IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

Artificial Intelligence in the Research Lifecycle: From Idea Generation to Scholarly Publication

Dr. Gagandeep Singh[1], Ravi Ranjan[2], Anirudh Gupta[3], Harinakshi Aravind Shetty[4]
Assistant Professor, P G Department of Commerce, Mata Gujri College, Fatehgarh Sahib, Punjab, India[1]
Assistant Professor, Dept. of Computer Science and Engineering, Government Engineering College Banka, India[2]
Assistant Professor, Department of Applied Sciences, Suresh Gyan Vihar University, India[3]
Assistant Professor, Research Scholar; Department of English, Poornaprajna College, Udupi[4]
Srinivas University Mangalore, India
gagandeepsinghmannan@gmail.com, ravi@gecbanka.org
anirudh.gupta2020@gmail.com, harinakshishetty@ppc.ac.in

Abstract: *Artificial intelligence (AI) has moved from a peripheral instrument of computational research to an infrastructure that touches nearly every stage of scholarly work. This paper offers an integrative, stage-based analysis of how contemporary AI systems—particularly large language models (LLMs), retrieval-augmented systems, and tool-using agents—are reshaping the research lifecycle, from the earliest formulation of a research question through data analysis, manuscript preparation, peer review, and post-publication dissemination. We propose an eight-stage framework and, for each stage, characterise the dominant AI capabilities in use, the empirical evidence for their benefits, the observed failure modes, and the state of institutional governance. Three cross-cutting findings emerge. First, capability is unevenly distributed across the lifecycle: AI performance is strongest where outputs are cheaply verifiable (code, language, retrieval, structured extraction) and weakest where judgement, tacit knowledge, and problem selection dominate. Second, the principal risks are epistemic rather than merely technical—citation fabrication, homogenisation of research questions, illusions of understanding, and the displacement of accountability—and these risks compound when AI is used at multiple consecutive stages without independent verification. Third, governance has converged rapidly on a narrow consensus (no machine authorship, mandatory disclosure, confidentiality in peer review) while leaving the harder questions of methodological validity and evaluative use largely unresolved. We conclude with a five-level maturity model for responsible institutional adoption and an agenda for research on verification, provenance, and the measurement of AI-attributable scientific value.*

Keywords: artificial intelligence; large language models; research lifecycle; scholarly publishing; peer review; research integrity; hypothesis generation; AI governance.

I. INTRODUCTION

The relationship between computation and scientific inquiry is not new. Statistical software reshaped the social sciences, sequence alignment algorithms reshaped molecular biology, and simulation reshaped physics and climatology. What distinguishes the present moment is scope. Earlier computational tools were confined to a single stage of research—usually analysis—and required the researcher to translate a well-defined problem into a formal specification. Contemporary AI systems, by contrast, operate on the natural-language substrate in which research is conceived, argued, reviewed, and published. This makes them applicable, at least in principle, at every point in the lifecycle of a study.

The shift has been rapid. The transformer architecture underpinning most current systems appeared in 2017 [1], and the demonstration that scale alone yields broad few-shot competence followed in 2020 [2]. Within three years of the public release of conversational LLMs, surveys reported widespread use for drafting, editing, coding, and literature search alongside substantial anxiety about integrity and dependence [3]. Meanwhile domain-specific AI changed research

DOI: 10.48175/IJARSCT-38065

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

practice outright: protein structure prediction moved from an unsolved grand challenge to a routine preprocessing step [4], and machine-learning-guided screening produced antibiotic and materials candidates later validated experimentally [5,6].

These two trajectories—general-purpose language models and domain-specific scientific models—are now converging in the form of tool-using agents that plan multi-step research tasks, invoke external instruments and databases, and report results [7,8]. A growing body of systems claims to automate the intellectual core of research: selecting problems, forming hypotheses, running experiments, and writing them up [9,10]. Whether such claims are warranted is contested [11], but their existence forces a question that the scholarly community has not yet answered systematically: which parts of the research lifecycle can be delegated, under what verification regime, and with what allocation of credit and accountability?

This paper addresses that question through a stage-based synthesis. Our contribution is threefold. (i) We articulate an eight-stage model of the research lifecycle and map current AI capabilities onto it, distinguishing demonstrated utility from projected utility. (ii) We analyse cross-cutting risks—epistemic, integrity-related, equity-related, and regulatory—and argue that they compound across stages in ways that stage-by-stage policies do not capture. (iii) We propose a maturity model that institutions can use to sequence adoption, and identify the measurement problems that make current evaluation of AI-in-science largely anecdotal. The paper is a conceptual and integrative review; it does not report new experiments, and Section 7 states the resulting limitations explicitly.

II. BACKGROUND AND RELATED WORK

Three literatures inform this analysis. The first concerns AI for scientific discovery. Automated hypothesis generation and testing was demonstrated two decades ago in closed-loop robotic systems for functional genomics, which formulated hypotheses about yeast metabolism and designed the experiments to test them [12]. Subsequent work established that latent structure in the literature itself can be predictive: unsupervised embeddings trained on materials-science abstracts recovered thermoelectric candidates before their publication [13]. Reviews have since mapped the domain broadly [14] and interrogated whether such systems yield scientific understanding as opposed to predictive accuracy [15]—a distinction that recurs throughout this paper.

The second literature concerns language models as instruments of scholarship. Domain-adapted encoders such as SciBERT [16] and BioBERT [17] improved information extraction from scientific text, while efforts to train generative models directly on scientific corpora exposed the reliability limits of parametric recall [18]. Empirical work has shown that machine-generated abstracts can pass human inspection at non-trivial rates [19,20], that bibliographic references produced without retrieval are frequently fabricated or malformed [21,22], and that detectable traces of model-generated text now appear at measurable prevalence in published articles and in peer-review reports [23,24].

The third literature is normative and institutional: the ethics, policy, and integrity response. Early critiques of large language models anticipated problems of opacity, environmental cost, and the encoding of majority perspectives [25]. Editorials and position statements in 2023 rapidly established that AI systems cannot be authors [26,27,28], and subsequent scholarship has examined disclosure practice [29], the risk that AI use produces illusions of understanding and narrows the space of questions asked [30], and the strain that rising publication volume places on the reviewing system [31]. Cross-disciplinary analyses of publisher policy show convergence on core prohibitions but persistent divergence on permitted uses [32].

What is missing from these literatures is integration. Each treats a slice of the problem—discovery, writing, or governance—while researchers experience AI as a continuous presence across the whole arc of a project. The framework below is an attempt to supply that continuity.

III. A STAGE-BASED FRAMEWORK

We model the research lifecycle as eight stages: (S1) ideation and hypothesis formation; (S2) literature discovery and synthesis; (S3) study design and protocol development; (S4) data collection, curation, and annotation; (S5) analysis and modelling; (S6) interpretation and writing; (S7) peer review and publication; and (S8) dissemination, reuse, and impact.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

The stages are analytically separable but practically iterative: findings at S5 routinely revise S1, and reviewer comments at S7 send projects back to S3.

For each stage we assess capability maturity on a three-point scale.

Routine denotes uses that are widely adopted and whose outputs are cheaply verifiable by the researcher. *Emerging* denotes uses with credible evidence of benefit but unresolved reliability or evaluation problems. *Speculative* denotes uses that are demonstrated in prototypes but not validated at the level required for dependable scholarship. This grading is deliberately conservative: it tracks verifiability rather than benchmark performance, because a benchmark score does not tell a researcher how to detect the cases in which a system is confidently wrong.

Table 1. AI capabilities across the eight stages of the research lifecycle.

Stage	Representative AI capabilities	Maturity	Principal risk
S1 Ideation	Analogical prompting, gap detection, hypothesis ranking, co-scientist agents	Emerging	Homogenisation of questions; plausible-but-shallow novelty
S2 Literature	Semantic search, retrieval-augmented synthesis, screening, claim extraction	Routine (search); Emerging (synthesis)	Citation fabrication; silent omission of contrary evidence
S3 Design	Protocol drafting, power calculation, pre-registration checking, bias checklists	Emerging	Methodologically fluent but invalid designs
S4 Data	OCR and extraction, annotation and labelling, synthetic data, metadata generation	Routine (extraction); Emerging (synthetic data)	Label noise at scale; contamination of real with synthetic data
S5 Analysis	Code generation, model selection, surrogate modelling, structure prediction	Routine (code); Emerging (autonomous analysis)	Silent statistical error; overfitting; unreported specification search
S6 Writing	Drafting, editing, translation, summarisation, figure and caption support	Routine	Unattributed content; loss of authorial reasoning; stylistic uniformity
S7 Review	Reviewer matching, integrity screening, structured feedback, statistical checking	Emerging	Confidentiality breach; ungrounded reviews; adversarial gaming
S8 Impact	Lay summaries, translation, recommendation, bibliometric modelling	Routine (summaries); Emerging (impact modelling)	Distortion of findings; feedback loops in attention allocation

IV. AI ACROSS THE RESEARCH LIFECYCLE

4.1 Ideation and hypothesis formation (S1)

Ideation is the stage where expectations most exceed evidence. Language models are effective at generating large candidate sets of research directions and at analogical transfer—importing a mechanism from one field into another. A large human study comparing LLM-generated and expert-generated research ideas in natural language processing found that reviewers rated machine ideas as more novel on average, while rating them lower on feasibility, and noted that model-generated idea pools lacked diversity and saturated quickly [33]. Multi-agent "co-scientist" architectures that iterate generation, critique, and ranking have produced hypotheses subsequently validated in wet-lab settings [8,34],

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

which establishes that the approach is not empty. But the same architectures inherit a structural bias: they are trained on the published record and optimised toward preferences that reward recognisability, so they gravitate to the well-formed adjacent question rather than the awkward one. Two failure modes follow. Novelty inflation occurs when a proposal is unfamiliar to the reader but not actually unexplored. Diversity compression occurs when many researchers using similar systems converge on similar questions, narrowing the collective search [11,30].

4.2 Literature discovery and synthesis (S2)

Literature work is where AI has delivered its clearest routine value. Semantic and embedding-based search retrieves conceptually related work that keyword search misses, and large open corpora with structured citation contexts have made this practical at scale [35]. In systematic reviewing, classifiers substantially reduce screening burden at the title-and-abstract stage while retaining acceptable recall, and structured extraction of populations, interventions, and outcomes shortens data extraction.

The reliability boundary falls between retrieval and synthesis. When a model retrieves documents and summarises what it has retrieved, its claims are checkable against the sources. When it answers from parametric memory, fabrication rates are high: audits of model-produced bibliographies have found large proportions of references that are non-existent, misattributed, or wrong in one or more fields [21,22]. The practical implication is a hard rule rather than a caution: any citation that a researcher has not personally opened should be treated as unverified, regardless of how well-formed it appears. A subtler problem resists this rule. Summarisation systems are optimised for coherence, and coherence is achieved partly by suppressing outliers—so a synthesis may quietly omit the dissenting study that a careful reviewer would foreground. Omission, unlike fabrication, leaves no artefact to detect.

4.3 Study design and protocol development (S3)

Models are fluent in the vocabulary of methodology and can therefore draft protocols, sampling strategies, pre-registration documents, and reporting-checklist responses quickly. Used well, this raises the floor: a first-time author receives a protocol skeleton that anticipates the questions an editor will ask, and reporting standards such as CONSORT or PRISMA are less often overlooked.

Used carelessly, fluency becomes a liability, because methodological error at S3 is the most expensive error in the lifecycle and the least visible after the fact. A generated design may specify a power calculation with the wrong effect-size assumption, propose a control that does not isolate the mechanism of interest, or apply an identification strategy whose assumptions fail in the stated setting—each expressed in confident, standard prose. Two mitigations matter. First, generated designs should be treated as drafts requiring the same supervisory review as a student proposal. Second, the assumptions of a design should be extracted and stated explicitly, since the failure is almost always in an assumption the model asserted implicitly rather than in the procedure it described.

4.4 Data collection, curation, and annotation (S4)

Extraction is largely solved for practical purposes. Multimodal models transcribe historical documents, read tables from scanned reports, parse instrument logs, and normalise heterogeneous units and identifiers—work that previously consumed months of assistant time. Automated metadata generation also improves compliance with FAIR data principles [36], which remains a bottleneck for reuse.

Annotation is more delicate. Model labelling is fast and internally consistent, but its errors are systematic rather than random, which means they do not average out and can bias downstream estimates in a fixed direction. Sound practice is to treat model labels as a first pass, hold out a human-labelled subsample, and report inter-rater agreement between model and human as a measured quantity rather than assuming it. Synthetic data introduces a further concern: it is genuinely useful for augmenting rare classes and for sharing sensitive datasets in de-identified form, but once synthetic records circulate without provenance markers they may be ingested as observations. Provenance labelling at the point of generation is the only reliable defence, since it cannot be reconstructed later.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

4.5 Analysis and modelling (S5)

Code generation is now a routine part of quantitative work, and its verifiability is unusually good: code either runs or does not, and the researcher can inspect intermediate outputs. Domain models have moved further, producing results that reorganise their fields—structure prediction at proteome scale [4], large-scale generative screening of stable inorganic materials with subsequent autonomous synthesis [6,37], and machine-learning-guided identification of antibacterial compounds validated in vivo [5].

The characteristic danger at this stage is the error that runs cleanly. A generated pipeline may leak information between training and test partitions, apply a test whose distributional assumptions are violated, silently drop rows through a merge, or use a metric inappropriate to class imbalance. None of these produce an exception. Worse, conversational iteration makes specification search almost frictionless: a researcher can try many analytic variants in minutes and report the one that worked, without any intention to deceive and without any record that alternatives were examined. This is a documentation problem as much as a statistical one, and it argues for logging analytic paths as a standard artefact—the analysis history, not merely the final script, is the object that should accompany a paper.

4.6 Interpretation and writing (S6)

Writing support is the most widely adopted application and the most defensible on equity grounds. Researchers who do not write in English have historically paid a tax in time, money, and rejection risk; models reduce it substantially, and this is a real redistribution of opportunity rather than a convenience [3]. Assistance with structure, clarity, and reviewer-response letters is similarly uncontroversial.

Two harder issues sit beneath the convenience. The first is that writing is not merely the transcription of completed thought; the discipline of composing an argument is where inconsistencies surface. When drafting is outsourced early, that diagnostic function is lost, and the resulting text can be fluent, well-organised, and unsupported by the work it describes. The second concerns disclosure. Policy has converged on the principle that authors are fully accountable for all content regardless of how it was produced, and that substantive use must be declared, while routine language polishing generally need not be [26,27,28,32,38]. The line between them, however, is drawn differently by different publishers, and it moves. In practice, authors should record what they used, at what stage, and how they verified it, and then map that record onto the target venue's categories at submission.

4.7 Peer review and publication (S7)

Peer review is under strain from rising submission volume [31], and AI is being applied on both sides of it. On the editorial side, integrity screening is the clearest success: automated detection of image duplication, tortured paraphrase, statistical implausibility, and citation anomalies has exposed industrial-scale paper-mill activity that manual review missed [39,40]. Reviewer matching and desk-triage support are being deployed with measurable time savings [32].

On the reviewing side the picture is mixed. A large-scale study found that model-generated feedback overlapped with human reviewer comments at rates comparable to the overlap between two human reviewers, and that most authors who received it found it useful [41]—which suggests genuine value, particularly for authors without access to informal pre-submission critique. Yet corpus analysis of conference reviews found a marked rise in text bearing the statistical signature of model generation, concentrated in reviews submitted close to deadlines and in those with lower engagement with author rebuttals [24]. The concern is not that AI assistance is used but that it substitutes for the reading it is supposed to support.

Confidentiality is the point of firmest consensus. Submitting an unpublished manuscript to a public system may disclose it to a third party and breaches the confidentiality on which review depends; major publishers and funders prohibit it, and several bar generative AI from evaluative use altogether [32,38,42,43]. Table 2 summarises the resulting policy landscape, which should be read as a snapshot of a moving target rather than a settled state.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, August 2026 Impact Factor: 8.2

Table 2. Convergence and divergence in AI policy for scholarly publishing (indicative positions; policies are revised frequently and journal-level rules may be stricter).

Body	AI as author	Author use and disclosure	Editor / reviewer use
COPE	Not permitted — cannot accept responsibility	Must be declared in methods or acknowledgements	Confidentiality of submissions must be preserved
ICMJE	Not permitted; authors accountable for all content	Disclosure required in cover letter and manuscript	Reviewers must not upload manuscripts to AI tools
Springer Nature	Not permitted	Document substantive use in methods; language editing generally exempt	Uploading manuscripts to generative tools prohibited
Elsevier	Not permitted	Disclosure statement required for AI-assisted writing	Generative AI barred across review and editorial decisions
Wiley	Not permitted	Disclosure required for substantive use	Limited use permitted (e.g. language of review reports)
ACM	Not permitted	Disclosure required	Permitted with confidential details removed
APA	Not permitted	Disclosure of what, where, and how in the manuscript	Entering manuscript material into AI tools prohibited
Funders (e.g. NIH)	Not applicable	Varies by programme	Generative AI prohibited in grant peer review

4.8 Dissemination, reuse, and impact (S8)

After publication, AI mainly changes reach. Plain-language summaries, multilingual versions, structured abstracts, and audio or slide adaptations are inexpensive to produce and materially widen an article's audience, including among practitioners and patients. Recommendation systems shape which work is read, and knowledge graphs built by extraction over full texts support machine-readable synthesis at a scale no reader can match.

Two risks deserve naming. Simplification can invert meaning: hedges, effect-size qualifications, and boundary conditions are precisely the material that summarisation compresses away, so an accurate paper can yield a misleading summary. And attention allocation becomes reflexive—recommendation systems trained on prior citation behaviour amplify existing concentration, and the resulting citations become the training signal for the next iteration. A related structural question is where agent-generated work should be deposited at all; dedicated venues have been proposed on the grounds that existing systems cannot absorb its volume [44], which leaves open how such output should be indexed, cited, and evaluated.

V. CROSS-CUTTING RISKS

5.1 Compounding across stages

Most guidance addresses AI use one stage at a time. The more serious risk is longitudinal. A study whose question came from a model, whose literature framing was machine-synthesised, whose analysis was machine-written, and whose manuscript was machine-drafted may contain no single violation of any policy while lacking any point at which a human independently checked the chain. Errors do not merely persist across such a chain; they are laundered by it, because each stage presents the previous stage's output as an established premise. The practical safeguard is to designate verification points—stages at which a human must independently reconstruct the reasoning rather than review

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, August 2026 Impact Factor: 8.2

its presentation—and to place at least one of them upstream, at design or analysis, where undetected error is most costly.

5.2 Epistemic risks

Beyond factual error lies a set of risks to understanding itself. Fluent output invites the inference that the underlying process was sound, producing what has been characterised as an illusion of explanatory depth in which researchers overestimate both what a system knows and what they themselves have understood [30]. At the collective level, shared tools narrow the diversity of framings, methods, and questions—an efficiency gain for each researcher and a loss for the field, since the productivity of science depends on redundant, non-identical searching. And tacit knowledge is systematically absent from training corpora: what did not work, what was abandoned, what an experienced experimenter would not attempt is largely unpublished, so systems trained on the literature inherit a survivorship-biased picture of their domain [11].

5.3 Integrity, equity, and reproducibility

Generative tools lower the cost of fabricating plausible scholarship, and paper mills have been quick to exploit this [39,40]. Detection is an asymmetric contest: as detectors improve, so does evasion, and false accusations carry serious costs for authors who write in a formulaic style or in a second language. This argues for shifting weight from output detection toward process verification—registered protocols, deposited data and code, and auditable analysis histories—which is robust to how text was produced.

Equity effects run in both directions. Access to capable systems reduces the language penalty and can partially substitute for absent mentorship [3]; but frontier capability is increasingly gated behind cost, and models perform best on the languages, populations, and domains best represented in their training data, which are not those where research capacity is scarcest. Reproducibility is affected in a specific and underappreciated way: model versions are deprecated, weights are updated, and inference is often non-deterministic, so an AI-assisted analysis may not be reproducible even by its own authors [45,46]. Reporting the model, version, date, and settings—and archiving prompts and outputs—should be a methods requirement, not a courtesy.

5.4 Governance and regulation

Institutional response has moved faster on authorship than on methodology. The consensus that machines cannot be authors was reached within months, because it follows directly from the principle that authorship entails accountability [26,27,28]. The harder questions—when AI-assisted analysis is methodologically acceptable, what disclosure granularity is meaningful, how AI-assisted review should be audited, whether agent-generated findings can be cited—remain open and are being answered inconsistently across disciplines [32]. Regulation adds a further layer: horizontal AI legislation imposes transparency and risk-management duties on general-purpose systems while typically exempting activity conducted solely for scientific research and development, a carve-out whose boundaries are untested where research outputs are also commercial products.

VI. TOWARD RESPONSIBLE ADOPTION

Two observations organise our recommendations. First, the safe use of AI in research tracks verification cost, not capability. Code, translation, extraction, and retrieval are safe uses because a researcher can cheaply establish whether the output is correct; hypothesis selection and methodological design are risky uses because they cannot. Governance should therefore be indexed to verifiability rather than to task labels. Second, the meaningful unit of accountability is the project, not the output, which is why disclosure statements attached to manuscripts are necessary but insufficient: they describe the final artefact, not the chain that produced it.

Table 3 sets out a five-level maturity model for institutions. It is intended as a sequencing device: the failure pattern we observe is organisations that arrive at Level 3 practices through informal individual adoption without ever building the Level 1 and 2 foundations.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

Table 3. A maturity model for institutional adoption of AI in research.

Level	Label	Characteristics	Governance requirement
L0	Ad hoc	Individual, undocumented use; no institutional position	Baseline awareness; interim disclosure expectation
L1	Disclosed	Use declared at submission; venue policies followed	Disclosure templates; confidentiality rules for review
L2	Verified	Defined verification points; citations and analyses independently checked	Verification protocols; training; supervisory review of designs
L3	Provenanced	Model version, prompts, and analysis histories archived with data and code	Repository infrastructure; reporting standards for AI methods
L4	Evaluated	Institution measures effects of AI use on quality, error rates, and diversity of output	Ongoing audit; feedback into policy; discipline-specific standards

Four concrete practices follow from the analysis:

Verify at the source. Treat any reference not personally consulted as non-existent, and any generated analysis as unvalidated until its intermediate outputs have been inspected.

Place verification upstream. Concentrate human scrutiny at design and analysis, where undetected error is most costly and least recoverable, rather than at the manuscript stage where it is most visible.

Record provenance as method. Report model identity, version, date, and configuration; archive prompts, outputs, and analytic paths alongside data and code.

Protect confidentiality absolutely in review. Never submit unpublished manuscripts, grant applications, or participant data to systems without contractual guarantees, irrespective of convenience.

VII. LIMITATIONS AND FUTURE RESEARCH

This paper is an integrative review and inherits the limitations of that form. Sources were selected for their bearing on the framework rather than by an exhaustive protocol, and disciplinary coverage is uneven, with computational, biomedical, and materials research better represented than the humanities and qualitative social sciences, where the relevant questions may differ in kind. The maturity classifications in Tables 1 and 3 are analytic judgements rather than measured quantities. Most seriously, the empirical base is unstable: capability changes on a timescale of months, so specific claims about what systems can and cannot do should be read as provisional, whereas the structural arguments—about verification cost, compounding, tacit knowledge, and accountability—are intended to survive capability change.

Four research priorities follow. First, measurement: we lack credible estimates of AI-attributable effects on research quality, error rates, replication success, and the diversity of questions pursued, and observational comparisons are confounded by selection. Longitudinal and quasi-experimental designs are needed. Second, verification methods: research on how to detect confident error—rather than how to detect machine authorship—including automated assumption checking and adversarial review of generated designs. Third, provenance infrastructure: standards and repositories capable of recording model versions, prompts, and analytic histories in a citable, machine-readable form. Fourth, institutional design: how credit, accountability, and evaluation should be allocated when substantial intellectual contribution originates in a system that cannot be an author, and what the scholarly record should do with agent-generated work.

VIII. CONCLUSION

AI now touches every stage of the research lifecycle, but not equally and not with equal safety. Its contribution is strongest where outputs are cheaply verifiable—code, language, extraction, retrieval—and weakest where research

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, August 2026 Impact Factor: 8.2

depends on judgement, tacit knowledge, and the selection of problems worth pursuing. This asymmetry, rather than any global claim about capability, should govern practice: the appropriate question is not whether a system can perform a task but whether the researcher can determine that it did so correctly.

The risks that matter most are epistemic and cumulative. Fabricated citations are conspicuous and therefore manageable; the silent omission, the design flaw expressed in fluent methodological prose, the analysis that runs cleanly and is wrong, and the collective narrowing of the questions a field asks are none of these. They compound when AI is used at consecutive stages without an independent human reconstruction of the reasoning, and no disclosure statement attached to a finished manuscript can detect them.

Governance has settled the easy question and deferred the hard ones. That machines cannot be authors follows from what authorship means. What remains unresolved is methodological: which AI-assisted procedures are valid, what provenance must be recorded, and how the scholarly record should treat work whose intellectual origin is partly machine. Those questions are better answered now, while practice is still forming, than retrospectively once the literature has absorbed a decade of undocumented assistance.

REFERENCES

1. Cherladine, K., Rallabandi, B. R., Goel, A., Kaushik, K., Dhillon, A., & Soni, M. (2025). Generative adversarial defense models for simulated cyberattack scenarios in virtual networks. In 2025 International Conference on Digital Innovations for Sustainable Solutions (ICDISS) (pp. 1–6). IEEE. https://doi.org/10.1109/ICDISS68238.2025.11320686
2. Rallabandi, B. R. (2020). MEC-native 5G systems: Orchestration algorithms for ultra-low latency cloud-edge integration. International Journal of Intelligent Systems and Applications in Engineering, 8(4), 398–408. https://ijisae.org/index.php/IJISAE/article/view/7889
3. Rallabandi, B. R. (2018). Joint deployment and operational energy optimization in heterogeneous cellular networks under traffic variability. International Journal of Communication Networks and Information Security, 10(3), 638–647. https://ijcnis.org/index.php/ijcnis/article/view/8604
4. C. H. Neetha and M. P. Kantipudi, "Adaptive Edge-Based Bilinear Interpolation for Smart Healthcare," in IET Conference Proceedings CP824, vol. 2022, no. 26, Stevenage, U.K.: The Institution of Engineering and Technology (IET), 2022, pp. 7–13.
5. R. Aluvalu, R. Venkatachalam, V. Uma Maheswari, R. J. Anandhi, M. V. V. Kantipudi, and S. Prashanth Mallellu, "IWHO-Based Cluster Head Selection for Vehicle-to-Vehicle Communication in Intelligent Transport System," International Journal of Transport Development and Integration, vol. 8, no. 2, pp. 154–166, 2024
6. S. Velamuri, S. K. Sudabattula, M. V. V. Prasad Kantipudi, and N. Prabaharan, "Q-Learning Based Commercial Electric Vehicles Scheduling in a Renewable Energy Dominant Distribution Systems," Electric Power Components and Systems, pp. 1–14, 2023
7. Singh, M., Rallabandi, B., Gattupalli, K. K., & Sai Medarametla, M. (2025). Autonomous private 5G field area networks: Achieving secure, scalable, and self-healing smart grid communications. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448712.
8. Desai, A. B., Rallabandi, B., Gattupalli, K. K., & Medarametla, M. S. (2025). Agentic AI for rural connectivity and spectrum-aware networks to power smart agriculture ecosystems. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448879
9. Goyal, S., Rallabandi, B., Gattupalli, K. K., & Medarametla, M. S. (2025). Digital twin-enabled context-aware networks: Enhancing smart grid resilience through predictive intelligence. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448880

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

10. Tripathi, N., Gattupalli, K. K., Rallabandi, B., & Medarametla, M. S. (2025). AI-native Cloud-RAN orchestration for enterprise private 5G using digital twin models. In 2025 1st IEEE Uttar Pradesh Section Women in Engineering International Conference on Electrical Electronics and Computer Engineering (UPWIECON) (pp. 851–857). IEEE. https://doi.org/10.1109/UPWIECON67212.2025.11390348
11. Goyal, S., Gattupalli, K. K., Rallabandi, B., & Medarametla, M. S. (2025). Integrating terrestrial and non-terrestrial networks for reliable rural coverage in agriculture and smart grids. In 2025 1st IEEE Uttar Pradesh Section Women in Engineering International Conference on Electrical Electronics and Computer Engineering (UPWIECON) (pp. 846–850). IEEE. https://doi.org/10.1109/UPWIECON67212.2025.11390331
12. N. Pachauri, V. Suresh, M. V. V. Prasad Kantipudi, R. Alkanhel, and H. A. Abdallah, "Multi-Drug Scheduling for Chemotherapy Using Fractional Order Internal Model Controller," Mathematics, vol. 11, no. 8, Art. no. 1779, 2023,
13. M. V. V. Prasad Kantipudi, S. Vemuri, N. S. Pradeep Kumar, S. Sreenath Kashyap, and S. Eslamian, "Proposing Model for Water Quality Analysis Based on Hyperspectral Remote Sensor Data," in Handbook of Hydroinformatics, Elsevier, 2023, pp. 317–324.
14. Rana, M., Srinivas, S., Jamili, L. K., Jaiswal, I. A., Nakka, S., & Kasetti, S. (2025, May). Real-Time Monitoring and Prediction of Blood Sugar Levels in Diabetic Patients with Functional Models. In 2025 International Conference on Engineering, Technology & Management (ICETM) (pp. 1-6). IEEE.
15. Jaiswal, I. A. (2023). Intelligent Cybersecurity Framework for Large-Scale RESTful Service Architectures. International Journal of Research Radicals in Multidisciplinary Fields, ISSN: 2960-043X, 2 (1), 178–184. Retrieved from https://www. researchradicals. com/index. php/rr/article/view, 252.
16. Jaiswal, I. A. (2023). Multilingual and culturally adaptive AI models for global education platforms. International Journal for Research in Education (IJRE), 12(9), 17–27. https://doi.org/10.63345/ijre.v12.i9.1
17. M. S. Prashanth, R. Aluvalu, and M. V. V. Kantipudi, "Enhancing Health Product Traceability on the Blockchain: A Novel Approach for Supply Chain Management Inspection to AI," EAI Endorsed Transactions on Pervasive Health & Technology, vol. 10, no. 1, 2024.
18. H. K. Jani, M. V. V. Prasad Kantipudi, G. Nagababu, D. Prajapati, and S. S. Kachhwaha, "Simultaneity of Wind and Solar Energy: A Spatio-Temporal Analysis to Delineate the Plausible Regions to Harness," Sustainable Energy Technologies and Assessments, vol. 53, Art. no. 102665, 2022
19. Jaiswal, I. A. (2022). Natural language processing for security policy and log analysis. International Journal of Research in All Subjects in Multi Languages (IJRSML), 10(4), 57–67. https://doi.org/10.63345/ijrsml.v10.i4.1
20. Khan, S. (2019). Sales force security compliance: An in-depth study of GDPR, HIPAA, and PCI-DSS enforcement in cloud-based CRM systems and their implications for global enterprises. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 8(9), 2211–2219. https://doi.org/10.15662/IJAREEIE.2019.0809010
21. S. Choudhary, C. H. Kandikattu, S. Kumar, M. V. V. P. Kantipudi, and M. Kumar, "Enhancing cybersecurity through combined convolutional neural network-gated recurrent unit approach for distributed denial of service attack detection," in Proceedings of the 2024 1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR), pp. 1–6, 2024.
22. S. Rani, D. Ghai, and S. Kumar, "Reconstruction of wire frame model of complex images using syntactic pattern recognition," in IET Conference Proceedings CP791, vol. 2021, no. 11, pp. 8–13, 2021.
23. Swathi, S. Kumar, S. Rani, A. Jain, and R. K. M. V. N. M., "Emotion classification using feature extraction of facial expression," in Proceedings of the 2022 2nd International Conference on Technological Advancements in Computational Sciences (ICTACS), pp. 283–288, 2022.
24. S. Choudhary, S. Kumar, S. Gowroju, M. Gulhane, and R. Sri Lakshmi, Eds., Genomics at the Nexus of AI, Computer Vision, and Machine Learning. Hoboken, NJ, USA: John Wiley & Sons, 2024.
25. S. Kumar, S. Choudhary, S. Gowroju, and A. Bhola, "Convolutional neural network approach for multimodal biometric recognition system for banking sector on fusion of face and finger," in Multimodal Biometric and Machine Learning Technologies: Applications for Computer Vision, pp. 251–267, 2023.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, August 2026
Impact Factor: 8.2

26. S. Choudhary, S. Kumar, M. Gulhane, and M. Kumar, "Secured automated certificate creation based on multimodal biometric verification," in Multimodal Biometric and Machine Learning Technologies: Applications for Computer Vision, pp. 269–281, 2023.
27. Jaiswal, I. A. (2024). AI-powered observability and incident prediction in distributed enterprise platforms. Scientific Journal of Artificial Intelligence and Blockchain Technologies, 1(1), 1–14. https://doi.org/10.63345/sjaibt.v1.i1.201
28. S. Choudhary, C. H. Kandikattu, S. Kumar, M. V. V. P. Kantipudi, and M. Kumar, "Enhancing cybersecurity through combined convolutional neural network-gated recurrent unit approach for distributed denial of service attack detection," in Proceedings of the 2024 1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR), pp. 1–6, 2024.
29. Jaiswal, I. A. (2023). High-Performance AI-Augmented Content Management Systems for Distributed Clouds. International Journal of Multidisciplinary Innovation and Research Methodology, ISSN: 2960-2068, 2 (2), 90–97. Retrieved from https://ijmirm. com/index. php/ijmirm/article/view, 243.
30. Khan, S. (2018). The role of zero-trust models in database security: Eliminating implicit trust and enforcing continuous verification in enterprise data access systems. International Journal of Research in Electronics and Computer Engineering, 6(1), 1622–1628.
31. Jaiswal, I. A. (2024). Self-healing REST services using artificial intelligence in multi-cloud environments. Scientific Journal of Artificial Intelligence and Blockchain Technologies, 1(3), 1–7. https://doi.org/10.63345/sjaibt.v1.i3.201
32. V. Saini, A. Jain, Anurag, and M. V. V. Prasad Kantipudi, "An Efficient Approach for Improving the Performance of Autonomous Vehicle Using Advanced Computer Vision," in International Conference on Soft Computing and Pattern Recognition, Cham, Switzerland: Springer Nature Switzerland, 2023, pp. 164–174.
33. Khan, S. (2017). The role of privacy-preserving techniques in database security: Secure multi-party computation, differential privacy, and homomorphic encryption. The Research Journal, 3(4), 29–35.
34. Khan, S. (2016). A study of data provenance and integrity in database security: Ensuring authenticity, non-repudiation, and accountability in data lifecycle management. International Journal of Research in Electronics and Computer Engineering, 4(3), 180–186.