Too Slow for the AI Age? Building a Dynamic Research Continuum

Ewan Simpson

Narxoz Business School, Narxoz University

December 4 2025

Abstract

Academic publishing is misaligned with the tempo of contemporary artificial intelligence. Large models alter capabilities, risks and regulatory questions on a quarterly cycle, while journal review and publication often take years. This tempo gap displaces epistemic authority from public, peer-reviewed knowledge to corporate laboratories, private platforms and informal commentary. The paper argues that closing this gap requires redesigning the research pipeline itself. It proposes a Dynamic Research Continuum (DRC): a continuous, versioned publication architecture that integrates preprints, open and AI-assisted review, and "living" synthesis articles into a single visible knowledge loop. Unlike piecemeal open-science reforms, DRC treats access, review and synthesis as one lifecycle explicitly designed for high-velocity AI domains. The architecture is operationalized through the Iterative Human–AI Co-Creation (IHACC) framework, which specifies how human researchers and AI systems cycle through scoping, drafting, critique and updating while preserving judgement, transparency and epistemic inclusion. Using a historical reconstruction of journal publishing, a diagnosis of current challenges in AI-relevant fields and an analysis of reform options, the paper argues that timely, trustworthy and openly updated research has become a precondition for legitimate governance in the AI age.

Keywords: artificial intelligence; dynamic research systems; higher education research; future of universities.

1. Introduction: a tempo mismatch

Since 2022, frontier artificial intelligence systems have moved from research prototypes to widely deployed infrastructure on a quarterly cycle. Models such as GPT 5, Claude, and Gemini are updated, expanded and integrated into products at a pace that changes what is technically and economically possible within months rather than decades. These systems now shape education, healthcare, finance, media and public administration, not as distant possibilities but as live inputs into everyday tools, services and decisions.

By contrast, the academic publishing system that still underpins most formal evidence and many policy processes moves much more slowly. Even in fast moving fields, it is common for two to three years to pass between a study being conceived and a final journal article appearing in print. Peer review remains opaque, editorial decisions are rarely explained, and most journals report only partial time metrics. Preprint servers have shortened the delay between completion and first disclosure, but they sit beside rather than inside the core publication workflow. In many fields preprints receive little formal review, are not systematically synthesized and are weakly linked to policy or regulatory debates.

The result is a widening tempo mismatch, an AI impact gap where systems reshape institutions faster than independent evidence can catch up. In domains that are central to AI deployment, such as medicine, education and law, decisions about adoption, regulation and risk management often have to be made before any mature, peer reviewed evidence base exists. Medical regulators have occasionally been able to move faster, because clinical trial cultures and preexisting infrastructures for evidence synthesis already exist. In education, labor markets and public administration, evidence arrives late, scattered across journals and conferences, and rarely integrated into living reviews that speak directly to policy questions. Where reliable synthesis is slow or absent, epistemic authority shifts by default toward corporate labs, consulting firms and informal online commentary.

Treating this situation as a temporary backlog underestimates the scale of the problem. If the underlying architecture of publishing remains organized around infrequent, static articles and slow, closed review, incremental speed ups will not restore alignment with AI development cycles. In an environment where capabilities, risks and use cases are continually recombined, a system that delivers isolated snapshots years after the fact cannot reliably support legitimate governance.

This paper argues that closing the gap requires redesigning the research pipeline itself rather than marginally accelerating existing workflows. Section 2 reconstructs how the current journal model emerged and why it is resistant to change. Section 3 sets out the specific failures that now matter in AI-relevant domains, including volume without synthesis, inequitable attention, and weak links between research and decision making. Section 4 analyses AI as both catalyst and stress test for the existing system. Sections 5 and 6 introduce the Dynamic Research Continuum (DRC) as a continuous, versioned publication architecture and the Iterative Human AI Co Creation (IHACC) (Simpson, 2025) framework as a practical human and AI workflow for operating within it. Section 7 connects DRC and IHACC to questions of temporal legitimacy and agenda setting in AI governance, and Section 8 outlines feasible implementation pathways for funders, publishers, universities and regulators. The conclusion summarises the implications for research, policy and institutional design.

The core claim is that only a versioned, openly reviewed and continuously synthesized system of living research, operated through explicit human and AI workflows, can restore temporal legitimacy to evidence informed governance in the AI age.

2 Historical Roots of the Academic Knowledge Pipeline

The current mismatch between AI’s tempo and the pace of academic publishing is not an accident or a temporary bottleneck. It is the predictable outcome of a publication architecture that was assembled for a print world, then layered with prestige incentives and partial digital fixes. This section traces how that architecture emerged and why it now struggles to provide timely, trustworthy knowledge about fast moving technologies and their social impacts.

2.1 Print era legacies

Modern scholarly publishing descends from the *Philosophical Transactions of the Royal Society* in the seventeenth century and the broader culture of learned journals that followed. Early journals existed to register priority, share correspondence, and curate small volumes of work that could be typeset and shipped a few times a year. The unit of production was the issue, not the individual article. Everything was constrained by the costs of paper, typesetting, and physical distribution.

This environment produced a naturally sequential workflow. Manuscripts were submitted by post, selected by editors and a small circle of trusted advisors, set in type, and batched into issues. That sequence made sense when the realistic alternatives were books, letters, or nothing at all. It also baked in a mental model of scholarship as a one way pipeline from author to editor to printer to reader, with significant delays at each point. The core structure has survived, often with surprisingly little change, even as the physical constraints disappeared (Suber, 2012).

2.2 Institutionalization of peer review

Anonymous external peer review only became widespread in the twentieth century, particularly after World War II, when government science funding, enrollments, and journal volumes expanded sharply. Peer review promised at least three things at once: quality control, triage under scarcity, and a way to distribute gatekeeping across a wider community rather than relying only on editors and their close networks.

In practice, the system has always been uneven. Different disciplines adopted peer review at different times and with different norms. Empirical studies show that review processes vary widely in rigor, reliability, and bias, and that much of the work is done by a relatively small pool of overloaded reviewers (Hosseini & Horbach, 2023; Squazzoni et al., 2021). Yet peer review became a symbolic gold standard. Journals and universities began to treat “peer reviewed” as a binary label that conferred legitimacy, without much concern for how long the process took or how transparent it was (Royal Society, 2017).

Crucially, the system was never designed for speed. It was designed to ration scarce pages and to filter submissions in a world where print was expensive and slow. That design choice now collides directly with the need to understand AI systems that change on quarterly cycles and that are already reshaping labor markets, education systems, and public administration.

2.3 The rise of the prestige economy

The next major layer was the metrics regime that grew around the journal impact factor. Garfield’s (2006) work on citation indexing and the history and meaning of the journal impact factor began as an attempt to map influence and help libraries make subscription decisions. Within a few decades, the impact factor had become a central currency in academic status contests. Hiring, promotion, and funding decisions increasingly turned on where papers were published rather than what they contained.

This created a prestige economy that rewards exclusivity and perceived selectivity. High impact factor journals accept a small fraction of submissions, build long pipelines of “revise and resubmit” cycles, and can maintain long delays because scarcity increases their signaling power. Authors learn to treat long review times as the price of admission to the top of the hierarchy. Metrics that were originally descriptive become targets. As Masum et al. (2013) and others have noted, this tends to concentrate attention and resources in a small subset of venues and reinforces conservative, low risk behavior by editors and reviewers.

The result is a structural misalignment between prestige and velocity. Publishing quickly in a lower impact venue can be rational for science and policy, but it is often irrational for individual careers. Journals of similar scope can have radically different turnaround times, yet slower venues often carry higher status. That pattern makes it difficult to treat timeliness itself as an indicator of quality or responsibility.

2.4 Digital migration without reform

From the 1990s onward, journals moved to electronic submission systems, PDF delivery, and online archives. In principle, digital tools could have enabled far more radical changes. Articles could have been published as soon as they cleared review, versioned as new data arrived, and linked to open reviews and replication packages. In practice, most publishers treated digital migration as a change in interface, not architecture.

Online portals replaced mailed manuscripts, but the underlying sequence of submission, confidential review, editorial decision, and batch publication remained in place. Continuous publication and “online ahead of print” reduced some delays but did not alter the assumption that an article is a fixed object that passes through a one-way pipeline. Experiments with open peer review, post-publication commentary, and more modular workflows remain marginal compared with the dominant closed model (Royal Society, 2017; Vale, 2015). At the same time, journals began to report performance using a small set of headline statistics, usually some combination of “submission to first decision” and “submission to acceptance.” These figures give an impression of efficiency but hide time lost to desk rejections, cascades through publisher portfolios, and resubmissions elsewhere, and they rarely disclose variation by article type or field.

Some publishers have shown that faster, more transparent pipelines are technically and organizationally possible. MDPI, for example, reports median review times on the order of a few weeks for many journals (MDPI, 2021). Yet these models are often treated with suspicion in prestige-driven disciplines, and they rarely shape the incentives that matter for tenure and grants. Digital infrastructure has increased volume and reach without fundamentally rethinking the way knowledge moves or how delay is measured.

2.5 Invisible delays in practice

Where researchers look beyond these headline metrics, they find much more complex and discouraging patterns. Powell (2016) documents long delays and large variance across disciplines and journals. Even within the same field there can be gaps of many months between the fastest and slowest journals, and these gaps are not consistently reflected in status or pricing. Work on open-source tools in machine learning shows that widely used software can shape practice and become de facto standards long before any formal publication or journal review process catches up (Sonnenburg et al., 2007).

For authors, these delays show up as months or years of uncertainty and the need to navigate multi-journal submission paths. For readers and policymakers, they show up as a knowledge environment where it is hard to tell which areas are stalled because there is no research and which are stalled because research is trapped in the pipeline. In the context of AI, where models and deployment patterns can change every few months, this opacity makes it very easy for corporate technical reports and consultancy white papers to dominate attention while careful academic work is still in transit.

2.6 Institutional lock in

Three intertwined forces now keep academic publishing trapped in slow cycles, even though the technology exists to move faster and in more transparent ways:

- Economic lock in. Subscription and article processing charge revenues are tied to existing journal portfolios and to large multi-year deals with libraries and consortia. Launching radically different workflows would require publishers and societies to take real financial risks, and it could undermine the pricing power of their flagship titles. Open access reforms have made some progress, but they have mostly been layered on top of the existing journal structure rather than replacing it (Suber, 2012).
- Cognitive and cultural heuristics. Researchers, evaluators, and university managers equate selectivity and brand strength with quality. A slow, highly selective journal is often seen as more rigorous than a faster, transparent one, even when studies find weak correlations between review time, impact factor, and methodological rigor (Garfield, 2006; Tennant et al., 2016; Vale, 2015). This status logic makes it individually rational to feed the slowest parts of the system, even when it collectively undermines responsiveness to urgent social questions.
- Workflow inertia. Editorial systems and referee norms are optimized for confidential, narrative reports handled in Word or PDF, not for modular, data-linked, versioned research objects. Moving to a more dynamic model would require retraining editors, redesigning incentives for reviewers, and investing in new platforms that can support open, AI-assisted workflows. That is a non-trivial organizational challenge, especially for smaller societies and journals that already depend on volunteer labor (Hosseini & Horbach, 2023; Squazzoni et al., 2021).

In sum, the problem is not technical incapacity. It is structural and incentive-driven. The core architecture of the academic knowledge pipeline was tuned for a print era in which slow, selective, and scarce publication could still claim to represent the cutting edge. In an AI-accelerated world, where the most powerful systems are deployed quickly and begin to reshape institutions long before journals can respond, those same legacies now leave public research struggling to keep pace. Section 3 turns to the current manifestations of this mismatch and to the specific ways in which it distorts our understanding of AI and its impacts on society, business, and education.

3 Current challenges

The historical legacies described above now show up as a set of concrete failures in how knowledge about AI is produced, organized, and made usable. None of these problems is entirely new, but the combination of speed, scale, and concentration in contemporary AI makes them much more consequential. In fields that matter for AI deployment, the result is not just delay but systematic distortion. Evidence arrives too late, in the wrong form, or with the wrong distribution of attention to inform decisions at the pace models are actually being deployed.

This section highlights six linked problems: volume without synthesis, inequitable citation and attention, disciplinary silos, incentives for performative relevance, new forms of epistemic risk as AI tools enter the workflow, and the way these combine into a structural crisis rather than a series of isolated glitches.

3.1 Volume without synthesis

The volume of AI related research has grown at an exponential rate. Mapping studies find that AI and machine learning now account for a large and rapidly rising share of outputs in computer science, medicine, and parts of the social sciences, with similar growth in adjacent areas such as data science and information systems. In education, labor economics, and governance there is a steady stream of pilots, case studies, and policy reports that examine AI based tools in schools, workplaces, and public services. On paper, there is no shortage of material.

Preprints and rapid communications help put findings into circulation earlier, as Section 1 already noted. During COVID 19, platforms such as medRxiv and bioRxiv showed that preprints can drastically shorten the time between study completion and public availability. In AI related fields, arXiv and similar repositories now serve as de facto first publication venues. What is missing is systematic synthesis. There are few maintained, living reviews that continuously integrate new evidence, distinguish between weak and strong studies, and present a coherent picture of what is known and what remains uncertain in a form that regulators, practitioners, and publics can use.

3.2 Citation, attention, and inequality

Citation practices further entrench inequality in who gets heard. Highly resourced universities and corporate labs are structurally better placed to publish quickly in visible venues, promote their work, and shape both technical and policy agendas. Their preprints and white papers are more likely to be shared, discussed in influential forums, and taken up in subsequent research, regardless of whether they are methodologically stronger than less visible work.

Metrics amplify these asymmetries. Journal impact factors, h indices, and altmetrics reward visibility and network position rather than the social value or contextual relevance of findings. Policy makers under time pressure often rely on highly cited or highly publicized sources as proxies for reliability. This tends to concentrate epistemic authority in a small number of institutions and vendors. For AI, that means a narrow set of perspectives and deployment contexts are over represented in the evidence base that shapes regulation and public debate, while work from low and middle income settings, smaller institutions, and marginalized communities struggles to register.

3.3 Disciplinary silos and the missing middle

The most important questions about AI often sit in the gaps between disciplines. Research on model architecture and training data is concentrated in computer science and corporate labs. Research on labor markets, education systems, health services, and public administration usually lives in separate disciplinary communities with their own journals, conferences, and evaluative norms. Work on ethics, law, and political theory circulates in yet another set of venues.

These silos slow translation. An education ministry that wants to understand how chatbots affect student learning and assessment has to navigate computer science papers on large language models, cognitive science studies on feedback and motivation, sociology of education work on inequality, and policy reports on assessment regimes. Few venues are designed to integrate this "missing middle" of interdisciplinary knowledge into timely, decision focused syntheses. As a result, system level questions about AI in schools, workplaces, or courts are often answered with narrow technical studies or abstract ethical arguments rather than with integrated, cross disciplinary evidence.

3.4 Incentive structures and performative relevance

Existing incentives push researchers toward work that is legible to high status journals rather than to the communities that most need guidance. Career progression, grant success, and institutional rankings are heavily tied to publication in a small set of prestige venues. These venues often prefer theoretically neat, generalizable contributions over messy, context specific studies of how AI actually interacts with real institutions and constraints.

As a result, much of the academic literature on AI impacts is retrospective and abstract. It describes "potential" effects or models hypothetical scenarios long after major deployment decisions have already been taken. Studies that would require close collaboration with practitioners, uncomfortable scrutiny of powerful vendors, or long term engagement with vulnerable populations are harder to fund and publish. They often demand interdisciplinary methods and shared credit that the current system does not reward. The appearance of relevance is sometimes maintained through framing and rhetoric rather than through direct engagement with live policy questions.

3.5 Trust, AI tools, and epistemic risk

AI itself is now embedded in research workflows, which creates new challenges for trust and validation. Large language models, code assistants, and automated screening tools are already woven into literature review, drafting, data analysis, and even aspects of peer review. They promise to save time and reduce cognitive load, and early studies suggest real productivity gains for some tasks. At the same time, they introduce opaque transformations into processes that were already only partially transparent.

Without clear norms and infrastructure, AI assistance can both amplify and mask existing problems. Models trained on biased or low quality corpora may reproduce citation distortions, overlook critical but less visible work, or hallucinate details in ways that are hard to detect under time pressure. Automated tools for screening, summarizing, and suggesting edits can quietly shape what counts as a "normal" paper style and which kinds of argument or evidence are treated as peripheral. In peer review, unacknowledged use of AI tools blurs responsibility for judgments and risks deepening confusion rather than clarifying truth claims (Naddaf, 2024; Su et al., 2025).

3.6 A structural crisis, not a series of glitches

Taken together, these challenges amount to a structural crisis rather than a sequence of isolated inefficiencies. Volume without synthesis, unequal attention, disciplinary fragmentation, misaligned incentives, and new epistemic risks from AI tools all stem from the same underlying architecture: a system built for infrequent, static articles and closed review, now asked to cope with high-velocity technologies that interact with every sector of society. The question is not only how to move manuscripts faster. It is how to redesign the knowledge pipeline so that it can produce timely, trustworthy, and socially grounded syntheses in domains where AI deployment is already under way. The next section examines AI itself as both catalyst and stress test for that redesign.

4 AI as catalyst and stress test

AI is not only an object of study for academic research. It is already embedded in the workflows through which research is produced, evaluated, and communicated. That dual role makes it both a catalyst for change and a stress test for the existing publishing architecture. On one side, AI tools promise to accelerate many research tasks, from literature review to drafting and coding. On the other, they increase the volume of material, introduce new forms of opacity and error, and raise the stakes of delay by accelerating the technologies that research is meant to interpret.

This section examines three aspects of that tension: AI in researcher workflows, AI in peer review and editorial work, and frontier models as a limit case for the current system.

4.1 AI in research workflows: acceleration and epistemic risk

Large language models and related tools are already integrated into everyday research practice. Code assistants help to write and debug scripts. Text models screen and summarize literature, generate outlines, suggest titles and abstracts, and draft sections of papers. Automated tools support data cleaning, descriptive analysis, and visualization. Early studies find substantial time savings for repetitive tasks and for researchers with less experience in coding or academic English, and many researchers report subjective gains in productivity and reduced cognitive load.

From the perspective of the knowledge pipeline, this is a powerful accelerator. If researchers can scope, prototype, and draft more quickly, then in principle more work should enter the publication stream and more variants of a question can be explored. That effect is amplified when funders and universities begin to assume AI use as a baseline competence. Over time, the idea of producing articles or syntheses without some form of AI assistance may start to look inefficient or even irresponsible.

The same tools, however, introduce new epistemic risks and amplify old ones. Language models trained on large, partially opaque corpora reproduce citation distortions and coverage gaps that privilege already visible institutions and languages. They hallucinate references, misattribute claims, and generate plausible but incorrect summaries that are hard to detect under time pressure. Automated screening and topic modelling can quietly frame which studies are seen as relevant or central. Drafting assistance tends to normalize particular styles of argument and presentation, which may narrow the range of acceptable contributions and marginalize nonstandard methods or voices.

These dynamics are particularly problematic in AI related fields, where much of the training data for models overlaps with the literature that needs to be evaluated. Without explicit versioning, disclosure norms, and validation procedures, it becomes difficult to know which parts of a paper reflect human judgment and which reflect opaque model outputs. The promise of acceleration is real, but in a pipeline that lacks robust mechanisms for tracking provenance and revision, acceleration can just as easily propagate error and noise.

4.2 AI in peer review and editorial decision making

AI tools are also entering peer review and editorial work. Journals and conferences already use automated systems for plagiarism detection, reference checking, and initial screening. Some are experimenting with language models to suggest reviewers, flag potential methodological issues, or generate structured summaries of submissions. Reviewers themselves may quietly use conversational models to help structure reports, rewrite feedback in more neutral language, or check aspects of the literature.

Used carefully, these tools could mitigate some long standing problems. They might help editors handle rising submission volumes, surface relevant prior work that reviewers have missed, or standardized certain parts of the review process. Structured AI generated summaries could make it easier to see patterns in reports across multiple reviewers, and automated checks could improve the consistency of basic methodological and reporting standards.

At the same time, AI assisted review raises unresolved questions about accountability, bias, and legitimacy. If models are used to triage submissions, they may embed and reinforce existing disciplinary and linguistic biases without clear avenues for appeal. If reviewers rely on AI to propose critiques or references, responsibility for the content of reviews becomes blurred. Authors may receive feedback that reflects the training distribution of a model as much as the considered judgment of a human expert. Editors face new pressures to move quickly based on machine generated signals that they may not fully understand.

Most importantly for this paper, AI in peer review interacts with a system that is already opaque and slow. Confidential reports, undisclosed use of tools, and limited post publication dialogue mean that errors or biases introduced by AI are hard to detect and correct. Rather than simply automating existing steps, a serious response would require redesigning review as part of a more open, traceable, and iterative process, in which both human and AI contributions to evaluation are visible and contestable.

4.3 Frontier models as a limit case for the publishing architecture

The interaction between AI and academic publishing is most visible at the frontier, where new systems can have system level effects on science, industry, and governance. The AlphaFold case is one often cited example: a model released by a corporate lab reshaped parts of structural biology and drug discovery far faster than traditional publication and synthesis mechanisms could respond. Something similar is now happening with general purpose language and multimodal models that affect education, law, media, and public administration.

Frontier foundation models are developed, deployed, and iterated in cycles measured in months. Their capabilities, risk profiles, and application ecosystems change rapidly as new versions are released, fine-tuned, or combined with other systems. Policy debates about safety, labor markets, misinformation, and education are already being shaped by vendor reports, benchmarks, and technical blog posts rather than by peer reviewed syntheses. By the time careful studies appear in journals, many of the most important deployment decisions have already been made and path dependencies are locked in.

In this environment, the current publishing architecture functions as a trailing indicator. It can document and critique past waves of deployment, but it struggles to shape live decisions. The limit case is not a speculative artificial general intelligence in the abstract. It is a world in which increasingly capable, general purpose models are continuously updated and embedded across sectors while the public knowledge system that is supposed to help govern them remains organized around slow, static outputs and closed review.

Seen this way, AI exposes the structural weakness of the existing system rather than simply stressing it at the margins. Accelerating individual stages of the pipeline will not be enough if the basic unit of publication remains an isolated article, and if synthesis is treated as an occasional, manual exercise rather than as a continuous public service. What is needed is an architecture that can absorb, version, and openly review research at the pace at which frontier systems change, and a set of human–AI workflows that can turn that flow into timely, trusted syntheses for decision makers. The next section introduces the Dynamic Research Continuum (DRC) as a proposal for such an architecture.

5 The Dynamic Research Continuum (DRC)

The problems identified in Sections 2 to 4 are not going to be solved by attaching a few AI tools to the existing journal pipeline. A system built around static articles, closed review, and weak synthesis cannot simply be "sped up" to match the tempo of contemporary AI. What is needed is a different publication architecture: one that treats research as a continuous, versioned public service rather than a sequence of isolated outputs.

By the Dynamic Research Continuum (DRC) we mean a continuous, versioned publication architecture in which:

- Research outputs appear first as structured preprints with rich metadata and links to code and data.
- Peer review is attached to those preprints in open or semi-open form, with both human and AI tools contributing.
- Curated "living" synthesis articles integrate and update evidence across overlapping questions under explicit update rules, drawing on pilots in medicine and policy such as Cannabis Evidence and F1000Research (Cannabis Evidence, n.d.; F1000Research, n.d.).
- All states and transitions are visible, traceable, and machine-readable, so that readers and tools can see not just conclusions but also the history of critique and revision.

The DRC generalizes ideas that already exist in specific corners of the system. eLife's recent model, living evidence programs, F1000Research, and cOAlition S's Publish–Review–Curate proposals all point toward more open, modular, and iterative publishing (eLife, 2023; cOAlition S, 2023, 2024; Cannabis Evidence, n.d.; F1000Research, n.d.). What is new here is not another local fix. The DRC integrates versioned preprints, open review, and living syntheses into a single lifecycle, and it does so explicitly for high-velocity domains where AI is already changing capabilities and institutions. In the next section, IHACC supplies the corresponding human–AI workflow; together they turn a set of scattered reforms into a coherent architecture for AI-relevant research.

5.1 Design principles

Four design principles guide the DRC.

First, temporal alignment. The system must be able to keep roughly in step with real-world change in AI-intensive domains. That does not mean instant consensus or real-time policy advice, but it does mean that high-quality syntheses should appear while questions are still live, and that updates should be routine as new evidence arrives rather than rewritten from scratch every few years.

Second, epistemic robustness. Faster is not useful if it simply amplifies noise. The DRC is designed to make uncertainty visible, reward careful methods, and support replication, audit, and contestation. Living syntheses should distinguish between strong and weak evidence and report uncertainty honestly, especially where evidence is thin or contested.

Third, transparency and traceability. In a world where AI tools are woven into research and review, it must be possible to see how claims were produced and revised. That requires public version histories, visible review traces, and machine-readable links between data, code, analyses, and narratives, not just a final PDF that hides the deliberation.

Fourth, inclusion and pluralism. The architecture should make it easier, not harder, for under-resourced institutions, researchers in low and middle income countries, practitioners, and affected communities to contribute. That includes support for multiple languages, qualitative and mixed methods, community-based research, and qualitative studies of lived experience alongside more conventional quantitative work.

The Dynamic Research Continuum is not a competing brand name for reviewed preprints, PRC models, or living systematic reviews. It treats those as partial fixes and recombines them into a single architecture that treats tempo as a first class design constraint. Reviewed preprints shorten the time to a first public version, but they still strand evidence in thousands of siloed manuscripts. PRC schemes improve the visibility of reviews, but do not guarantee that policy makers or practitioners ever see the synthesized state of knowledge on a question. Living reviews maintain a small number of summaries in near real time, but they are usually funded as exceptional projects rather than as a default layer of the system. DRC integrates these strands by requiring that new studies appear in a versioned, openly reviewed record and that they are continually pulled into structured syntheses for defined decision domains, with explicit cadence targets for each layer. In other words, it treats reviewed preprints, open review, and living syntheses as components of a single pipeline, not as independent reforms that can be adopted piecemeal without addressing the AI impact gap.

5.2 From linear pipeline to visible loop

The traditional model treats publication as a linear pipeline: manuscripts enter, are reviewed behind closed doors, and emerge months or years later as fixed articles. Section 2 showed how that sequence was optimized for print and scarce page space. In a high-velocity AI environment, the result is a long queue of disconnected snapshots with little visibility into what is happening inside the pipeline.

Conceptually, the DRC replaces the pipeline with a loop that has four recurring stages:

- Exposure. Authors release a structured preprint that contains the core claims, methods, and materials, with clear links to code, datasets, and preregistrations where relevant.
- Critique. Reviews, comments, and replications attach to that preprint. These may include traditional referee reports, signed commentaries, registered replications, or independent audits, some of which can be supported by AI tools for screening and cross-checking
- Synthesis. Living review articles and narrative syntheses incorporate the preprint and its critiques into a broader picture, updating under explicit rules drawn from living evidence methodologies (Cannabis Evidence, n.d.; F1000Research, n.d.).
- Feedback. Insights from syntheses feed back into new studies, policy questions, and practice changes, which then generate new preprints and data.

In a functioning DRC, these stages repeat. The same line of inquiry is not represented by a single "definitive" article but by an evolving record that shows how claims have been tested, refined, and sometimes overturned. For readers, the primary object becomes the visible loop: a cluster of preprints, reviews, and syntheses that makes it easier to see where consensus is emerging, where findings are fragile or contested, and where important gaps remain.

5.3 Core objects in the DRC

Within this loop, the DRC operates on three main kinds of objects.

The first is the versioned research object. This includes empirical studies, methodological notes, conceptual pieces, and data or code releases. Each object has a persistent identifier and a version history. Minor and moderate updates do not require "new" papers; they are recorded as new versions, with clear documentation of what changed and why. In technical fields some of this already happens informally through arXiv updates and code repositories; the DRC treats this as the default, not an exotic exception.

The second is the review trace. Instead of reviews disappearing into editorial inboxes, they are surfaced and credited, even when the final curation choices differ. Review traces can include anonymous reports, signed commentaries, methodological audits, and replication outcomes. They are linked to specific versions of research objects and can themselves be cited and evaluated. This makes the evaluative history of a claim visible and gives reviewers credit for intellectual labor that currently goes unrecognized.

The third is the living synthesis. These are not static narrative reviews updated every decade. They are maintained syntheses with their own identifiers, version histories, and credit structures. Editorial teams, including practitioners and affected communities where appropriate, are responsible for keeping them current according to predefined update rules. In areas where AI is reshaping practice, such as education and labor markets, the expectation is not that each study stands alone, but that relevant work is quickly pulled into at least one maintained synthesis that evolves with the field.

A simple example illustrates how these objects interact. Imagine a study of AI tutors in low-income secondary schools. In a DRC world, the first structured preprint appears with links to code, anonymized data where possible, and a clear description of context and implementation. Reviews and school-level commentaries attach as a visible trace, including critiques of equity and local feasibility. Within weeks, a living synthesis on "AI tutors in compulsory education" incorporates the study, weighing it alongside other trials in different countries. As follow-up cohorts are tracked or new schools are added, the research object is versioned and the synthesis updated. Regulators, school leaders, and parents can see not only headline effect sizes but also how conclusions have shifted across versions and contexts.

5.4 Versioning and delta updates

A central feature of the DRC is the use of explicit versioning and "delta" updates. Instead of treating every change as a new paper, the system distinguishes between:

- Minor updates, such as corrections, additional robustness checks, or clarified wording.
- Moderate updates, such as new data from a follow-up survey or an additional experimental condition that does not overturn the core argument.
- Major updates, such as a change in the theoretical framework or a reversal of the main conclusion.

Living review methodology already recognizes similar distinctions: small batches of new studies may trigger minor narrative adjustments, while larger shifts in the evidence base trigger a new round of review and a new major version (Elliott et al., 2017; Garner et al., 2016).

This approach does two things. It lowers the friction for incremental improvement, making it easier to correct and extend work without waiting for a separate publication slot. And it creates an auditable record of how claims evolve over time, which is essential when AI tools are involved in generation, analysis, or review. Readers can see not only the latest state but also when key turning points occurred, who contributed, and how review feedback affected the trajectory of the work.

5.5 Governance, incentives, and equity

The DRC is not a purely technical proposal. It requires corresponding changes in governance and incentives.

One element is recognition and credit. If peer review, synthesis, and curation are central to the new architecture, then they must count in hiring, promotion, and funding decisions. This implies explicit credit for review traces and living synthesis contributions, the inclusion of these outputs in CVs and assessment dossiers, and metrics that track their use and influence alongside primary publications (Tennant et al., 2016; Hosseini & Horbach, 2023).

Another element is infrastructure for equity. A dynamic system that simply amplifies existing inequalities in visibility and resources would make the situation worse. To avoid consolidating the dominance of a small set of elite actors, DRC implementations should:

- Support low and middle income institutions and independent scholars with shared platforms and subsidized participation.
- Encourage contributions from practitioners, community organizations, and affected groups, especially in living syntheses on topics that affect them directly.
- Use AI tools to surface relevant work from outside the usual citation networks, not only to optimize convenience for already prominent authors.

AI can help here, but only if it is trained and governed with equity in mind. Search, recommendation, and summarization tools need to be tuned to avoid simply reproducing existing attention hierarchies. Section 6 takes up this question through the Iterative Human–AI Co-Creation (IHACC) framework, which specifies how human judgment and AI assistance can be combined without erasing pluralism or undermining trust.

5.6 How the DRC addresses the AI impact gap

The point of the DRC is not abstract neatness. It is to close a practical gap between the real tempo of AI deployment and the availability of trustworthy, independent knowledge about its impacts.

For technical research, the DRC shortens the time between early prototypes, independent replication, and integrated assessment. Instead of waiting several years for scattered papers and occasional static reviews, developers, auditors, and regulators can consult living syntheses that update as new versions, benchmarks, and failure modes emerge.

For applied domains such as education, labor markets, healthcare, and public administration, the combination of versioned research objects, review traces, and living syntheses allows system-level questions to be addressed more quickly and with better contextual grounding. Studies from low and middle income countries, qualitative work on lived experience, and evidence from public services can be surfaced alongside high-profile results from corporate labs. Practitioners and affected communities can see how their settings are represented and have clearer routes to contribute.

In short, the DRC reorients the system from static prestige outputs toward a continuous, open record of inquiry, evaluation, and synthesis. It replaces a black-box pipeline with an architecture in which the evolution of knowledge about AI is visible, contestable, and responsive to changing conditions. Section 6 turns to the micro level. Where the DRC describes the public architecture, the Iterative Human–AI Co-Creation (IHACC) framework provides a concrete model of how researchers and AI tools work together inside that architecture so that acceleration does not come at the cost of judgment, transparency, or pluralism.

6 The Iterative Human–AI Co-Creation (IHACC) framework

The Dynamic Research Continuum provides a public architecture for versioned research objects, review traces, and living syntheses. To operate it in practice, we need a workflow that recognizes AI tools as part of research rather than as an afterthought. The Iterative Human–AI Co-Creation (IHACC) framework serves that role.

In this paper, the Iterative Human–AI Co-Creation (IHACC) model is not offered as a second, free-standing theoretical contribution. Its role is operational: IHACC is the workflow that makes a Dynamic Research Continuum node usable in practice without collapsing into either naïve automation or business-as-usual publishing. It specifies how human experts and AI systems cycle through scoping, drafting, critique, synthesis and disclosure, and it builds in guardrails around provenance, pluralism and human judgement. In other words, DRC defines the architecture and tempo of a reformed research system; IHACC describes how individual teams actually run that system day to day. A fuller treatment of IHACC as a general model of human–AI knowledge production is developed elsewhere (Simpson, 2025); here it is used instrumentally, as the process layer that allows DRC's institutional reforms to work.

6.1 Stages in the IHACC cycle

IHACC organizes research into six recurring stages, applied at different scales from rapid scans to multi-year programs.

1. Orientation. Clarify the question, context, and stakes. Researchers define the problem, identify who is affected, and specify constraints. AI tools can suggest alternative framings and related questions, but humans decide what matters and what would count as success or failure.
2. Mapping. Systematically explore the existing terrain. This includes scoping literature, scanning preprints and datasets, and identifying relevant syntheses. AI acts as a high-capacity scout and summarizer, proposing clusters and gaps. Humans refine the search, filter noise, and decide which lines of evidence deserve deeper work.
3. Design. Decide how to extend or test the knowledge base. This covers research design, method selection, sampling, and ethics, as well as choices about how outputs will enter the DRC. AI can compare design options to prior studies and flag obvious weaknesses. Human researchers weigh trade-offs, account for local constraints, and take responsibility for the final design.
4. Generation and critique. Produce analyses and drafts while actively testing them. AI tools assist with data cleaning, code, robustness checks, and provisional text, tables, or figures. They can also act as critics, proposing counter-arguments and alternative explanations.

Humans decide which outputs to accept, reinterpret machine suggestions in light of domain knowledge, and ensure context is not flattened by generic templates.

5. Validation and stress-testing. Challenge findings before public release. This includes internal replication, sensitivity analysis, triangulation with qualitative evidence, and consultation with practitioners or affected communities. AI can automate parts of checking and search for contradictory evidence. Humans decide whether standards of rigor and clarity about uncertainty have been met.
6. Integration and revision. Publish into the DRC and connect the work to existing loops. This may involve creating or updating a versioned research object, attaching review traces, or contributing to a living synthesis. AI assists with formatting, metadata, cross-linking, and translations. Human teams curate what is released, ensure appropriate credit, and decide when new IHACC cycles are needed in response to critique or new data.

In practice researchers move back and forth between stages, and several cycles may run in parallel. The point is not rigid sequencing but making these moves explicit and repeatable rather than idiosyncratic and invisible.

6.2 Human and AI roles

Within these stages, IHACC distinguishes recurring AI roles. The same model can play several roles, but making them explicit helps to avoid silent delegation.

- Scout. Searches across large corpora of articles, preprints, policy reports, and code repositories to surface potentially relevant material, including low-visibility venues and non-English sources.
- Summarizer. Produces structured summaries of studies, clusters, datasets, or stakeholder inputs using agreed templates.
- Drafting partner. Helps generate provisional text, figures, and tables or alternative formulations of arguments, always under human editing with clear documentation of assistance.
- Pattern spotter. Assists in identifying patterns, anomalies, and gaps in evidence that humans may wish to explore, for example under-studied populations or jurisdictions.
- Red-teamer and critic. Generates questions, counter-arguments, and failure scenarios that test the robustness of methods and conclusions.
- Translator and bridge. Supports adaptation of outputs for different audiences and languages, including policy briefs and practitioner guides, subject to human review.
- Archivist and connector. Helps maintain links between versions, review traces, and syntheses inside the DRC, updating metadata so that provenance and change histories are visible.

Humans define goals, interpret evidence, handle ethical and contextual judgment, and own the final claims. AI roles support breadth, speed, and consistency; they do not replace domain expertise or institutional responsibility.

6.3 Guardrails: disclosure, provenance, pluralism

To contribute to trust rather than erode it, IHACC depends on three guardrails.

First, disclosure. Projects should document where and how AI tools were used at each stage, including which models and prompts were employed for key outputs. Within a DRC implementation, this information is part of the research object's metadata and links to review traces.

Second, provenance and reproducibility. Wherever feasible, code, data, prompts, and version histories should be shared or at least auditable. This includes noting when AI-generated text or analysis was heavily edited or rejected so that others can see how conclusions were formed.

Third, pluralism and contestation. IHACC should not quietly converge everything toward a single model-shaped style of reasoning. The architecture must support qualitative and participatory work, minority perspectives, and local knowledges. That implies engaging practitioners and affected communities in design and validation, and encouraging critical engagement with model outputs rather than passive acceptance.

These guardrails echo wider debates on AI in science but, in IHACC, they are tied to a concrete workflow and to the public architecture of the DRC.

6.4 An IHACC cycle inside the DRC: education example

Consider the example of AI tutors in low-income secondary schools, introduced in Section 5.

In orientation, a research team works with teachers, students, and local officials to define the question: not simply whether AI tutors raise test scores, but how they affect learning, motivation, equity, and teacher workload. AI tools suggest related questions and surface international examples, but the team decides which outcomes matter locally.

During mapping, AI scouts search across arXiv, education journals, and policy reports for trials of AI tutoring systems, classroom pilots, and qualitative studies. Humans review and filter these results, identify genuine gaps, and note that most existing work is from high-income settings.

In design, the team chooses a mixed-methods approach that combines quantitative measures of learning with classroom observations and interviews. AI helps compare design options and simulate power, but the design reflects practical constraints in the schools and ethical requirements around consent and risk.

During generation and critique, AI tools assist with data cleaning, code, and preliminary drafting, and act as critics by suggesting alternative explanations for observed effects and pointing to related evidence. Human researchers interpret these suggestions in light of classroom realities and adjust analyses.

In validation and stress-testing, the team checks robustness, triangulates findings with teacher and student accounts, and uses AI to search for contradictory evidence elsewhere. Humans decide which conclusions are strong and which remain tentative.

In integration and revision, the team releases a structured preprint into the DRC with documentation of AI use and links to code and anonymized data where possible. Review traces build as others respond. A living synthesis on AI tutors in compulsory education incorporates the study and updates its conclusions. As new cohorts are followed or additional schools join, the IHACC cycle repeats and the research object is versioned, with each iteration visible inside the DRC.

In this way, IHACC and the DRC operate together. The DRC supplies an open, versioned architecture; IHACC provides a disciplined way for humans and AI tools to generate and update evidence inside that architecture. The next section connects this combined model to questions of governance, temporal legitimacy, and agenda setting in AI policy.

7 Governance, temporal legitimacy, and the AI impact gap

So far the argument has focused on the internal mechanics of research: how studies are produced, reviewed, and synthesized. The deeper issue is political. In fields shaped by AI, the speed and structure of the knowledge pipeline now determine who can legitimately claim to know enough to set rules, allocate resources, and define acceptable risk. When research moves too slowly and in the wrong format, formal institutions lose their ability to anchor public decisions. That is a problem of temporal legitimacy.

By temporal legitimacy we mean the capacity of institutions to make and justify decisions on the basis of evidence that is not only rigorous and independent, but also timely enough to bear on the technologies and practices actually in play. A regulator who can only cite five year old studies about defunct models, or a ministry that relies mainly on vendor reports because public research is absent, may be procedurally correct on paper but substantively hollow. In AI-intensive domains, the gap between the tempo of model deployment and the tempo of public knowledge production translates into an AI impact gap, where systems reshape labor, education, and administration faster than independent evidence and oversight can respond.

The DRC and IHACC are not attempts to eliminate this gap, which would be unrealistic. They are attempts to shrink it to a scale where democratic institutions, professional standards, and public deliberation can still operate with some integrity.

7.1 Temporal legitimacy in AI governance

Existing governance frameworks assume a background of relatively stable knowledge. Laws, standards, and policies can be updated periodically on the basis of reports, inquiries, and expert reviews that lag reality by years. That assumption breaks when general purpose AI systems are updated quarterly and are immediately embedded into high stakes decisions in finance, education, healthcare, welfare, and policing.

In that setting, temporal legitimacy has at least three components.

First, alignment of cycles. The cadence of serious evidence syntheses has to bear some relation to the cadence of major deployment decisions. If new model generations are being rolled out every six to twelve months, then waiting three to five years for peer reviewed syntheses is tantamount to abandoning the field to vendor narratives and ad hoc commentary.

Second, independence and contestability. Timely knowledge that is wholly produced and interpreted by the same actors who build and sell AI systems does not provide a stable basis for governance. What matters is not just speed but the ability of independent researchers, practitioners, and affected communities to test and challenge claims while they still matter.

Third, visibility of uncertainty and disagreement. In fast moving domains, genuine uncertainty and reasonable disagreement are inevitable. Legitimacy depends on making those uncertainties visible, not on pretending that consensus exists. A system that can show what is known, what is contested, and where evidence is thin is better suited to support careful regulation than one that offers a string of belated, overconfident articles.

The DRC and IHACC speak directly to these elements. They are designed to shorten cycles, widen participation in knowledge production, and make the evolution of claims visible. The question is what that means in practice for specific domains.

7.2 Labor, education, and public administration as test beds

The AI impact gap is most visible in three areas that are central to social contract debates: labor markets, education systems, and public administration.

In labor markets, vendors and consultancies have been quick to publish sweeping projections about job loss, productivity gains, and new occupational categories. Public research on actual displacement, task change, wage effects, and bargaining power often arrives later and in fragmented form. Without timely syntheses that integrate firm level evidence, worker experience, and institutional constraints, labor ministries and social partners are left reacting to narratives framed elsewhere.

In education, platforms and tools promise personalized tutoring, automated grading, and adaptive curricula. School systems face immediate decisions about procurement, pedagogy, assessment, and equity. Evidence of actual learning effects, distributional impacts, and unintended consequences tends to lag behind, scattered across pilots, vendor case studies, and a small number of academic trials. As a result, high stakes judgments about what counts as valid learning and fair assessment are often made with partial or biased information.

In public administration, AI tools are being used for triage, eligibility assessments, fraud detection, and resource allocation. The epistemic and ethical risks here are significant: errors can translate into denial of benefits, wrongful sanctions, or opaque shifts in priorities. Yet independent evaluations of algorithmic systems in welfare, taxation, policing, and migration are few, often retrospective, and jurisdiction specific. Courts and oversight bodies are being asked to rule on systems that have never been subject to systematic, ongoing public evaluation.

In each of these domains, a functioning DRC plus IHACC combination would look different. Labor market syntheses would track not only task level exposure but also institutional responses and bargaining outcomes. Education syntheses would integrate experimental and qualitative evidence and give voice to teachers and students, not just vendors and ministries. Public administration syntheses would connect technical audits with legal analysis and lived impacts, updating as systems are modified or withdrawn.

7.3 Frontier models and system level risk

Frontier models raise the stakes further because their effects can be system level. The issue is not a speculative artificial general intelligence in the abstract, but the cumulative impact of widely deployed models that are repeatedly updated, fine-tuned, and integrated into critical infrastructure and decision pipelines.

From a governance perspective, frontier models stress the publishing architecture in three ways.

- They create moving targets. Capability profiles, failure modes, and deployment patterns can change with each new release, making once accurate assessments obsolete.
- They blur sectoral lines. A single model may affect law, education, media, and security simultaneously, while traditional research and regulation are organized by sector.
- They amplify agenda setting power. In the absence of timely, independent syntheses, policy debates default to vendor reports, benchmark results, and selectively curated case studies.

In such an environment, a journal system that functions as a trailing indicator cannot anchor risk governance. The role of a DRC style architecture is not to certify frontier systems as safe, but to provide a continuously updated, multi-perspective record of what is known about their capabilities, failures, and social impacts, and to make that record legible to regulators, courts, and publics.

7.4 What success would look like

If DRC and IHACC are more than attractive metaphors, they should make a measurable difference to how AI governance operates. Several indicative criteria would signal progress toward closing the AI impact gap.

- Reduced lag between deployment and synthesis. In priority domains such as education, labor markets, and public administration, the median time between major deployment decisions and the appearance of high quality, DRC grade syntheses would fall from several years to something closer to twelve to eighteen months, with incremental updates in between.
- Shift in citation and attention patterns. Key policy documents, regulatory impact assessments, and judicial opinions would increasingly cite DRC based syntheses and visible review traces, not only vendor reports or single high profile studies. Bibliometric work would show a higher share of references to living syntheses and to research from a broader range of institutions and countries.
- Improved coverage and equity. Living syntheses in AI-relevant domains would include evidence from low and middle income settings and from marginalized groups as a matter of routine, not as an afterthought. Gaps would still exist, but they would be documented explicitly rather than hidden by aggregation.
- Transparent audit trails for contested decisions. When disputes arise about AI use, regulators, courts, and affected communities would be able to reconstruct which bodies of evidence were available at the time, how they were synthesized, and where uncertainties or disagreements were acknowledged. That does not guarantee consensus, but it provides a basis for meaningful accountability.

These indicators will not be controlled solely by any one actor. They depend on funders, publishers, universities, professional bodies, and regulators making aligned choices. But they provide a way of assessing whether DRC and IHACC type reforms are doing more than producing new jargon.

7.5 Risks and failure modes

Any attempt to redesign the knowledge pipeline carries its own risks.

One is formalism without substance. Institutions might adopt DRC language and minimal technical infrastructure while leaving incentives, credit, and participation unchanged. The result would be an additional layer of bureaucracy on top of the existing system, with little impact on temporal legitimacy or equity.

Another is capture by incumbents. If only the largest publishers and best resourced institutions can afford to run DRC style platforms and IHACC informed workflows, the architecture could tighten rather than loosen their grip on epistemic authority. AI tools used in scouting and synthesis could be tuned, consciously or not, to amplify familiar voices and topics.

A third is overconfidence in automation. There is a real risk that AI assisted syntheses, especially if presented with clean graphics and authoritative branding, will be treated as more definitive than they are. That would undermine the pluralism and contestation that genuine legitimacy requires.

These are not arguments against the DRC and IHACC. They are reminders that architecture and workflow are necessary but not sufficient conditions for legitimate governance in the AI age. Implementation choices, power relations, and political contestation will decide whether the new structures actually narrow the AI impact gap or merely rename it.

The next section turns to those implementation questions directly. It considers what DRC and IHACC mean for funders, publishers, universities, and regulators, and how reform could proceed under real world constraints rather than in an idealized institutional vacuum.

8 Implementation pathways and practical obstacles

The Dynamic Research Continuum (DRC) and the Iterative Human–AI Co-Creation (IHACC) framework are systemic proposals. They cannot be implemented by individual researchers acting alone, and they will not emerge automatically from existing “open science” reforms. Movement toward a DRC-style architecture requires coordinated shifts by at least four groups of institutions: funders, publishers and platforms, universities and research organizations, and regulators and public bodies. Each group holds different levers, faces different constraints, and will move at different speeds.

This section outlines concrete pathways for each of these actors and then sketches a realistic phased approach. The emphasis is on changes that are materially possible within current institutional structures, rather than on idealized blueprints that assume away political and economic constraints.

8.1 Funders: buying temporal legitimacy

Funders are in the strongest position to change behavior quickly, because they can attach conditions to money. They also have a direct interest in closing the AI impact gap: delayed, fragmented evidence erodes the value of their own investments. The main levers are summarized below.

First, funders can treat DRC-grade outputs as first-class deliverables. Calls in AI-relevant domains can explicitly fund and require: (a) structured preprints with public review traces; (b) contributions to living syntheses; and (c) documentation of IHACC-style human–AI workflows. These outputs should be counted alongside conventional articles in reporting and evaluation.

Second, funders can seed dedicated DRC pilots. This means financing topic-specific nodes (for example on AI in education, labor markets, or public administration) that maintain versioned research objects, open reviews, and living syntheses on shared infrastructure. Grants can support editorial teams, technical staff, and community engagement, not just individual projects.

Third, funders can retool assessment and audit. Panels can be asked to value contributions to review traces and syntheses, not just first-author papers in prestige journals. Periodic audits of funded portfolios can examine how quickly results enter DRC-like loops and how visible they are to practitioners and regulators.

The main obstacles are political and cultural. Funders are risk-averse, and many evaluation systems are still tightly coupled to journal brands and citation metrics. There are also equity questions: if DRC participation becomes a de facto requirement for funding, low and middle income institutions may struggle without targeted support. None of these issues is trivial, but they are about allocation of resources and attention, not about technical feasibility.

8.2 Publishers and platforms: from journals to DRC nodes

Publishers and platform providers are central to any transition because they control much of the infrastructure through which manuscripts, reviews, and metrics currently flow. The point is not to abolish journals, but to re-purpose them as nodes in a DRC architecture rather than as endpoints. Four changes are central.

First, set reviewed preprints as the default. Journals move toward workflows in which peer review is attached to preprints, with public or semi-public review traces. Final "versions of record" become one visible state in an ongoing evolution rather than the only recognized output. Second, ensure support for versioning and linkage. Platforms support explicit version histories, link reviews and commentaries to specific versions, and expose structured metadata that allows DRC objects to be harvested and recombined across sites.

Third, assure integration of living syntheses. Publishers host and recognize maintained syntheses with their own identifiers, editorial teams, and credit structures. These can be linked to, but are not constrained by, any single journal title. Fourth, ensure interoperability and open APIs. To avoid new silos, DRC objects must be discoverable and usable across platforms. That requires common metadata standards and open interfaces that allow funders, universities, and independent aggregators to build on top of publisher systems.

The obstacles are obvious. Established publishers have strong financial incentives to defend their existing business models and brand hierarchies. Some experiment with open review or rapid tracks, but often at the margins. There is also coordination risk: interoperability only works if enough major actors agree on standards. A realistic pathway is to start with consortia and high-priority domains. Major funders and societies can negotiate with a subset of publishers to host DRC-compliant pilots on AI and labor, education, or public administration, pooling risk and cost. If these nodes demonstrate value to authors, readers, and regulators, they create pressure for broader adoption.

8.3 Universities and research organizations: aligning careers with public value

Universities and research organizations control the careers of most researchers. They can reinforce the current publication pipeline or help staff move toward DRC and IHACC practices. There are three particularly important shifts.

First, criteria for hiring, promotion, and tenure need to recognize DRC-aligned work. This includes explicit credit for: leading or contributing to living syntheses; providing high-quality open reviews; maintaining shared datasets and code; and documenting human–AI workflows in ways that improve transparency and reuse. Institutions do not have to abandon journal metrics overnight, but they can downgrade their centrality and upgrade recognition of DRC outputs.

Second, universities can provide time and infrastructure. Participation in DRC nodes and IHACC workflows takes effort: engaging with community partners, contributing to syntheses, and maintaining open artefacts. Departments and centres can allocate formal roles (for example, "synthesis editor," "data custodian," or "AI methods lead") and protect time for them, rather than leaving this work as unpaid labor.

Third, universities can build capacity for human–AI collaboration. Training in IHACC-style workflows, model limitations, and disclosure norms should be treated as part of core research education, not as an optional side course. Internal ethics and governance bodies can set expectations about responsible AI use in research, aligned with DRC principles.

The main obstacles are inertia and perceived risk. Senior scholars who built careers in a journal-centric world may distrust new metrics and formats, and many managers remain focused on league tables and rankings that reward traditional outputs. Early movers risk being penalized if external recognition lags. This is precisely where funders and consortia can help by sending clear signals that DRC-style work will be valued.

8.4 Regulators and public bodies: creating demand for DRC outputs

Regulators, standard-setting bodies, and public agencies do not control the production of research, but they strongly influence which evidence counts in practice. They can either reinforce the current AI impact gap or help close it by creating demand for DRC outputs. Regulators and public bodies can start with thee critical actions.

First, they can state explicit expectations that high-stakes decisions about AI systems in labor, education, healthcare, and public administration should, where possible, be informed by up-to-date, independently produced syntheses. Guidance documents, regulatory impact assessments, and oversight reports can signal a preference for DRC-grade syntheses and visible review traces over vendor white papers or isolated studies.

Second, they can embed DRC references into processes. Regulatory consultation papers, calls for evidence, and judicial guidance notes can point to specific DRC nodes or synthesis programs as relevant sources, while inviting critique and supplementation. Over time, this can normalize the idea that contested AI deployments require reference to a public, versioned evidence record.

Third, regulators and public bodies can co-fund and co-govern living syntheses in critical domains, alongside independent funders and research organizations. This does not mean controlling conclusions. It means providing stable support, access to administrative data, and formal channels for feeding findings into policy cycles.

Obstacles include budget constraints, political pressure, and limited in-house expertise. Many public agencies are already stretched and may be wary of sponsoring infrastructures that could later criticize their own decisions. There is also a risk of capture if DRC nodes become closely tied to particular ministries or regulators. These risks need to be managed through governance design, for example by multi-stakeholder boards and transparency requirements.

8.5 Phased implementation and realistic starting points

DRC and IHACC represent a significant shift in how knowledge is produced and used. Full, simultaneous adoption across all actors is unrealistic. A phased approach is more plausible, with different sectors moving at different speeds but reinforcing one another.

A near-term phase (roughly one to three years) could focus on targeted pilots in clearly defined AI-impact domains:

- Funders commission a small number of topic-focused DRC pilots on AI and labor, education, and public administration, hosted on existing preprint platforms with added review-trace features.
- Selected projects in these domains are required and supported to use IHACC-style workflows, with documentation treated as part of deliverables.
- One or two living syntheses per domain are established and maintained by international consortia, with joint funding from agencies and foundations.

A medium-term phase (three to seven years) would scale and normalize these practices:

- More journals adopt reviewed-preprint workflows and publish review traces in AI-relevant fields.
- Universities integrate DRC/IHACC participation into evaluation criteria, workload models, and research training.
- Regulators and public bodies begin routinely referencing DRC-based syntheses in consultation papers, guidance documents, and oversight reports.
- Shared metadata standards and APIs for DRC objects are agreed by consortia of funders, publishers, and platforms.

A longer-term phase (beyond seven years) would focus on entrenching and protecting the resulting infrastructures:

- Open, interoperable DRC nodes become the default for high-impact AI domains, with stable mixed funding.
- National and international research assessment frameworks recognize DRC-aligned outputs, review labor, and living syntheses as core contributions.
- For high-stakes AI systems, statutory or regulatory frameworks expect reference to DRC-grade evidence, while preserving room for challenge and minority views.

Each step is individually modest. Together, they amount to a reorientation of the knowledge system from static prestige outputs to a visible continuum of inquiry, critique, and synthesis.

8.6 Political economy and resistance

Reforming knowledge infrastructures is never purely technical. DRC and IHACC will encounter organized resistance and subtler forms of drift.

Large commercial publishers have strong incentives to shape any transition in ways that preserve their margins and control over branding and metrics. Preprint platforms and indexers will seek rent-extraction opportunities of their own. Universities may endorse open science in principle while continuing to reward traditional journal counts and rankings in practice. Within disciplines, influential actors may perceive open review and living syntheses as diluting their gatekeeping roles.

Internationally, there is a risk that DRC infrastructures become centred in a handful of wealthy countries and institutions, reproducing the same geopolitical asymmetries that characterize much current AI development. Without deliberate design choices, DRC nodes may mostly reflect concerns and contexts from the global North, with low and middle income settings appearing as occasional case studies rather than as equal partners.

These dynamics cannot be solved by architecture alone. They require coalitions of funders, societies, universities, and civil society organizations willing to push for governance arrangements that include diverse stakeholders, transparent decision making, and strong safeguards against capture.

8.7 Limits and open questions

Finally, it is important to be clear about what DRC and IHACC cannot do. They cannot guarantee that AI systems will be deployed ethically or that political actors will follow the best available evidence. They cannot remove value conflicts or eliminate uncertainty. What they can do is change the default conditions under which judgments are made, by making relevant evidence more timely, traceable, and contestable.

Several open questions deserve further work:

- Funding and governance. How should living syntheses and DRC nodes be funded and governed so that they are robust to political and commercial pressure, especially in contested domains?
- Disciplinary variation. How can DRC and IHACC be adapted to fields with very different norms and artefacts, such as law, the humanities, or design research, without forcing them into a single template?
- Evaluation of human–AI choreography. How should the quality of human–AI collaboration itself be assessed as part of research evaluation, beyond crude bans or uncritical enthusiasm?
- Concentration and control. What safeguards are needed to prevent DRC infrastructures from becoming new sites of centralized control over what counts as "legitimate" knowledge about AI?

The conclusion returns briefly to these themes. Its focus is not on offering guarantees, but on arguing that in an AI-accelerated world, a versioned, openly reviewed, and continuously synthesized research system is a necessary condition for any serious attempt at legitimate governance, even though it will not be a sufficient one.

9. Conclusion

This paper began from a simple but awkward observation. In domains reshaped by contemporary AI, the tempo of model development, deployment, and integration into institutions now runs on a quarterly cycle, while the academic publishing system that is meant to provide independent knowledge still moves on a scale of years. The result is an AI impact gap. Systems alter labor markets, education, and public administration faster than public, peer-reviewed evidence can accumulate, be synthesized, and enter decision making. In that environment, authority over "what is happening" defaults to corporate labs, consultancies, and informal commentary. The problem is not only one of rigor or access, but of temporal legitimacy.

Sections 2 to 4 showed how we arrived here. The journal model was optimized for a print era, scarce page space, and relatively slow-moving technologies. Digital migration kept the same architecture while increasing volume. The combination of high output, weak synthesis, attention inequality, disciplinary silos, and misaligned incentives now interacts with AI in ways that make delay and distortion systemic rather than incidental. AI tools inside research workflows amplify both the potential for acceleration and the risk of opaque error. Frontier models expose the limits of a system that functions largely as a trailing indicator.

The Dynamic Research Continuum (DRC) is proposed as an architectural response. It treats research not as a sequence of isolated, static articles, but as a visible loop of versioned research objects, review traces, and living syntheses. Structured preprints, open or semi-open review, and maintained syntheses are not separate reforms; they are integrated into one lifecycle, with explicit rules for versioning and delta updates. The goal is not speed for its own sake, but a knowledge pipeline that can produce timely, traceable, and revisable assessments in domains where AI deployment is already under way.

The Iterative Human–AI Co-Creation (IHACC) framework (Simpson, 2025) is the corresponding workflow. It organizes research into recurring stages of orientation, mapping, design, generation and critique, validation, and integration, with clearly defined roles for AI tools as scouts, summarizers, drafting partners, critics, translators, and archivists. Humans retain responsibility for problem framing, ethical and contextual judgment, and final claims. IHACC does not try to eliminate AI from research; it makes its involvement explicit, constrained, and auditable.

Taken together, DRC and IHACC address three deficits that currently undermine temporal legitimacy in AI-relevant research.

- They respond to the temporal deficit by shortening the distance between early results, independent critique, and integrated synthesis, and by making updating a normal part of the publication lifecycle rather than an occasional event.
- They respond to the opacity deficit by exposing version histories, review traces, and human–AI divisions of labor, so that downstream users can see how claims were produced, challenged, and revised.
- They respond to the agenda deficit by creating structures in which evidence from a wider range of institutions, jurisdictions, and affected communities can be surfaced and integrated, rather than relying solely on outputs from a small set of well-resourced actors.

Sections 7 and 8 argued that these changes are not only technical but institutional. Funders, publishers and platforms, universities, and regulators each hold different parts of the system in place and each would have to move. The paper sketched a phased path: targeted pilots in high-impact domains such as labor, education, and public administration; gradual normalization of reviewed preprints, living syntheses, and IHACC-style workflows; and, over time, embedding DRC-grade evidence into regulatory and assessment practices. None of these steps is individually dramatic. Cumulatively, they amount to a shift from a prestige-oriented article economy toward a public service model of knowledge.

There are obvious risks. DRC rhetoric could be adopted without changing incentives or credit structures. New infrastructures could be captured by incumbents and concentrated in a few countries. AI-assisted syntheses could be granted more authority than they merit. These are reasons to treat implementation as a matter of governance and power, not as a purely technical upgrade. They are not reasons to resign the field to existing pipelines whose limitations are already visible.

The claim of this paper is deliberately modest and deliberately hard. A versioned, openly reviewed, and continuously synthesized research system, operated through explicit human–AI workflows, is not sufficient for legitimate AI governance. It will not resolve value conflicts, prevent all harms, or ensure that political actors follow the best available evidence. But without something like the DRC and IHACC, institutions will continue to rely on slow, opaque, and unevenly distributed knowledge in a context where the systems they are meant to oversee are evolving much faster. Under those conditions, claims to govern in the public interest will become steadily less credible.

If we take seriously the idea that AI is now part of the basic infrastructure of social and economic life, then the way we produce and organize knowledge about it is not a secondary technical matter. It is a core component of the social contract. The Dynamic Research Continuum and the IHACC framework are offered as one concrete route toward re-aligning that contract with the realities of the AI age: not by promising perfect control, but by raising the floor under how quickly, transparently, and inclusively we can learn what these systems are doing to us. Temporal legitimacy is not an abstract property of institutions; it shows up in whether people experience decisions as informed and fair. In labor markets, slow and fragmented evidence makes it easier to justify automation strategies that concentrate gains and externalize risks onto workers. In education, curricular and funding choices can lock cohorts of students into skills that were already obsolete when the relevant reports were written. In public administration, outdated syntheses give cover to risk-averse or politically convenient choices, even when better evidence exists in preprints and grey literature. A Dynamic Research Continuum is therefore not only a technical upgrade to the knowledge pipeline, but a condition for sustaining trust that governments, employers and universities are making choices on the basis of the best available knowledge in real time.

AI Use Statement

During the preparation of this work the author used ChatGPT 5.1 in order to refine grammar and carry out formatting and referencing consistency checks. The author takes full responsibility for the content of the publication.

References

Cannabis Evidence. (n.d.). *Living systematic review methods*. Cannabis Evidence. https://www.cannabisevidence.org/about/living-systematic-review-methods/
cOAlition S. (2023). *Towards responsible publishing: A proposal from cOAlition S*. Plan S. https://www.coalition-s.org/wp-content/uploads/2023/11/2023_11_09_Responsible_Publishing_proposal.pdf

cOAlition S. (2024, October 28). *Peer-reviewed preprints and the Publish–Review–Curate model*. Plan S. https://www.coalition-s.org/blog/peer-reviewed-preprints-and-the-publish-review-curate-model/
eLife. (2023, February 14). *2023 eLife Community Ambassadors: New year, new model*. Inside eLife. https://elifesciences.org/inside-elife/9433fc59/2023-elife-community-ambassadors-new-year-new-model
Elliott, J. H., Synnot, A., Turner, T., Simmonds, M., Akl, E. A., McDonald, S., ... Thomas, J. (2017). Living systematic review: 1. Introduction—the why, what, when, and how. *Journal of Clinical Epidemiology, 91*, 23–30. https://doi.org/10.1016/j.jclinepi.2017.08.010
F1000Research. (n.d.). *Preparing a living systematic review article*. F1000Research. https://f1000research.com/for-authors/article-guidelines/living-systematic-reviews
Garfield, E. (2006). The history and meaning of the journal impact factor. *JAMA, 295*(1), 90–93. https://doi.org/10.1001/jama.295.1.90
Garner, P., Hopewell, S., Chandler, J., MacLehose, H., Schünemann, H. J., Akl, E. A., ... Panel for Updating Guidance for Systematic Reviews. (2016). When and how to update systematic reviews: Consensus and checklist. *BMJ, 354*, i3507. https://doi.org/10.1136/bmj.i3507
Hosseini, M., & Horbach, S. P. J. M. (2023). Fighting reviewer fatigue or amplifying bias? Considerations and recommendations for the use of ChatGPT and other large language models in scholarly peer review. *Research Integrity and Peer Review, 8*, 4. https://doi.org/10.1186/s41073-023-00133-5
Masum, H., Rao, A., Good, B. M., Todd, M. H., Edwards, A. M., & Chan, L. (2013). Ten simple rules for cultivating open science and collaborative R&D. *PLOS Computational Biology, 9*(9), e1003244. https://doi.org/10.1371/journal.pcbi.1003244
MDPI. (2021, October 27). *MDPI submission process: Your questions answered*. MDPI Blog. https://blog.mdpi.com/2021/10/27/mdpi-submission-process-your-questions-answered/
Naddaf, M. (2024). AI workloads demand new metrics. *Nature, 627*(8002), 20–22.
Powell, K. (2016). Does it take too long to publish research? *Nature, 530*(7589), 148–151. https://doi.org/10.1038/530148a
Royal Society. (2017). *Transparency in peer review*. Royal Society. https://royalsociety.org/topics-policy/projects/peer-review/transparency/
Simpson, E. (2025). Flow acceleration and the IHACC model: Human–AI co-creation in epistemology. *SSRN*. https://doi.org/10.2139/ssrn.5742245
Sonnenburg, S., Braun, M. L., Ong, C. S., Bengio, S., Bottou, L., Holmes, G., ... Williamson, R. C. (2007). The need for open source software in machine learning. *Journal of Machine Learning Research, 8*, 2443–2466.
Squazzoni, F., Bravo, G., Grimaldo, F., García-Costa, D., Farjam, M., & Mehmani, B. (2021). Peer review and journal policy are the two main pillars of editorial decision-making at specialty journals. *Science Advances, 7*(2), eabd0299. https://doi.org/10.1126/sciadv.abd0299
Su, H., Sun, Y., Li, R., Zhang, A., Yang, Y., Xiao, F., Duan, Z., Chen, J., Hu, Q., Yang, T., Xu, B., Zhang, Q., Zhao, J., Li, Y., & Li, H. (2025). Large language models in medical diagnostics: Scoping review with bibliometric analysis. *Journal of Medical Internet Research, 27*, e72062. https://doi.org/10.2196/72062
Suber, P. (2012). *Open access*. MIT Press.
Tennant, J. P., Waldner, F., Jacques, D. C., Masuzzo, P., Collister, L. B., & Hartgerink, C. H. J. (2016). The academic, economic and societal impacts of Open Access: An evidence-based review. *F1000Research, 5*, 632. https://doi.org/10.12688/f1000research.8460.3
Vale, R. D. (2015). Accelerating scientific publication in biology. *Proceedings of the National Academy of Sciences, 112*(44), 13439–13446. https://doi.org/10.1073/pnas.1511912112