Radensky et al., 2026)	α : literature retrieval, facet extraction, cross-domain recombination, novelty checking γ : research problem-concept-facet-domain alignment ω : inspirations, provenance, novelty signals, feedback, revision paths	Weakly grounded ideas; misleading analogies; vague research directions
Falsification & feedback systems	  Baby-AIGS (Liu et al., 2024b)   InternAgent (Zhang et al., 2025a)   Kosmos (Mitchener et al., 2025)   Robin (Ghareeb et al., 2026)   AI Co-Scientist (Gottweis et al., 2026)   CellVoyager (Alber et al., 2026)	α : hypothesis generation, critique, validation, data analysis, revision loops γ : objective-evidence-data/code-claim alignment ω : falsification/review agents, rankings, experiment results, citations, feedback	Amplified weak hypotheses; proxy overfitting; unsupported long-horizon claims
Recursive AI scientist workflows	  The AI Scientist-v2 (Yamada et al., 2025)   AI-Researcher (Tang et al., 2025)   DeepScientist (Weng et al., 2026)   The AI Scientist (Lu et al., 2026)	α : ideation, code, experimentation, writing, result analysis, recursive refinement γ : research goal-hypothesis-result-manuscript alignment ω : experiment logs, errors, plots, reviews, manuscript drafts, feedback	Ungrounded ideas; recursive errors; result misinterpretation; hidden methodological gaps
Verification & search systems	  FunSearch (Romera-Paredes et al., 2024)   LLM-SR (Shojaee et al., 2025)   AlphaEvolve (Novikov et al., 2025)	α : program/equation search and evaluation γ : problem spec-evaluator-code alignment ω : scores, evaluator outputs, program databases, best candidates	Proxy-score overfitting; high-scoring but weakly interpretable outputs
Infrastructure & agent ecosystems	  ToolUniverse (Gao et al., 2025)   OmniScientist (Shao et al., 2025)	α : tool, protocol, and provenance workflows γ : tool-data-knowledge-graph-norm alignment ω : traces, outputs, provenance, review signals, rankings	Hidden tool misuse; provenance gaps; weak governance; misaligned collaboration

Table 2: Full system-level coding of representative human-agent scientific collaboration systems. Systems are grouped by epistemic mode and workflow cluster. Icons indicate artifact type:  model,  framework,  single-agent or agentic system,  multi-agent system,  benchmark,  method,  interface, and  infrastructure. The coordination-demand column summarizes articulation work (α), common ground (γ), and awareness (ω).

Survey Theme	Survey Papers	Scientific Epistemic Context	Human Role Framing	Explicit Failure Modes
AI4Science	Luo et al. (2025b)	✓ Addressed	~ Addressed for peer review/writing; no human-role taxonomy	✗ Absent
	Chen et al. (2025a)	✗ Not addressed; focuses on tasks and workflows	✓ Addressed	✗ Absent
	Wei et al. (2025)	✓ Addressed	~ Discusses human roles in relation to agents; no human-role taxonomy	✗ Absent
	Zheng et al. (2025)	✓ Addressed	✓ Addressed	✗ Absent
	Tie et al. (2026)	✓ Addressed	~ Discusses human roles in relation to agents; no human-role taxonomy	✗ Absent
Agentic System	Wang et al. (2024a)	~ Not central; general application focus	✗ Not addressed in scientific epistemic contexts	✗ Absent
	Guo et al. (2024)	~ Not central; general application focus	~ Weakly addressed; no human-role taxonomy	✗ Absent
	Tran et al. (2025)	✗ Not addressed	~ Weakly addressed; no human-role taxonomy	✗ Absent
	Xu et al. (2025)	✗ Not addressed	~ Weakly addressed; no human-role taxonomy	✗ Absent
	Jiang et al. (2026)	✓ Addressed	✗ Not addressed	✗ Absent
Human-Agent Collaboration	Zahedi and Kambham- pati (2021)	✗ Not addressed	✓ Addressed	✗ Absent
	Vaccaro et al. (2024)	✗ Not addressed; general applica- tion focus	~ Weakly addressed; no human-role taxonomy	✗ Absent
	Gomez et al. (2024)	✗ Not addressed; healthcare appli- cation	✓ Broadly addressed in decision-making settings	✗ Absent
	Vats et al. (2025)	✗ Not addressed; healthcare appli- cation	✓ Addressed	~ Not explicit
	Zhang et al. (2025b)	✗ Not addressed	✓ Addressed	✗ Absent
	Shi et al. (2026)	✓ Addressed	✓ Addressed	~ Not explicit
	Zou et al. (2026)	✗ Not addressed	✓ Addressed	~ Not explicit
This Survey	Ours	✓ Addressed	✓ Addressed	✓ Explicit

Table 3: Comparison of existing survey traditions. The table contrasts prior surveys and this work by whether they explicitly address scientific epistemic context, human-role framing, and failure-mode analysis. These dimensions highlight the analytical gap addressed by CSCW×Kuhn.

Conceptual Block	Term	Operational Definition
Human-Agent Scientific Collaboration		
Collaboration scope	Human-Agent Collaboration	A collaborative arrangement in which humans and agentic systems jointly contribute to a scientific goal. We exclude simple retrieval tools and static classifiers, and include agents that participate in planning, judgment, execution, or revision within scientific workflows.
	Agentic system	A system that integrates one or more AI agents capable of planning, reasoning, tool use, memory, workflow execution, or autonomous task decomposition.
Coordination structure	Human-agent coordination	The organization of interdependent work between humans and agents, including task allocation, state tracking, feedback, and intervention.
	Agent-agent coordination	The organization of interdependent work among agents in multi-agent systems, including role assignment, communication, delegation, and aggregation.
Coordination Dimensions		
Coordination demand	Coordination demand	The type and intensity of coordination required by a scientific workflow under a given epistemic context.
	<i>Articulation work (α)</i>	The work required to organize, assign, sequence, integrate, and revise interdependent tasks so that the workflow remains coherent.
	<i>Common ground (γ)</i>	The shared understanding of goals, assumptions, constraints, evidence, methods, task state, and responsibilities among collaborators.
	<i>Awareness (ω)</i>	The visibility of actions, intermediate outputs, decisions, errors, provenance, and consequences during collaboration.
Scientific Context		
Scientific activity	Scientific innovation	The advancement of scientific knowledge through research, discovery, validation, and refinement of methods or explanations.
	<i>Scientific research</i>	Systematic inquiry that formulates questions, applies methods, analyzes evidence, and produces scientific claims, models, tools, or artifacts.
Epistemic condition	<i>Scientific discovery</i>	The production of new, externally verifiable scientific knowledge or explanations.
	Epistemic context	The scientific conditions that shape collaboration, including uncertainty, assumptions, goals, evidential standards, and validation criteria.
	Anomaly	A finding, observation, or claim that does not fit prevailing assumptions and raises uncertainty about whether existing explanations remain sufficient.
Epistemic Modes		
	Epistemic mode	A form of scientific work characterized by distinct goals, uncertainty levels, evidential standards, and coordination demands.
	<i>Puzzle-solving in normal science</i>	Routine scientific work conducted within accepted paradigms, methods, assumptions, and evaluation criteria.
	<i>Anomaly evaluation</i>	The assessment of unexpected findings, uncertain claims, or contested contributions to determine their credibility, significance, and relation to existing assumptions.
	<i>Paradigm-shift exploration</i>	Exploratory work that searches for alternative concepts, frameworks, or research directions when existing assumptions may be insufficient.
Coordination Failure Modes		
	Failure mode	A recurring breakdown associated with misalignment between coordination structure and epistemic demands in human-agent scientific collaboration.
	<i>Procedural drift</i>	Deviation from the intended workflow, protocol, constraint set, or task state.
	<i>Grounding substitution</i>	Replacement of evidential support or domain validity with plausibility, fluency, or apparent novelty.
	<i>Evaluative flattening</i>	Reduction of contested scientific judgment into simplified scores, rankings, criteria, or majority opinions.
	<i>Proxy optimization</i>	Optimization of benchmark metrics or proxy objectives while losing alignment with the underlying scientific goal.
	<i>Epistemic black-boxing</i>	Limited inspectability into provenance, assumptions, rationale, intermediate states, tool use, or responsibility.

Table 4: Operational definitions of key concepts used in this survey. Terms are grouped by conceptual function to clarify how CSCW×Kuhn connects human-agent collaboration, coordination demand, epistemic context, and coordination failure.

Research Work	Science Domain	Scope	Target Task	Metric / Approach	Coverage Gap			
		Seg. Task			C	O	J	V
Execution and Automated Workflow								
MLE-Bench [B] (Chan et al., 2025)	AI / ML	✓	Train models and submit prediction files for Kaggle competitions	Any-medal rate; medal rates; above-median rate; valid submission rate; pass@k; raw competition score				✗
DiscoveryBench [B] (Majumder et al., 2025)	General	✓	Search and verify hypotheses from datasets and discovery goals	Hypothesis Matching Score; context F1; variable F1; relationship accuracy	✗	✗		
ScienceAgentBench [B] (Chen et al., 2025d)	General	✓	Generate self-contained Python code to solve tasks from published papers	Valid execution rate; success rate; CodeBERTScore; API cost; rubric-based human rating				✗
Auto-Bench [B] (Chen et al., 2025b)	General	✓	Infer causal graphs through iterative oracle-guided interventions	Success rate; average iterations; adjacency-matrix match; trajectory accuracy			✗	
DSBench [B] (Jing et al., 2025)	Data Science	✓	Answer data-analysis questions and build predictive models	Task-level accuracy; competition-level accuracy; task success rate; performance gap		✗	✗	
AAAR-1.0 [B] (Lou et al., 2025)	General	✓	Validate equations, design experiments, detect paper weaknesses, identify human-review deficiencies	F1; entailment precision/recall/F1; semantic similarity; ROUGE/BERTScore; TF-IDF		✗		
RE-Bench [B] (Wijk et al., 2025)	AI / ML	✓	Solve open-ended ML research-engineering environments	Normalized score; score@k; time budget; human-expert comparison			✗	
Gravity-Bench-v1 [B] (Koblischke et al., 2025)	Physics	✓	Infer gravitational dynamics from budgeted observations in simulated binary-star systems	Task success rate; percent-error threshold; expert-reference comparison; observations used		✗		
MLAlgo-Bench [B] (Wang et al., 2025a)	AI / ML	✓	Implement ML algorithms and workflows from instructions	Pass rate; performance; time overhead; instruction-following score				✗
PaperBench [B] (Starace et al., 2025)	AI / ML	✓	Replicate AI research papers	Replication Score; rubric pass/fail scores; JudgeEval F1; human baseline comparison			✗	✗
MLGym-Bench [B] (Nathani et al., 2025)	AI / ML	✓	Improve baselines on open-ended AI research tasks	AUP score; performance profiles; Best Attempt@4; Best Submission@4; raw task metrics			✗	✗
SciReplicate-Bench [B] (Xiang et al., 2025)	NLP	✓	Generate code from algorithm descriptions in recent NLP papers	Execution accuracy; CodeBLEU; reasoning graph accuracy; dependency/API recall				✗
Scholarly Judgment and Evidential Support								
AgentReview [E] (Jin et al., 2024)	AI / ML	✓	Simulate peer-review process	Decision agreement; Jaccard index; Cohen's κ ; rating mean/variance; review-text length; BERTScore/embedding similarity				✗
TruthHypo [B] / KnowHD [E] (Xiong et al., 2025)	Biomed.	✓	Generate and verify biomedical entity-relation hypotheses	Link-level precision/recall/F1; relation accuracy; groundedness score; human/GPT-4o truthfulness preference			✗	✗
SchNovel [B] / RAG-Novelt [E] (Lin et al., 2025)	General	✓	Compare paired papers to identify the more novel work	Pairwise accuracy; field-level accuracy; year-gap accuracy; RAG-Novelt accuracy				✗
ChemRAG-Bench [B] (Zhong et al., 2025)	Chemistry	✓	Evaluate RAG on chemistry question-answering, calculation, molecule, and reaction tasks	Accuracy; exact match; validity; molecular fingerprint similarity; BLEU/ROUGE; error accuracy				✗
MLRC-Bench [B] (Zhang et al., 2025d)	AI / ML	✓	Propose and implement novel ML methods for research competitions	Improvement to human solution; effectiveness; efficiency; code simplicity; pass@k				✗
Wu et al. (2025) [E] (Wu et al., 2025)	AI / ML	✓	Predict method novelty from novelty sentences in reviews and LLM-generated method summaries	Accuracy; weighted F1				✗
CLAIMCHECK [B] (Ou et al., 2025)	AI / ML	✓	Analyze and verify review weaknesses in disputed paper claims	Claim-association F1; weakness-label agreement; specificity/groundedness preference; weakness similarity		✗	✗	
ScholarQABench [B] / OpenScholar [E] (Asai et al., 2026)	General	✓	Retrieve relevant literature and synthesize citation-backed answers to scientific queries	Rubric accuracy; citation F1; correctness accuracy; ROUGE-L; LLM-judge quality scores; expert preference		✗		✗
PaperWritingBench [B] (Song et al., 2026a)	AI / ML	✓	Generate LaTeX manuscripts, synthesize targeted related work, construct citations and visuals	Citation F1/recall; literature-review/manuscript quality; simulated acceptance; human pairwise preference		✗		✗
Exploratory Ideation and Research-Direction Value								
MOOSE [B/E] (Yang et al., 2024)	Social Science	✓	Generate hypotheses from raw web corpus	Validity score; novelty score; helpfulness score; expert ratings; feedback-iteration gains				✗
Si et al. (2025) [E] (Si et al., 2025)	NLP	✓	Compare novelty in LLM-generated versus expert-written research ideas through blind expert review	Novelty score; excitement score; feasibility score; expected-effectiveness score; overall score; reviewer agreement				✗
MOOSE-Chem [B/E] (Yang et al., 2025)	Chemistry	✓	Rediscover chemistry hypotheses by inspiration retrieval, composition, and ranking	Hit ratio; Matched Score; average rank ratio; Top MS; Average MS		✗		✗
AI Idea Bench 2025 [B] (Qiu et al., 2025)	AI / ML	✓	Generate research ideas using target and inspiration papers	Idea-to-topic score; idea-to-idea score; multiple-choice success; idea-competition score; novelty score; feasibility score; feasibility-per-step score			✗	✗
IdeaBench [B] (Guo et al., 2025)	General	✓	Generate research ideas from target-paper reference abstracts	Insight Score; semantic similarity; idea-overlap rating; human-ranking alignment				✗
LiveIdeaBench [B] (Ruan et al., 2026)	General	✓	Generate scientific ideas from single-keyword prompts	Originality; feasibility; fluency; flexibility; clarity; overall score			✗	✗
ResearchBench [B] (Liu et al., 2026)	General	✓	Retrieve inspirations, compose hypotheses, and rank hypotheses	Hit ratio; normalized hypothesis-composition score; pairwise ranking accuracy		✗		✗
SDE [B/E] (Song et al., 2026b)	General	✓	Answer scenario-grounded questions; generate, test, and refine hypotheses in discovery loops	Question-level accuracy; normalized project-level subjective scores			✗	✗

Table 5: Benchmark and evaluation studies for human-agent scientific collaboration. [B] denotes benchmark and [E] denotes evaluation study. Seg. denotes workflow-segment evaluation, and Task denotes task-specific evaluation. Coverage gaps are marked for workflow coordination (C), human oversight (O), scientific judgment (J), and downstream validation (V).