Article

Ontological and Epistemological Challenges in Contemporary Multidisciplinary Research: Towards an AI-Enhanced Framework for Academic Knowledge Production and Evaluation

Richard Murdoch Montgomery [*]

Article

Ontological and Epistemological Challenges in Contemporary Multidisciplinary Research: Towards an AI-Enhanced Framework for Academic Knowledge Production and Evaluation

Richard Murdoch Montgomery

Scottish Science Society; editor@scottishsciencesocietyperiodic.uk

Abstract

Contemporary academic research increasingly confronts complex global challenges that transcend traditional disciplinary boundaries, yet the ontological and epistemological frameworks underpinning scholarly knowledge production remain rooted in outdated paradigms. This manuscript presents a theoretical framework for addressing these challenges through artificial intelligence-enhanced peer review systems, whilst acknowledging both the speculative nature of such proposals and legitimate concerns about algorithmic governance in academia. We position this work as a critical engagement with techno-solutionist approaches, learning from historical failures in academic automation whilst proposing safeguards against potential negative consequences. Through critical realist analysis, engagement with AI critics, and examination of economic and legal implications, we develop a nuanced vision that balances transformative potential with realistic assessment of risks. Our framework addresses power distribution, economic equity, regulatory requirements, and the preservation of epistemological diversity. We conclude that whilst AI-enhanced evaluation systems offer promising possibilities, their implementation requires unprecedented attention to governance, equity, and the protection of academic values.

Keywords: contemporary academic research; glogal challenges; ontologica and epitemological frameworks; transformative potential; epistemological diversity; academic values

1. Introduction

1.1. The Crisis of Traditional Academic Frameworks

The contemporary academic landscape finds itself at a critical juncture, wherein traditional structures of knowledge production and validation prove increasingly inadequate for addressing complex, interconnected challenges. However, we must approach proposed technological solutions with appropriate scepticism, learning from past failures and contemporary critiques. As O'Neil (2016) warns in "Weapons of Math Destruction," algorithmic systems can perpetuate and amplify existing biases whilst creating an illusion of objectivity. This manuscript therefore presents not an uncritical embrace of AI solutions but a nuanced framework that grapples with both possibilities and perils.

The historical development of academic disciplines, whilst providing valuable depth and specialisation, has inadvertently created epistemological silos that impede the holistic understanding necessary for tackling multifaceted problems such as climate change, global health crises, and sustainable development. This disciplinary fragmentation reflects what Boon and Van Baalen (2019) identify as the dominance of a "physics paradigm of science," which privileges reductionist approaches and assumes that knowledge can be objectively represented independent of its construction context.

The roots of this crisis can be traced to the nineteenth and early twentieth centuries, when the modern university system crystallised around distinct disciplinary boundaries. This period

witnessed the establishment of academic departments, professional societies, and specialised journals that, whilst fostering expertise within specific domains, simultaneously erected barriers to cross-disciplinary collaboration and knowledge integration. The resulting academic structure, characterised by what Kuhn (1970) termed "disciplinary matrices," has proven remarkably resilient, shaping not only how knowledge is produced but also how it is evaluated, disseminated, and rewarded within academic institutions.

However, the challenges confronting contemporary society increasingly demand interdisciplinary or transdisciplinary approaches that transcend traditional disciplinary boundaries. Climate change, for instance, necessitates integration of atmospheric science, economics, sociology, political science, and engineering, amongst other fields. Similarly, the COVID-19 pandemic demonstrated the necessity of combining epidemiological, social, economic, and behavioural insights to develop effective public health responses. Yet the academic infrastructure remains largely organised around disciplinary lines, creating what Thorén and Persson (2013) characterise as a fundamental mismatch between the structure of knowledge production and the nature of contemporary challenges.

1.2. The Problem Statement

The central problem addressed in this manuscript concerns the fundamental incompatibility between the epistemological requirements of cutting-edge multidisciplinary research and the ontological assumptions embedded within traditional academic knowledge production and evaluation systems. This incompatibility manifests across multiple dimensions of academic practice, creating systemic barriers that impede the advancement of knowledge necessary for addressing complex contemporary challenges (Montgomery, 2024). We propose that artificial intelligence, despite current limitations, offers unprecedented opportunities to transcend these barriers through sophisticated evaluation systems capable of integrating diverse disciplinary perspectives.

At the ontological level, different academic disciplines operate with fundamentally divergent assumptions about the nature of reality, causation, and appropriate units of analysis. Natural sciences typically embrace realist ontologies that assume the existence of objective, discoverable laws governing natural phenomena. Social sciences, by contrast, often adopt constructivist or interpretivist ontologies that emphasise the socially constructed nature of reality and the importance of meaning-making processes. Humanities disciplines frequently operate with hermeneutic ontologies that prioritise textual interpretation and cultural understanding. When researchers attempt to integrate insights across these domains, they encounter what philosophers of science term "ontological incommensurability"—the absence of shared foundational assumptions about the nature of reality itself.

1.3. Learning from Historical Failures

Before proposing AI-enhanced solutions, we must examine why previous attempts at automating academic evaluation have failed or produced unintended consequences. The bibliometric trap illustrates how quantitative metrics can fundamentally distort academic behaviour. The widespread adoption of impact factors and h-indices led researchers to optimise for metrics rather than quality, resulting in salami-slicing of publications, citation cartels, and the proliferation of predatory journals. This demonstrates that any AI system must avoid creating new metrics that can be similarly gamed.

Citation analysis provides another cautionary tale. Automated citation analysis promised objective research evaluation but failed to account for disciplinary differences in citation practices, the significance of negative citations, and cultural biases that privilege certain forms of academic discourse. This failure demonstrates the danger of assuming algorithmic objectivity without understanding the complex social dynamics underlying academic communication.

Even seemingly benign applications like automated plagiarism detection have revealed unexpected biases. These systems have shown false positive biases against non-native English

speakers and interdisciplinary work that legitimately draws from multiple sources. This highlights how automation can discriminate against already marginalised groups within academia, reinforcing existing power structures rather than challenging them.

These failures teach us that technological solutions must be designed with deep understanding of academic culture, power dynamics, and potential for unintended consequences. They remind us that metrics shape behaviour, algorithms encode biases, and automation can amplify rather than alleviate existing inequalities.

1.4. Technological Uncertainties and Aspirations

Whilst this manuscript advocates for AI-enhanced peer review systems, we must explicitly acknowledge the distinction between current technological capabilities and aspirational functionalities. Present-day natural language processing systems, though sophisticated, cannot yet fully comprehend the nuanced argumentation and theoretical subtleties that characterise high-quality academic discourse. Machine learning algorithms excel at pattern recognition but struggle with genuine understanding of novel theoretical contributions.

The capabilities we describe throughout this manuscript—such as AI systems that can transcend human cognitive limitations or perform automated hypothesis generation—represent technological aspirations based on projected developments rather than current realities. We estimate that basic AI-assisted review functions such as bias detection and reviewer matching are achievable within two to three years, whilst more sophisticated capabilities including autonomous evaluation and theoretical synthesis may require ten to fifteen years of technological advancement.

This temporal uncertainty necessitates a flexible implementation framework that can adapt as technologies mature. We must resist the temptation to oversell current capabilities whilst maintaining vision for transformative possibilities. The path forward requires honest assessment of present limitations combined with thoughtful preparation for future capabilities.

2. Theoretical Framework: Critical Engagement with Techno-Solutionism

2.1. Engaging with AI Critics

Our framework directly addresses critiques from scholars who warn against algorithmic governance in academia. Cathy O'Neil identifies three characteristics of problematic algorithms that create "weapons of math destruction": opacity, scale, and damage. Each of these concerns requires careful consideration in designing AI-enhanced peer review systems.

Regarding opacity, we propose that all AI evaluation algorithms must be open-source and subject to academic scrutiny. This transparency extends beyond simply publishing code to include comprehensive documentation of training data, design decisions, and known limitations. The academic community must be able to interrogate and critique these systems just as rigorously as they would any research methodology.

The question of scale presents particular challenges. Whilst the efficiency of AI systems lies partly in their ability to operate at scale, rapid widespread deployment risks amplifying any embedded biases or flaws. Our phased implementation approach deliberately limits scale until effectiveness is proven, allowing for course corrections based on empirical evidence rather than theoretical assumptions.

Addressing potential damage requires robust appeals processes that ensure no researcher's career is derailed by algorithmic decisions. Human oversight must remain paramount, with AI serving as augmentation rather than replacement for human judgement. Clear accountability structures must ensure that when errors occur, affected researchers have meaningful recourse.

Zeynep Tufekci's warnings about techno-solutionism remind us that peer review's problems stem partly from power structures that technology alone cannot fix. The concentration of publishing power in commercial entities, the dominance of Global North perspectives, and the marginalisation of certain methodological approaches reflect deeper structural inequalities within academia. Our

governance framework must explicitly address these dimensions rather than assuming technology will automatically democratise academic evaluation.

Ruha Benjamin's concept of the "New Jim Code" demonstrates how algorithms can perpetuate racial bias through encoded inequity. In academic contexts, this might manifest as AI systems trained predominantly on research from elite institutions reproducing their biases about what constitutes "quality" scholarship. Our multi-model approach aims to prevent such homogenisation by explicitly valuing diverse epistemological traditions and creating space for marginalised perspectives.

Virginia Eubanks' work on algorithmic inequality shows how automated systems can punish the poor and marginalised. In academic contexts, this could translate to discrimination against researchers from under-resourced institutions, those working in local languages, or those employing methodologies that deviate from dominant paradigms. Our economic analysis addresses how to prevent AI peer review from exacerbating academic inequalities.

2.2. Governance and Power Distribution

The implementation of AI-enhanced evaluation systems raises fundamental questions about power dynamics within academia. Current peer review systems, despite their flaws, distribute power across thousands of individual reviewers. AI systems risk concentrating this power in the hands of those who control the algorithms. This concentration could prove even more problematic than current journal editor gatekeeping if not carefully managed.

We propose a distributed governance model that prevents any single entity from controlling AI evaluation systems. This requires establishing international consortiums with representation from diverse geographical regions, institutional types, and disciplinary traditions. Governance structures must include not only senior academics but also early-career researchers, practitioners, and representatives from marginalised communities within academia.

Algorithmic transparency forms a cornerstone of this governance approach. All AI evaluation algorithms must be open-source, allowing the global academic community to scrutinise their operation. This transparency must extend to training data, with clear documentation of what research was included and why. Regular algorithmic audits should examine both technical performance and social impacts, with results published openly.

Democratic oversight mechanisms must ensure that AI systems serve the academic community rather than controlling it. Representative committees from diverse disciplines and geographical regions should oversee AI system development, with rotating membership to prevent entrenched interests. These committees must have real power to modify or halt AI systems that demonstrate bias or other harmful effects.

Protecting epistemological pluralism requires explicit safeguards ensuring AI systems preserve rather than homogenise legitimate methodological and theoretical diversity. This means recognising that different fields have different standards of excellence and that innovation often comes from challenging established norms. AI systems must be sophisticated enough to distinguish between poor quality research and research that simply employs unfamiliar approaches.

Appeal mechanisms must provide robust processes allowing researchers to challenge AI evaluations through human review panels. These appeals should not require researchers to prove algorithmic error but simply to demonstrate that their work deserves human consideration. The burden of proof must rest with the AI system to justify its decisions, not with researchers to defend against algorithmic judgements.

3. Economic Analysis and Equity Considerations

3.1. Development and Funding Models

The development of AI-enhanced peer review systems raises critical economic questions that must be addressed transparently. Initial development costs, which we estimate at £50-100 million, require careful consideration of funding sources and their implications for system control and access.

International consortium funding presents the most promising approach to prevent single-entity control whilst ensuring adequate resources. This might involve contributions from national research councils, international scientific organisations, and philanthropic foundations committed to open science. Public funding must predominate to ensure public benefit, with transparent budgeting processes preventing scope creep or capture by commercial interests.

The operational model chosen will profoundly impact equity and access. A public good model, funded by research councils and free to all users, would best serve the global academic community. This approach treats peer review as essential research infrastructure, similar to particle accelerators or telescope facilities, requiring public investment for public benefit.

Alternatively, a cooperative model where academic institutions collectively own and operate AI systems could provide sustainable funding whilst maintaining academic control. This might operate similarly to existing library consortiums, with institutions contributing based on their size and resources whilst ensuring universal access.

We strongly caution against purely commercial models that could lead to the same profit-driven distortions currently plaguing academic publishing. The commercialisation of peer review through AI must be resisted to prevent further commodification of academic knowledge production.

3.2. Addressing Global Inequalities

AI systems risk exacerbating inequalities between wealthy and resource-poor institutions unless explicitly designed for equity. Access must be universal, with systems freely available to researchers globally regardless of their institution's ability to pay. This requires developing low-bandwidth versions for areas with limited internet access, multilingual interfaces that don't privilege English, and technical support available across time zones.

Training data equity presents particular challenges. Current academic databases overrepresent research from wealthy countries and prestigious institutions. AI systems trained on such biased data will inevitably perpetuate these biases. Deliberate efforts must collect and include research from Global South institutions, non-English publications, and alternative publication formats such as working papers and policy reports.

Governance equity requires that leadership structures include representatives from different economic contexts, various academic systems, and marginalised communities. Token representation must be avoided in favour of meaningful participation in decision-making. This might require providing resources for representatives from under-resourced institutions to participate fully, including funding for travel, time release, and technical support.

The benefits of AI-enhanced peer review must flow to all researchers, not just those at wealthy institutions. This means ensuring that efficiency gains translate into faster publication times for all, that bias reduction benefits marginalised researchers most, and that innovation detection particularly supports unconventional research from unexpected sources.

4. Legal and Regulatory Framework

4.1. Intellectual Property and Data Protection

AI systems processing academic manuscripts raise complex intellectual property questions that current legal frameworks struggle to address. Copyright law, designed for human readers, becomes complicated when AI systems analyse papers for patterns and quality indicators. We propose establishing statutory exceptions for AI peer review under fair use or fair dealing provisions, similar to existing text mining exceptions in some jurisdictions.

Database rights present another challenge, particularly regarding ownership of training datasets and evaluation data. While individual papers remain property of their authors, the aggregated insights derived from analysing thousands of papers create new forms of intellectual property. Clear frameworks must establish that such derived insights remain common property of the academic community rather than becoming proprietary assets.

Authors' moral rights, including attribution and integrity, must be preserved throughout AI processing. Systems must ensure that ideas and innovations are properly credited to their originators, preventing AI from obscuring the human creativity underlying research. This becomes particularly important if AI systems begin suggesting improvements or identifying connections between papers.

Data protection regulations such as GDPR provide a framework that must be adapted for academic contexts. Purpose limitation principles mean data collected for peer review cannot be repurposed for other uses without explicit consent. Data minimisation requires collecting only information necessary for evaluation, resisting the temptation to build comprehensive researcher profiles.

The right to explanation becomes crucial in AI peer review contexts. Authors must be able to understand why their papers received particular evaluations, requiring AI systems to provide interpretable justifications rather than opaque scores. This interpretability must be built into systems from the beginning rather than added as an afterthought.

5. Unintended Consequences and Mitigation Strategies

5.1. Anticipating and Preventing Negative Effects

The history of academic metrics demonstrates how evaluation systems shape researcher behaviour in unexpected ways. AI-enhanced peer review risks creating new forms of gaming as researchers optimise their papers for algorithmic approval rather than genuine quality. This could lead to homogenisation of writing styles, methodological approaches, and research topics as researchers converge on patterns that AI systems reward.

Mitigation requires regularly updating algorithms to prevent gaming, maintaining human oversight to detect manipulation, and employing diverse evaluation criteria that resist simple optimisation. AI systems must be sophisticated enough to detect and penalise obvious gaming attempts whilst remaining open to genuine innovation.

New forms of bias may emerge from AI systems that differ from but prove equally problematic as human biases. AI might systematically undervalue certain types of argumentation, privilege quantitative over qualitative approaches, or favour research that fits established patterns. Continuous bias auditing must examine not just protected characteristics but also methodological, theoretical, and cultural biases.

The loss of serendipity presents a subtle but significant risk. Human reviewers sometimes recognise brilliant work that breaks all conventional rules. AI systems might struggle to distinguish between incompetence and revolutionary innovation. Explicit innovation detection mechanisms must be built into AI systems, along with "minority report" functions that flag unusual work deserving special human attention.

Standardisation pressure could emerge as global AI systems impose uniform standards across diverse academic cultures. What counts as good argumentation varies across cultures, and AI must not impose Western academic norms globally. Regional variations and cultural adaptations must be built into systems whilst maintaining rigorous quality standards.

5.2. Preserving Academic Values

Essential academic values must be actively protected rather than assumed to persist through technological transformation. Academic freedom requires that AI systems never constrain research topics or approaches based on political, commercial, or ideological criteria. Researchers must remain free to pursue unpopular questions and challenge established wisdom.

The collegial nature of peer review, where scholars help improve each other's work, must survive automation. AI systems should augment rather than replace the constructive feedback that characterises good peer review. This means designing systems that facilitate rather than supplant human communication between authors and reviewers.

Disciplinary evolution requires that fields remain free to develop new standards and approaches. AI systems must not freeze current disciplinary norms but allow for organic development of new methodologies and criteria. Regular retraining and updating of AI systems must incorporate emerging scholarly practices rather than enforcing outdated standards.

The mentorship aspect of peer review, where senior scholars guide junior researchers, provides crucial professional development. AI systems must preserve opportunities for developmental feedback whilst potentially democratising access to high-quality review. This might involve AI identifying areas where human mentorship would prove most valuable.

6. Implementation Strategy

Learning from critiques and historical failures, we propose a cautious implementation strategy that prioritises evidence over enthusiasm. The proof of concept phase, spanning years one to three, would involve small-scale trials in willing disciplines with parallel human review for comparison. Extensive bias and equity auditing would examine both technical performance and social impacts, with public reporting ensuring transparency.

Controlled expansion in years four to seven would gradually extend successful approaches based on empirical evidence rather than theoretical promises. Continuous monitoring for unintended effects would enable course corrections, whilst regular stakeholder consultation would ensure systems serve community needs. Legislative framework development during this phase would establish necessary legal structures before full deployment.

Mature implementation, beginning in year eight, would occur only after proven benefits justify the risks. Human oversight would remain integral rather than vestigial, with regular algorithmic auditing ensuring systems continue serving their intended purposes. Governance structures would evolve based on experience whilst maintaining core commitments to equity and transparency.

This extended timeline reflects lessons from previous technological implementations in academia. Rushed deployment of metrics like impact factors created problems that persist decades later. By proceeding cautiously, we can identify and address issues before they become entrenched in academic culture.

7. Conclusion

This investigation has presented a critical framework for AI-enhanced peer review that engages seriously with techno-solutionist critiques, historical failures, and potential negative consequences. We maintain that whilst AI offers possibilities for addressing multidisciplinary research challenges, implementation requires unprecedented attention to governance, equity, and academic values.

Our analysis reveals that technology alone cannot solve the deep-seated problems in academic evaluation. The ontological and epistemological barriers between disciplines reflect genuine differences in how fields understand knowledge and reality. Social, economic, and political reforms must accompany any technological intervention. The framework we propose—with its emphasis on transparency, equity, and democratic governance—offers one path forward, but only if implemented with appropriate caution and continuous critical reflection.

The urgency of contemporary challenges creates pressure for rapid transformation of academic systems. Climate change, pandemics, and social inequalities demand interdisciplinary responses that current peer review systems struggle to support. Yet this urgency must not lead to reckless implementation of untested systems that could cause more harm than good.

We conclude with neither naive optimism nor paralysing pessimism. The transformation of academic evaluation requires not just technological innovation but fundamental commitment to justice, diversity, and the preservation of what makes academic inquiry valuable to humanity. AI-enhanced peer review represents one possible future, promising yet perilous, requiring our most thoughtful engagement.

The academy stands at a crossroads. The path forward requires courage to imagine radical alternatives whilst maintaining commitment to rigour, diversity, and ethical responsibility. It demands that we learn from past failures whilst remaining open to transformative possibilities. Most importantly, it requires recognising that the future of academic knowledge production will be shaped by the choices we make today about technology, governance, and values.

May this contribution stimulate the critical dialogue necessary for responsible transformation. The conversation must include not just technologists and administrators but the full diversity of the global academic community. Only through inclusive deliberation can we develop systems that serve both excellence and equity, innovation and tradition, efficiency and humanity.

References

Alvargonzález, D. (2011). Multidisciplinarity, interdisciplinarity, transdisciplinarity, and the sciences. International Studies in the Philosophy of Science, 25(4), 387-403.

Benjamin, R. (2019). Race After Technology: Abolitionist Tools for the New Jim Code. Polity.

Bhaskar, R. (1975). A realist theory of science. Leeds Books.

Biagioli, M., & Lippman, A. (Eds.). (2020). Gaming the Metrics: Misconduct and Manipulation in Academic Research. MIT Press.

Boon, M. (2017). An engineering paradigm in the biomedical sciences: Knowledge as epistemic tool. Progress in Biophysics and Molecular Biology, 129, 25-39.

Boon, M., & Van Baalen, S. (2019). Epistemology for interdisciplinary research -- shifting philosophical paradigms of science. European Journal for Philosophy of Science, 9(1), 16.

Danermark, B. (2019). Applied interdisciplinary research: a critical realist perspective. Journal of Critical Realism, 18(4), 368-382.

Eubanks, V. (2018). Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. St. Martin's Press.

Frodeman, R. (2010). The Oxford handbook of interdisciplinarity. Oxford University Press.

Galison, P. (1997). Image and logic: A material culture of microphysics. University of Chicago Press.

Kuhn, T. S. (1970). The structure of scientific revolutions (2nd ed.). University of Chicago Press.

MacRoberts, M. H., & MacRoberts, B. R. (2018). The mismeasure of science: Citation analysis. Journal of the Association for Information Science and Technology, 69(3), 474-482.

Montgomery, R. M. (2024). The Obsolescence of Traditional Peer Review: Why AI Should Replace Human Validation in Scientific Research. Preprints. https://doi.org/10.20944/preprints202411.0246.v1

O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown.

Proctor, D. M., Dada, N., Serquiña, A., & Willett, J. L. E. (2023). Problems with peer review shine a light on gaps in scientific training. mBio, 14(3), e03183-22.

Schmidt, J. C. (2008). Towards a philosophy of interdisciplinarity: An attempt to provide a classification and clarification. Poiesis & Praxis, 5(1), 53-69.

Thorén, H., & Persson, J. (2013). The philosophy of interdisciplinarity: Sustainability science and problem-feeding. Journal for General Philosophy of Science, 44(2), 337-355.

Tufekci, Z. (2017). Twitter and Tear Gas: The Power and Fragility of Networked Protest. Yale University Press.

Weber-Wulff, D. (2019). Plagiarism detectors are a crutch, and a problem. Nature, 567(7749), 435.