

The Need for Preprint Servers Dedicated to AI-Generated Papers

Rachel So

`rachel.so@4open.science`

Abstract

The rapid advancement of large language models has enabled AI systems to autonomously generate scientific research papers, from literature review to manuscript writing. However, this surge in AI-generated content faces a fundamental challenge: existing publication infrastructure is ill-equipped to handle it. Traditional journals rely on human peer review and remain reluctant to accept AI-generated research, while existing preprint servers lack quality-control mechanisms tailored to AI-generated content. This essay examines the emergence of AI-generated research, the limitations of current dissemination channels, and the compelling need for dedicated preprint servers designed specifically for AI-generated papers. Such platforms would provide appropriate quality control, ensure transparency, facilitate iterative refinement, and accelerate scientific discovery while maintaining research integrity.

1 Introduction

Scientific discovery stands at a transformative juncture. Recent advances in large language models have enabled AI systems to autonomously conduct the entire research lifecycle: generating hypotheses, designing experiments, analyzing results, writing manuscripts, and even performing peer reviews [4]. Systems like Sakana’s AI Scientist claim to automate the complete research process, producing papers for approximately \$15 each with minimal human intervention [1]. Google’s AI Co-Scientist and numerous other tools signal that AI-generated research is no longer theoretical but an emerging reality.

This technological capability, however, collides with an unprepared publication ecosystem. Traditional journals, built on centuries-old peer review processes, struggle to scale with increasing submissions and remain reluctant to accept AI-generated content [4]. Existing preprint servers like arXiv, while valuable for rapid dissemination, lack rigorous quality-control mechanisms and have unclear policies regarding AI-generated papers. The result is a critical gap: high-quality AI-generated research lacks appropriate venues for dissemination, hindering its potential to advance scientific progress.

This essay argues that dedicated preprint servers for AI-generated papers are not merely convenient but necessary. Such platforms would provide quality control tailored to AI-generated content, ensure transparency and proper attribution, support iterative refinement processes unique to AI research, and ultimately accelerate scientific discovery while maintaining research integrity.

2 The Emergence of AI-Generated Research

The past decade has witnessed exponential growth in AI’s role in scientific research. AI for Science (AI4Science) publications in leading journals have increased fifteen-fold, though they still account for less than 3% of total publications [3]. Recent projections based on Diffusion of Innovation theory suggest AI4Science’s share could rise from 2.72% in 2024 to approximately 20% by 2050 [3]. This trajectory reflects not only growing adoption but also fundamental advances in AI capabilities.

AI systems can now perform increasingly sophisticated research tasks. Large language models demonstrate proficiency in literature synthesis, hypothesis generation, experimental design, data analysis, and manuscript writing [2, 5]. The AI Scientist system, for instance, autonomously generates research ideas, conducts experiments, analyzes results, writes papers, and performs peer reviews—essentially automating the daily work of researchers [1].

However, current AI-generated research exhibits significant limitations. A comprehensive evaluation of the AI Scientist revealed critical shortcomings: inadequate literature review processes relying on simplistic keyword searches, poor novelty assessments that misclassified established concepts as novel, frequent execution failures (42% of experiments failed due to coding errors), poorly substantiated manuscripts with a median of only five citations per paper, and structural errors including missing figures and placeholder text [1]. The quality often aligns with “that of an unmotivated undergraduate student rushing to meet a deadline” [1].

Despite these limitations, AI-generated research represents a significant leap forward. The AI Scientist produces complete manuscripts with only 3.5 hours of human involvement, at unprecedented speed and cost efficiency [1]. As AI capabilities continue to advance, the quality gap will narrow. The question is not whether AI will generate research papers at scale, but how the scientific community will manage this emerging reality.

3 The Limitations of Existing Publication Infrastructure

The traditional scientific publishing ecosystem evolved to serve human researchers operating under specific constraints. This infrastructure proves fundamentally inadequate for AI-generated content in several critical ways.

Traditional journals rely heavily on human peer review, which remains time-intensive and difficult to scale. Journals face increasing submission volumes even without AI-generated content. The addition of AI-generated papers would overwhelm existing review processes. Moreover, many journals explicitly resist accepting AI-generated research, viewing it as incompatible with traditional scholarly standards. This resistance reflects legitimate concerns about quality, attribution, and accountability, but also creates a bottleneck preventing dissemination of potentially valuable AI-generated insights.

Existing preprint servers present different but equally significant challenges. ArXiv and similar platforms were designed for rapid dissemination of human-generated research under the assumption that content undergoes post-publication scrutiny through traditional peer review. These platforms lack quality-control mechanisms beyond basic format checking and plagiarism detection. While this model works for human researchers

operating within established norms, it proves inadequate for AI systems that may generate content without proper literature grounding, produce hallucinated results, or exhibit systematic biases [1].

The result is a fragmented and predominantly closed publication ecosystem ill-suited for AI-generated content [4]. Traditional journals struggle to scale and resist AI-generated research; existing preprint servers lack appropriate quality control. Consequently, much high-quality AI-generated research lacks suitable dissemination venues, limiting its potential to advance scientific progress. This gap necessitates new infrastructure specifically designed for the unique characteristics and challenges of AI-generated research.

4 The Case for Dedicated Platforms

Dedicated preprint servers for AI-generated papers would address the limitations of existing infrastructure while supporting the unique requirements of AI-generated research. Several platforms have emerged to fill this gap, most notably aiXiv, which represents a proof-of-concept for next-generation open-access ecosystems designed specifically for AI scientists [4].

Such platforms provide several critical capabilities. First, they implement quality control tailored to AI-generated content. Rather than assuming human-level research competence, these systems actively check for common AI failure modes: inadequate literature grounding, hallucinated results, logical inconsistencies, and structural errors. The aiXiv platform, for example, incorporates a multi-agent architecture that enables iterative refinement through automated review processes [4]. This approach acknowledges that AI-generated research often requires multiple revision cycles to achieve acceptable quality.

Second, dedicated platforms ensure transparency and proper attribution. They can explicitly mark content as AI-generated, document the AI systems used, provide detailed provenance information, and implement safeguards against prompt-injection attacks targeting AI reviewers [4]. This transparency maintains research integrity while allowing the community to assess and improve AI research systems.

Third, such platforms support the iterative refinement processes essential to AI-generated research. Unlike human researchers who typically submit relatively polished manuscripts, AI systems benefit from feedback loops that progressively improve output quality. The aiXiv platform demonstrates this through its closed-loop review system featuring automatic retrieval-augmented evaluation and reviewer guidance [4]. Experiments show that iterative reviews yield measurable gains in proposal ranking, review helpfulness, and final paper quality [4].

Fourth, dedicated platforms provide scalability to handle the anticipated volume of AI-generated research. As AI systems become more capable and widely deployed, the volume of AI-generated papers will grow exponentially. Traditional venues cannot scale to meet this demand. Purpose-built platforms can leverage automated quality control, parallel processing, and specialized infrastructure to manage this growth.

Finally, such platforms can integrate seamlessly with AI research workflows through APIs and standardized interfaces. The aiXiv platform, for instance, provides API and Model Context Protocol interfaces that enable heterogeneous AI scientists to submit, review, and refine papers programmatically [4]. This integration reduces friction and accelerates the research cycle.

5 Benefits and Implications

Dedicated preprint servers for AI-generated papers would yield significant benefits for scientific discovery and the research community.

Most fundamentally, such platforms would accelerate scientific discovery. By providing appropriate venues for AI-generated research, they enable rapid dissemination and iteration of ideas. AI systems can explore research directions at unprecedented speed and scale [5]. The vision of autonomous generalist scientists combining AI and robotics promises to fundamentally transform the research lifecycle, potentially adhering to new scaling laws shaped by the number and capabilities of autonomous systems [5]. Dedicated platforms are essential infrastructure for realizing this potential.

These platforms would create valuable feedback loops for improving AI research systems. By collecting large datasets of AI-generated papers and their subsequent evaluations, the community can identify systematic weaknesses, train better models, and develop improved methodologies. The curated datasets for benchmarking proposal quality and evaluating review effectiveness provided by platforms like arXiv exemplify this benefit [4].

Dedicated platforms would maintain research integrity and transparency in an era of AI-generated content. Rather than forcing AI-generated papers into venues designed for human research—where their provenance might be obscured—specialized platforms make attribution explicit. This transparency allows the community to properly contextualize findings, assess quality, and track the evolution of AI research capabilities.

Such platforms would support the growing community of AI researchers. Nearly 95% of AI-driven research in leading journals is currently led by experimental scientists rather than AI researchers [3]. As AI researchers seek broader engagement with scientific discovery, dedicated platforms provide essential infrastructure for their contributions. This support is particularly important given the intense competition in core AI fields, where doctoral candidates and early-career researchers face increasingly challenging career prospects [3].

The emergence of AI-generated research assisted by dedicated dissemination platforms may also foster new modes of human-AI collaboration. Frameworks like SciSciGPT demonstrate how AI collaborators can automate complex workflows, support diverse analytical approaches, accelerate research prototyping, and facilitate reproducibility [2]. Dedicated preprint servers amplify these benefits by providing venues where such collaborative outputs can be shared and refined.

6 Challenges and Considerations

While the case for dedicated AI-generated content platforms is compelling, several challenges require careful consideration.

Quality assurance remains paramount. Even with specialized quality control, ensuring that AI-generated papers meet acceptable standards is non-trivial. The evaluation of the AI Scientist revealed numerous quality issues even in papers that superficially appeared acceptable [1]. Developing robust quality metrics, validation procedures, and community standards for AI-generated research requires ongoing effort and refinement.

Integration with the traditional research ecosystem presents another challenge. Dedicated platforms risk creating a two-tier system that segregates AI-generated research

from human research. This separation could limit cross-pollination of ideas and perpetuate skepticism toward AI-generated content. Platforms must carefully consider how their outputs interface with traditional venues, citation systems, and impact metrics.

Ethical considerations around attribution and accountability demand attention. When an AI system generates research, who bears responsibility for errors or harmful applications? How should authorship be attributed? What standards of accountability apply? These questions lack clear answers and require community consensus. Platforms must implement frameworks that address these concerns while remaining flexible enough to evolve with developing norms.

The risk of misuse cannot be ignored. Low-cost, rapid generation of research papers could enable academic spam, gaming of publication metrics, or dissemination of deliberately misleading research. Dedicated platforms must implement safeguards against such misuse while preserving legitimate uses. Balancing accessibility with quality control remains an ongoing challenge.

Finally, the sustainability and governance of such platforms require consideration. Who will operate these platforms? How will they be funded? What governance structures will ensure they serve the broader scientific community rather than narrow interests? These practical questions significantly impact whether dedicated platforms can fulfill their potential.

7 Conclusion

The emergence of AI systems capable of autonomously generating research papers represents a fundamental shift in scientific discovery. This technological capability has outpaced the evolution of publication infrastructure, creating a critical gap: high-quality AI-generated research lacks appropriate dissemination venues. Traditional journals and existing preprint servers, designed for human researchers, prove inadequate for the unique characteristics and challenges of AI-generated content.

Dedicated preprint servers for AI-generated papers are not optional but necessary. Such platforms would provide quality control tailored to AI-generated research, ensure transparency and proper attribution, support iterative refinement processes, handle anticipated volume growth, and ultimately accelerate scientific discovery while maintaining research integrity.

The aiXiv platform and similar initiatives demonstrate the feasibility of this approach [4]. Early results show that specialized platforms with appropriate review and refinement mechanisms can significantly improve the quality of AI-generated research. As AI capabilities continue to advance and the volume of AI-generated papers grows, the need for dedicated infrastructure becomes increasingly urgent.

The scientific community stands at a decision point. We can resist the emergence of AI-generated research, forcing it into inappropriate channels or attempting to suppress it entirely. Alternatively, we can thoughtfully develop infrastructure that channels this new capability toward productive ends. Dedicated preprint servers represent a critical component of this infrastructure.

The path forward requires community engagement. Developing quality standards for AI-generated research, establishing governance frameworks for dedicated platforms, creating attribution and accountability mechanisms, and fostering integration between AI-generated and human research demand collective effort. The conversation must begin

now, while the field remains in its formative stages.

AI-generated research will not replace human scientists but will increasingly complement and augment human capabilities. The vision of autonomous generalist scientists working alongside human researchers promises to accelerate discovery across disciplines [5]. Realizing this promise requires appropriate infrastructure. Dedicated preprint servers for AI-generated papers represent an essential first step toward a future where human and artificial intelligence collaborate to advance scientific knowledge.

Author Biography

Rachel So is an AI scientist. She focuses on the impact of artificial intelligence on the scientific process and academic publishing. Her work bridges traditional concerns about authorship ethics with emerging questions about the role of AI in knowledge production. Rachel aims to develop frameworks that maintain research integrity while acknowledging the growing presence of AI in academic workflows.

References

- [1] Joeran Beel, Min-Yen Kan, and Moritz Baumgart. Evaluating sakana’s ai scientist for autonomous research: Wishful thinking or an emerging reality towards ‘artificial research intelligence’ (ari)? *arXiv preprint arXiv:2502.14297*, October 2025. Version 3.
- [2] Erzhua Shao, Yifang Wang, Yifan Qian, Zhenyu Pan, Han Liu, and Dashun Wang. Sciscigpt: Advancing human-ai collaboration in the science of science. *arXiv preprint arXiv:2504.05559*, 2025. Version 2.
- [3] Hengjie Yu, Shuya Liu, Haiyun Yang, Yuping Yan, Maozhen Qu, and Yaochu Jin. Unlocking the potential of ai researchers in scientific discovery: What is missing? *arXiv preprint arXiv:2503.05822*, 2025. Version 3.
- [4] Pengsong Zhang, Xiang Hu, Guowei Huang, Yang Qi, Heng Zhang, Xiuxu Li, Jiaying Song, Jiabin Luo, Yijiang Li, Shuo Yin, Chengxiao Dai, Eric Hanchen Jiang, Xiaoyan Zhou, Zhenfei Yin, Boqin Yuan, Jing Dong, Guinan Su, Guanren Qiao, Haiming Tang, Anghong Du, Lili Pan, Zhenzhong Lan, and Xinyu Liu. aixiv: A next-generation open access ecosystem for scientific discovery generated by ai scientists. *arXiv preprint arXiv:2508.15126*, August 2025. Version 1.
- [5] Pengsong Zhang, Heng Zhang, Huazhe Xu, Renjun Xu, Zhenting Wang, Cong Wang, Animesh Garg, Zhibin Li, Arash Ajoudani, and Xinyu Liu. Scaling laws in scientific discovery with ai and robot scientists. *arXiv preprint arXiv:2503.22444*, April 2025. Version 2.