**Flow Acceleration and the IHACC Model: Human–AI Co-Creation in Epistemology**

**Ewan Simpson**

**Narxoz Business School**

**November 13 2025**

**Abstract**

This paper introduces the Iterative Human–AI Co-Creation (IHACC) model as a conceptual framework for knowledge production in the age of generative AI. IHACC describes a recursive process in which human agents, AI systems and feedback loops jointly generate, test and refine ideas. At its core is "flow acceleration": the AI-mediated compression of the classical data–information–knowledge–wisdom progression into rapid human–AI interaction cycles, while leaving humans in charge of epistemic goals, standards and responsibility. Drawing on actor–network perspectives and the notion of the infosphere, IHACC treats knowledge not as a static product but as an emergent property of socio-technical processes, even as it acknowledges Heideggerian concerns about reduction of the world to mere resource. The model synthesizes empirical findings on productivity gains and error risks from large language models, interpreting them as evidence of reconfigured epistemic agency and temporality rather than simple efficiency effects. The paper argues that IHACC is a candidate paradigm shift in epistemology because it relocates the locus of "knowing" into governed human–AI systems and opens new possibilities for epistemic justice, while also posing normative questions about dependence, opacity and control in accelerated knowledge infrastructures. We position IHACC as a base model for accelerated human–AI inquiry and specify propositions that invite direct testing.

**Keywords**: epistemology; human–AI collaboration; flow acceleration; DIKW; actor-network theory; infosphere

## 1. Introduction

The rapid diffusion of generative AI, particularly large language models, is reshaping the landscape of knowledge production and straining the epistemological frameworks that have long structured how humans understand and create knowledge. Classical and social epistemologies presuppose relatively slow, human-paced movement from data, through structured understanding, to more reflective forms of judgment (Descartes, 1641/1996; Kant, 1781/1998; Kuhn, 1962; Ackoff, 1989; Rowley, 2007). Inquiry unfolds as a gradual process constrained by human memory, attention and institutional tempo. Human agents and institutions are treated as the primary locus of judgment, validation and responsibility.

Generative AI disrupts this temporal and structural picture. A single interaction can now surface candidate literatures, generate alternative framings, simulate critiques and draft provisional syntheses. Tasks previously separated by time, expertise and organizational boundaries are compressed into one interface. What changes, therefore, is not only the speed of output but the temporal structure of inquiry and the distribution of epistemic labor ('knowledge work') between humans and machines. AI's capacity to synthesize patterns and iterate with users compresses much of the traditional progression into tight cycles, raising questions about who or what counts as the "knower", how epistemic legitimacy is established and how ethical responsibility is allocated.

This paper introduces the Iterative Human–AI Co-Creation (IHACC) model as a response to those questions. IHACC treats knowledge production as a recursive process in which human agents, AI systems and structured feedback loops jointly generate, test and refine ideas. At the core of the model is "flow acceleration": the AI-mediated compression of multiple stages of inquiry into rapid human–AI interaction cycles. Within such cycles, data, candidate interpretations, provisional claims and higher-order reflection can be generated and revised within a single extended episode of engagement. The human role does not disappear, but it changes. Rather than acting solely as originator and final arbiter of content, the human becomes an orchestrator of prompts, constraints and evaluations inside a highly accelerated epistemic environment.

This reframing is both descriptive and normative. Descriptively, it suggests that the locus of "knowing" may no longer reside exclusively in individual subjects or even in human communities, but in governed human–AI systems whose dynamics unfold over time. Normatively, it forces a reconsideration of what responsible agency looks like when acceleration and delegation are built into the architecture of inquiry. Heidegger's critique of technology as an enframing that reduces beings to a "standing reserve" warns that mechanization may strip epistemic practices of reflective depth (Heidegger, 1977). By contrast, actor–network perspectives and information-centric accounts portray knowledge as emerging from networks of human and non-human actors within a digitally saturated environment (Latour, 2005; Floridi, 2014). IHACC is proposed as a way of reconciling these tensions: a model that acknowledges the risks of reduction while treating AI as a co-constitutive epistemic partner within a networked infosphere.

The argument proceeds in three steps. First, the paper revisits selected strands of modern epistemology and philosophy of technology to locate the present moment, contrasting classical subject-centered views, communal paradigms and critical reflections on technology with networked, information-theoretic perspectives (Descartes, 1641/1996; Kant, 1781/1998; Kuhn, 1962; Heidegger, 1977; Latour, 2005; Floridi, 2014). Second, it develops IHACC and the concept of flow acceleration as a structured description of how human goals, AI capabilities and feedback mechanisms interact in practice, emphasizing the distribution of agency, the compression of temporal cycles and the relocation of

knowledge from static products to evolving processes. Third, it interprets recent empirical work on generative AI's effects on productivity, error patterns and distributional outcomes as evidence of reconfigured epistemic agency and temporality rather than mere efficiency gains (Dell'Acqua et al., 2023; Noy & Zhang, 2023; Brynjolfsson, Li, & Raymond, 2023).

On this basis, the paper argues that IHACC marks a significant epistemic reconfiguration. It suggests that the primary site of "knowing" in an AI-augmented reality is the governed human–AI system rather than the isolated human subject or institution. This shift opens possibilities for broader participation in knowledge work yet sharpens existing concerns about dependence, opacity and epistemic injustice highlighted by critical AI scholarship (Noble, 2018; Crawford, 2021; Coeckelbergh, 2020). Section 2 situates IHACC within historical and contemporary debates on knowledge, technology and networks; Section 3 presents the IHACC model and the notion of flow acceleration in more formal terms; Section 4 uses empirical findings on generative AI to illustrate how IHACC helps make sense of observed productivity and error patterns, as well as their distributional effects; and Section 5 draws together the philosophical and empirical threads, assessing the implications of accelerated, AI-mediated knowledge production. We position IHACC as a foundational lens for subsequent empirical tests and policy design in human–AI knowledge work

A note on scope: IHACC models accelerated, human-framed knowledge work where tasks admit iteration, interpretation, and qualitative evaluation (analysis, synthesis, drafting, exploratory coding) and where humans retain goal-setting and normative control. It is not a claim about safety-critical pipelines with hard real-time constraints, fully codified transactional processes that can be automated end-to-end, or tasks whose quality is exhausted by a single deterministic answer. IHACC's predictions are most informative when (i) process observability is feasible (prompts, versions, review gates), (ii) outputs can be blind-rated for quality, and (iii) workers can vary loop granularity as a design choice.

## 2. Philosophical Context

The Iterative Human–AI Co-Creation (IHACC) model sits against a long tradition that has treated knowledge as the outcome of relatively slow, human-paced processes. Classical epistemology assumes staged progression from data, through structured understanding, to more reflective forms of judgment (Ackoff, 1989; Rowley, 2007; Descartes, 1641/1996; Kant, 1781/1998; Kuhn, 1962). These models presuppose that cognitive labor is constrained by human memory, attention and institutional tempo; depth is purchased at the cost of speed. They therefore struggle to accommodate an environment in which large language models can compress many of these stages into a single extended episode of human–AI interaction.

Heidegger's philosophy of technology provides a powerful, if ambivalent, frame for this disruption. In his account, modern technology "enframes" beings as a standing reserve, ready for extraction and control, and thereby threatens to hollow out more reflective, meaningful relations to the world (Heidegger, 1977). Applied to AI-mediated inquiry, this worry translates into the risk that cognition itself becomes mechanized: knowledge work reduced to the rapid generation of outputs optimized for instrumental criteria. Yet Heidegger also insists on the possibility of a "free relation" to technology in which humans remain capable of disclosing and questioning rather than merely consuming. IHACC is designed precisely in this space of tension: it treats AI as a powerful epistemic resource while insisting that human agents retain the role of setting goals, evaluating outputs and holding systems answerable to standards of wisdom rather than mere efficiency.

More recent accounts of knowledge and information systems offer resources for articulating this position in positive terms. Latour's actor–network theory (ANT) reframes knowledge as the emergent product of heterogeneous networks of human and non-human actors, from instruments and datasets to institutions and now AI systems themselves (Latour, 2005). Floridi's notion of the "infosphere" similarly portrays contemporary reality as an informational environment in which human and artificial agents co-construct an ever-denser web of data and meanings (Floridi, 2014). Taken together, these perspectives legitimize treating AI not simply as a tool but as a participant in epistemic networks. IHACC builds directly on this move: it conceptualizes the locus of "knowing" as the governed human–AI system operating within the infosphere, thereby aligning flow acceleration with a networked, relational epistemology rather than with an individualist ideal.

The proliferation of practical human–AI collaboration models underscores the need for such a reframing. Following Lakhani (2024), we use centaur for a task-splitting strategy (humans decide what to hand to AI versus retain) and cyborg for a tightly integrated co-editing strategy (human and AI interleave at a fine grain). (Lakhani, 2024) Mollick's "co-intelligence" framework emphasizes synergies that boost creativity and productivity, treating AI as an amplifier of existing human capabilities (Mollick, 2024). Holzinger's work on "human-in-the-loop" learning focuses on domains where safety, interpretability and regulatory oversight demand continuous human supervision of algorithmic systems (Holzinger, 2016). All three perspectives have proved influential in design and management practice, but their primary focus is operational: who should do what, and under which governance arrangements. By contrast, IHACC appropriates their strongest elements: division of labor; synergistic interaction; and structured oversight, but shifts the unit of analysis from individual tasks to the epistemic process itself, asking how cycles of human–AI interaction reconfigure agency, temporality and the very status of knowledge.

In this sense, IHACC is not presented as a repudiation of historical epistemology or of existing collaboration metaphors, but as a synthesis and extension that makes the epistemic stakes explicit. Heidegger's critique keeps the risks of reduction and loss of meaning squarely in view; Latour and Floridi supply a vocabulary for understanding AI as a co-constitutive actor in informational networks; centaur, co-intelligence and human-in-the-loop models demonstrate that hybrid arrangements are already operative in practice. IHACC weaves these strands into a single philosophical backdrop: knowledge production as a networked, recursive process in which accelerated human–AI loops generate, test and refine claims under human-set norms. On that basis, Section 3 turns from context to construction, formalizing the IHACC model and the notion of flow acceleration in more precise terms.

## 3. Flow Acceleration and the IHACC Model

This section develops the core conceptual apparatus of the paper. It first defines flow acceleration as a specific form of AI-mediated compression of the data–information–knowledge–wisdom (DIKW) progression (Ackoff, 1989; Rowley, 2007). It then sets out the structure of the Iterative Human–AI Co-Creation (IHACC) model as a recursive human–AI system operating within this accelerated environment. Finally, it positions IHACC relative to existing collaboration metaphors in order to clarify its distinctive epistemic focus.

### 3.1 Flow acceleration as DIKW compression

In classical accounts, movement from data to wisdom is staged and temporally extended. Data must be collected and cleaned; information extracted through structuring and analysis; knowledge stabilized by validation, replication and integration into a disciplinary corpus; and wisdom associated with reflective judgment about meaning and consequences (Ackoff, 1989; Rowley, 2007). Each step is constrained by human cognitive limits, institutional procedures and the slow circulation of texts and arguments.

Generative AI alters this dynamic by making it possible to traverse large parts of the DIKW hierarchy within a single extended human–AI interaction. A user can supply minimal input data (e.g., a prompt, a document fragment or a dataset), request structured summaries, ask for critical evaluations or alternative framings, and iterate repeatedly in real time. The system can generate synthetic literature reviews, propose hypotheses, simulate objections and produce candidate interpretations that the human can accept, reject or modify.

"Flow acceleration" is the term used here to describe this AI-mediated compression of the DIKW progression into tightly coupled interaction cycles. The term captures three related shifts. First, throughput increases: more candidate interpretations and formulations can be generated per unit time

than would be feasible manually. Second, iteration speed increases: cycles of proposal, critique and revision are shortened, allowing epistemic moves that would previously have taken weeks or months to unfold in hours. Third, the temporal structure of inquiry changes: stages that were once sequentially ordered now appear as overlapping operations within a single loop. Data, information and knowledge are not abandoned as categories, but their boundaries become more porous in practice as AI systems help to generate, transform and recombine them on demand.

Flow acceleration is therefore not simply "more speed" applied to an otherwise unchanged process. It reconfigures how the stages associated with DIKW are combined and revisited. The same interaction can involve the generation of new information from data, the articulation of provisional knowledge claims, and higher-order reflection on their meaning. Human agents remain responsible for epistemic standards, but they now operate within an environment in which many of the intermediate steps are instantiated through rapid, AI-supported transformations.

### 3.2 The structure of the IHACC model

Within this accelerated landscape, the IHACC model conceptualizes knowledge production as a recursive system involving three principal components: (1) human intentional frames via the exercise of agency, (2) AI processing capabilities and (3) structured feedback loops through multiple iterations. Figure 1 illustrates conceptually the stages of two iterations. In reality the number of iterations will be far greater.

**Figure 1: The IHACC Model – Accelerating Knowledge Production through Human-AI Co-Creation**

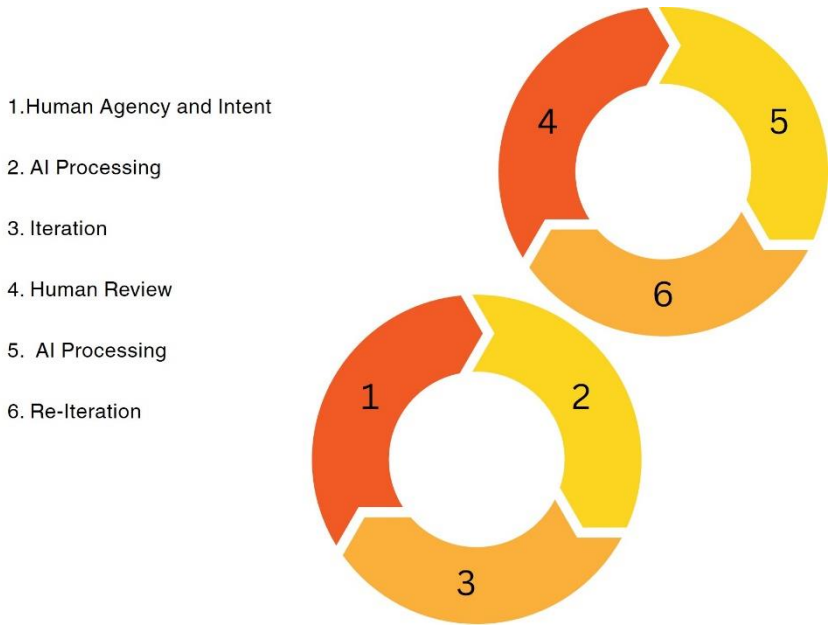

First, human intentional frames encompass the goals, questions, values and background commitments that initiate and guide inquiry. These include explicit research questions, constraints on acceptable evidence and implicit disciplinary norms. In IHACC, the human agent is responsible for setting and revising these frames over time.

Second, AI capabilities refer to the specific operations that generative models can perform: pattern recognition, text generation, transformation, summarization, code production and so forth. In the IHACC model, these capabilities are not treated as independent sources of knowledge but as operators that transform inputs provided by humans and by prior cycles of the loop.

Third, structured feedback loops describe the temporal pattern through which human frames and AI operations are linked. A simple IHACC cycle can be represented as:

1. The human specifies an intentional frame (goal, question, constraint) and provides initial inputs (data, text, examples).
2. The AI system generates candidate outputs consistent with that frame (structures, summaries, explanations, options).
3. The human evaluates these outputs against epistemic and practical criteria, revises the frame where necessary and issues refined prompts or corrections.
4. The cycle repeats, with each iteration modifying both the content under consideration and, potentially, the guiding frame.

Over time, extended sequences of such cycles produce emergent knowledge: not because the AI "possesses" knowledge in a substantive sense, but because the human–AI system jointly constructs, tests and stabilizes claims. What is distinctive in IHACC is the emphasis on the loop as the primary epistemic unit. The model locates "knowing" in the governed dynamics of human framing, AI transformation and human evaluation, rather than in any single node. This resonates with networked accounts of knowledge in actor–network theory and the infosphere (Latour, 2005; Floridi, 2014), while preserving a clear role for human agents as bearers of responsibility and meaning.

The IHACC model can thus be schematically described as a function on sequences of interactions rather than on isolated acts. Let H denote human frames and evaluations, A denote AI transformations and F denote the feedback relation linking them. An IHACC process is a sequence $\{(H_0, A_0), (H_1, A_1), \ldots, (H_n, A_n)\}$ generated under F, where each $H_i$ depends on prior $H_j$, $A_j$ and external inputs, and each $A_i$ depends on $H_i$ and the model's internal state. Knowledge claims emerge when patterns in this sequence stabilize under human-endorsed criteria (e.g., coherence, empirical adequacy, normative acceptability). The

formalism is intentionally minimal, but it highlights that in IHACC the unit of analysis is an evolving trajectory rather than a single decision or output.

Flow acceleration enters this model by shortening the temporal distance between successive elements of the sequence and by increasing the density of transformations $A_i$ that can be explored for a given investment of human time. Because AI can produce multiple candidate outputs in parallel, human agents can interrogate and refine more potential pathways through the space of possibilities than would otherwise be feasible. The epistemic challenge is to ensure that this accelerated exploration remains tethered to reliable standards and does not simply amplify superficial plausibility or existing biases.

### 3.3 Relation to existing collaboration metaphors

IHACC is deliberately positioned in relation to existing human–AI collaboration models such as the "centaur", "co-intelligence" and "human-in-the-loop" paradigms. Lakhani's centaur model emphasizes a division of labor in which humans handle tasks requiring strategic judgment, creativity and complex social understanding, while AI systems handle computation-heavy or routine tasks (Lakhani, 2024). Mollick's co-intelligence framework highlights complementarities that arise when humans and AI systems are combined to enhance performance and creativity, treating the pair as a team whose combined output exceeds what either could achieve alone (Mollick, 2024). Holzinger's work on human-in-the-loop learning stresses the need for continuous human oversight in high-stakes domains, particularly to ensure interpretability, safety and compliance (Holzinger, 2016).

IHACC incorporates elements from all three. Like the centaur model (Lakhani, 2024), it recognizes that different tasks are differentially suited to human or AI capabilities, and it assumes that effective collaboration requires thoughtful allocation of work. Like co-intelligence, it emphasizes the possibility of genuinely generative synergies in which AI helps humans to see patterns, possibilities and framings they might not otherwise have considered. Like human-in-the-loop approaches, it regards human oversight not as optional but as constitutive of legitimate practice.

The distinctive contribution of IHACC lies in its epistemic reframing. Whereas centaur, co-intelligence and human-in-the-loop models are primarily operational and managerial in focus, concerned with who should do what under which governance arrangements, IHACC foregrounds the structure of the knowledge-producing process itself. It treats the recursive loop of framing, transformation and evaluation as the primary object of analysis and asks how flow acceleration changes the distribution of agency, the temporality of inquiry and the locus of epistemic authority. In that sense, IHACC can be seen as an epistemological counterpart to the more practice-oriented collaboration metaphors: it provides a

language for describing what is happening, in epistemic terms, when those operational arrangements are instantiated in an AI-augmented environment.

### 3.4 Epistemic implications of the model

Taken together, flow acceleration and the IHACC model suggest three principal shifts. First, agency becomes hybrid. Human agents remain responsible for initiating, steering and ratifying inquiry, but AI systems play a direct role in proposing, structuring and recombining candidate contents. The epistemic subject is therefore neither the human nor the machine alone but the governed human–AI system operating over time.

Second, temporality is compressed. The cycles through which data are interpreted, claims are formulated and objections are tested become shorter and more numerous. This can support more extensive exploration and critical refinement, but it also risks encouraging superficial engagement if human involvement collapses into mere acceptance of plausible outputs under time pressure.

Third, the locus of knowledge shifts from static products to evolving processes. In an IHACC setting, what counts as "known" is less a single settled statement and more the stability of a trajectory of human–AI interactions under scrutiny and constraint. This aligns with networked and processual accounts of knowledge (Latour, 2005; Floridi, 2014), while keeping Heideggerian concerns in view: without a "free relation" to the technology, there is a real risk that accelerated loops simply mechanize cognition and hollow out the very wisdom that the DIKW hierarchy places at its apex (Heidegger, 1977).

Section 4 turns to recent empirical studies of generative AI and uses the IHACC model to interpret their findings. The aim is not to test IHACC in a narrow sense, but to show that flow acceleration and hybrid agency provide a coherent lens through which to understand observed productivity gains, error patterns and distributional effects.

## 4. Empirical Insights Supporting IHACC

The conceptual framework developed so far gains traction when read alongside emerging empirical work on generative AI in applied settings. Rather than treating these studies as tests in a narrow sense, this section uses them as lenses on three dimensions of accelerated, hybrid inquiry: (a) the speed and structure of work, (b) the redistribution of epistemic agency and access, and (c) the ethical and experiential burdens of AI integration. The claim is that the reported patterns are better understood as changes in epistemic process and locus of "knowing" than as simple productivity effects.

### 4.1 Acceleration of inquiry

A first cluster of findings concerns the speed and quality of task performance when workers are supported by large language models. Dell'Acqua et al. (2023) report that consultants using AI tools complete complex writing and analysis tasks more quickly and with higher-rated outputs than control groups without AI. Brynjolfsson, Li and Raymond (2023) find that novice customer support agents assisted by AI achieve substantial gains in both productivity and quality, converging towards the performance of more experienced peers. Similar patterns appear in other domains, where generative AI is reported to increase throughput while maintaining or improving perceived quality. A field experiment with professional consultants explicitly names and describes "Centaur" and "Cyborg" practices and links them to performance differences on complex tasks. We adopt those labels for observed workflow patterns. (Dell'Acqua et al., 2023)

Recent work on AI agents reinforces this picture while complicating simple automation narratives. Wang et al. (2025) compare the workflows of 48 human workers with those of LLM-based agents across sixteen realistic tasks in data analysis, engineering, computation, writing and design. Agents are substantially faster and cheaper but produce lower-quality outputs and frequently fabricate or misuse tools. Human–AI "augmentation" workflows yield moderate speed gains, whereas "automation" workflows tend to slow humans down because of the additional verification and debugging burden placed on human workers (Wang et al., 2025). Complementary studies of human–AI workflows in user research show that the most effective patterns often involve humans structuring and grouping data while AI systems handle summarization and drafting (Shin et al., 2024). From the perspective developed here, these results illuminate flow acceleration rather than undermining it. In all of these studies, participants do not simply automate a single step in an existing pipeline. They engage in rapid sequences of prompting, evaluation, revision and re-prompting: supplying outlines, data or drafts; requesting alternative framings or summaries; inspecting suggestions; and iterating in real time. Tasks that would traditionally be separated across roles and time—search, synthesis, drafting, basic critique—are compressed into a single extended episode of interaction. Acceleration attaches to well-structured interaction loops in which humans retain framing and evaluative roles; naive automation that sidelines human oversight simply displaces effort into downstream checking.

Crucially, human evaluation remains central. Consultants and agents are responsible for steering the interaction, setting standards and deciding which outputs to accept or discard. In Heideggerian terms, the system does not necessarily "enframe" them as passive consumers of machine-generated content; they can maintain a "free relation" to the technology by interrogating, redirecting and constraining it (Heidegger, 1977). The studies therefore support the claim that generative AI modifies not just output

levels but the temporal structure of inquiry and the distribution of cognitive labor within accelerated loops of interaction.

### 4.2 Democratization and epistemic access

A second theme concerns who benefits most from AI support. Noy and Zhang (2023) show that less-skilled writers assisted by generative AI produce higher-quality texts than control participants, while also reporting greater task enjoyment. Brynjolfsson et al. (2023) similarly find that novice customer support workers gain more from AI assistance than experienced staff. In both cases, AI narrows performance gaps by lifting the lower end of the distribution rather than primarily augmenting already high performers.

Viewed through the present framework, these findings speak to the redistribution of epistemic agency. If less-experienced actors can, with AI support, participate more effectively in tasks that previously required specialist expertise, the composition of the epistemic community can expand. Latour's actor–network theory anticipates this by emphasizing heterogeneous assemblages of human and non-human participants in the production of knowledge (Latour, 2005). Floridi's account of the infosphere likewise suggests that digital technologies can widen the set of agents able to contribute meaningfully to informational processes (Floridi, 2014). Flow-accelerated interaction loops help explain how this happens in practice: novices draw on modeled patterns, structures and framings, but still apply their own contextual knowledge and values in evaluation and selection.

This has implications for epistemic justice. Critical studies of algorithmic systems show how data infrastructures and search engines can reproduce or intensify existing inequalities, marginalizing already under-represented voices (Noble, 2018; Crawford, 2021). The Noy and Zhang (2023) and Brynjolfsson et al. (2023) results point in a more hopeful direction: when designed and governed appropriately, hybrid human–AI systems can lower barriers to participation in knowledge-intensive work. The model developed here captures this as a shift in who can occupy the role of "epistemic agent" within accelerated loops. It does not guarantee just outcomes—data, models and institutional settings remain contested—but it clarifies the mechanisms through which AI might support, rather than erode, more inclusive epistemic practices.

### 4.3 Risks, burdens and ethical constraints

The same studies expose tensions that reinforce the normative concerns running through the earlier analysis. Brynjolfsson et al. (2023) note that gains are uneven: highly skilled workers sometimes experience smaller improvements or even performance reductions when AI suggestions conflict with

established practices. Dell'Acqua et al. (2023) observe that AI assistance can lead to over-reliance or anchoring on plausible but flawed outputs. Pew Research Center (2025) finds that a third of U.S. workers feel overwhelmed about future workplace AI, and among those using AI chatbots, the tools are seen as better at speeding work than improving quality, a pattern consistent with added human oversight to check, correct, and contextualize machine output.

Workflow-level evidence sharpens this point. Wang et al. (2025) find that when humans remain responsible for checking and debugging agent outputs, overall task time can increase even when the automated segments of the workflow are substantially faster. On the present model, this is not an anomaly but a structural feature: accelerated loops that lack adequate human framing and division of labor simply displace effort into downstream verification, undermining both epistemic reliability and working conditions.

These patterns resonate with critical accounts of algorithmic systems. Crawford (2021) documents how opaque, data-hungry AI architectures can entrench power imbalances and shift burdens onto workers tasked with cleaning, labeling and monitoring data. Noble (2018) shows how search and recommendation systems can encode racial and gender biases, undermining epistemic trust. Coeckelbergh (2020) argues for a relational ethics that treats human–AI interactions as ongoing negotiations of responsibility and vulnerability rather than one-off design choices.

Within the present model, these concerns appear as constraints on legitimate flow acceleration. Human oversight is not an optional add-on but constitutive of acceptable practice. The emphasis on recursive feedback and human framing aligns with "human-in-the-loop" approaches (Holzinger, 2016) while making the epistemic stakes explicit: without sustained, critical human engagement, accelerated loops risk collapsing into mechanization and bias amplification. Survey evidence indicates that a sizable share of workers feel overwhelmed by workplace AI, consistent with added verification and oversight burdens (Pew Research Center, 2025). Heidegger's warning about the reduction of beings to standing reserve reappears here as the risk that human knowers themselves become resources for propping up automated infrastructures (Heidegger, 1977).

### 4.4 From productivity to epistemic reconfiguration

Taken together, these empirical insights support a reading of generative AI as a driver of epistemic reconfiguration rather than a mere source of efficiency gains. Acceleration studies show that substantial portions of inquiry can be traversed within single interaction loops under human guidance (Dell'Acqua et al., 2023; Brynjolfsson et al., 2023; Peng et al., 2023). Distributional findings indicate that novices and

less-skilled actors can, with AI support, participate more fully in knowledge-intensive work, aligning with networked and infospheric accounts of agency (Noy & Zhang, 2023; Latour, 2005; Floridi, 2014). Critical work on bias, opacity and workload demonstrates that these gains are contingent and fragile, reinforcing the need for governed human–AI systems that prioritize transparency, accountability and relational ethics (Crawford, 2021; Noble, 2018; Coeckelbergh, 2020).

On this view, knowledge increasingly emerges from dynamic interactions between human framing and AI transformation within accelerated loops. It is not simply a product possessed by either humans or machines, but a property of trajectories of interaction that remain under human normative control. At the same time, the limitations of the current empirical base are evident: most studies are situated in specific organizational contexts in the Global North and were not designed to test these claims directly. Future work could examine similar dynamics in other cultural and institutional settings, and across different epistemic traditions, to probe the scope conditions of accelerated, hybrid systems. Section 5 draws out the philosophical implications of this shift and asks how far it warrants revising our account of contemporary epistemology.

## 5. Discussion and Further Steps

### 5.1 Is IHACC a Paradigm Shift?

This paper began from a simple observation with far-reaching consequences: generative AI compresses the temporal and structural pathways through which data are turned into understanding and, more rarely, into reflective judgment (Ackoff, 1989; Rowley, 2007). Classical and social epistemologies assumed that this progression unfolded at human pace and under human control (Descartes, 1641/1996; Kant, 1781/1998; Kuhn, 1962). By contrast, large language models now participate directly in the construction, testing and recombination of candidate claims within tightly coupled loops of interaction. The IHACC model and the notion of flow acceleration were proposed as a way of making this shift explicit, drawing on philosophy of technology and information (Heidegger, 1977; Latour, 2005; Floridi, 2014) to reinterpret empirical findings on AI-assisted work (Dell'Acqua et al, 2023; Noy & Zhang, 2023; Brynjolfsson, Li, & Raymond, 2023).

In what sense, then, does IHACC amount to a paradigm shift? In a constrained Kuhnian sense, it does not overthrow prior epistemologies so much as re-center the object of analysis. Where earlier frameworks located the epistemic subject primarily in individuals or communities, IHACC treats the governed human–AI system as the primary unit of "knowing". Flow-accelerated loops of framing, transformation and evaluation become the locus where data, information and knowledge are produced and stabilized. Latour's insistence that knowledge emerges from networks of heterogeneous actors and Floridi's portrayal of the infosphere as an environment of co-constituted informational entities both

anticipate this move (Latour, 2005; Floridi, 2014). IHACC extends their insights by specifying the recursive structure of human–AI interaction under conditions of acceleration and by foregrounding the role of human agents as orchestrators of prompts, constraints and norms.

The model also sharpens the ethical and political stakes of AI-mediated inquiry. Critical work on algorithmic systems has shown how infrastructures of data and computation can entrench power asymmetries, encode racial and gender bias and shift burdens of surveillance and correction onto workers and marginalized groups (Noble, 2018; Crawford, 2021). Empirical findings that AI assistance disproportionately benefits novices suggest real possibilities for widening epistemic participation, but these gains are fragile and contingent (Noy & Zhang, 2023; Brynjolfsson et al., 2023). IHACC makes clear that flow acceleration is compatible with both epistemic justice and epistemic injustice. If accelerated loops are governed by inclusive data practices, transparent model use and meaningful human oversight, they can support more equitable access to knowledge work. If they are deployed without such safeguards, they risk realizing Heidegger's warning that technology reduces beings including human knowers to 'standing reserve', mere resources for throughput (Heidegger, 1977). Coeckelbergh's call for a relational ethics of AI is therefore not an optional add-on but a condition of legitimacy: responsibility and vulnerability are negotiated within the very loops IHACC describes (Coeckelbergh, 2020).

These considerations point to several directions for further work. Conceptually, IHACC could be refined into a more explicit normative framework, specifying criteria for acceptable configurations of human framing, AI capability and feedback in different domains, and integrating insights from human-in-the-loop safety research (Holzinger, 2016). Empirically, a systematic program could examine IHACC-like dynamics across sectors, task types and institutional contexts, including settings outside the 'Global North' and within epistemic traditions that do not share Western assumptions about the individual knower. Such work would test the scope conditions of flow acceleration and hybrid agency, as well as their impact on workload, burnout and the quality of deliberation. Philosophically, IHACC invites dialogue with non-Western and processual epistemologies that already treat knowing as relational and dynamic, and with existential questions about meaning and identity in an environment where inquiry itself is subject to relentless acceleration.

If IHACC has a central claim, it is not that AI has made knowledge obsolete, but that in an AI-augmented world knowledge is less a finished product held by isolated subjects than an emergent property of governed interactions in time. To speak of "knowing" under conditions of flow acceleration is to speak about how humans design, inhabit and sustain these recursive loops. Knowledge, on this view, is no longer merely something to be stored and possessed; it is a relation to be maintained, criticized and renewed within the accelerated infrastructures we are now building.

### 5.2 Further Steps: Testable propositions and empirical cues

**P1. Distributional lift.**

Access to assistive models should raise performance more for weaker baseline performers than for stronger ones. In practice, we expect quality scores to bunch closer to the mean, with the lower tail improving most. A simple falsification would be no change in dispersion or a deterioration among lower-baseline participants.

**P2. Process, not presence.**

Outcomes should depend less on whether AI is used and more on how the human–model process is organized. Workflows that alternate in small units of work, with explicit review points, should outperform coarse one-shot hand-offs on tasks that require interpretation. If a binary "AI used" indicator explains results as well as, or better than, process measures, this proposition fails.

**P3. Throughput–verification trade-off.**

Tightly interleaved "cyborg" work should raise throughput while also increasing the time and attention required for checking. That added burden should decline when the process is traceable and disciplined, for example through prompt logs, version histories and checklists. If no additional verification effort is observed, or if traceability does not moderate the burden, the claim is not supported.

These propositions can be tested by using blind rating of outputs against pre-registered rubrics; recording simple process traces that capture who worked when and how often hand-offs occurred; and track time spent generating versus checking. Average effects should observed and reported combined with distributional changes and interactions with baseline ability.

### Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used ChatGPT 4.0, ChatGPT 5, Deepseek and Grok 4 in order to refine grammar, carry out formatting consistency checks, and to enhance readability. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

**References**


Ackoff, R. L. (1989). From data to wisdom. *Journal of Applied Systems Analysis, 16*, 3–9.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research.

Coeckelbergh, M. (2020). *AI ethics*. MIT Press.

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). https://doi.org/10.2139/ssrn.4573321

Descartes, R. (1641/1996). *Meditations on first philosophy* (J. Cottingham, Trans.). Cambridge University Press.

Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. Oxford University Press.

Heidegger, M. (1977). *The question concerning technology, and other essays* (W. Lovitt, Trans.). Harper & Row.

Holzinger, A. (2016). Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Brain Informatics, 3*(2), 119–131. https://doi.org/10.1007/s40708-016-0042-6

Kant, I. (1781/1998). *Critique of pure reason* (P. Guyer & A. W. Wood, Eds. & Trans.). Cambridge University Press.

Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.

Lakhani, K. R. (2024). Centaurs and cyborgs: Navigating the jagged edge of generative AI productivity [Video transcript]. CXOTalk. https://www.cxotalk.com/episode/centaurs-and-cyborgs-navigating-the-jagged-edge-of-generative-ai-productivity

Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford University Press.

McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio.

Nielsen, J. (2023). AI improves employee productivity by 66%. *Nielsen Norman Group*. https://www.nngroup.com/articles/ai-tools-productivity-gains/

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *arXiv*. https://arxiv.org/abs/2304.12244

Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv*. https://arxiv.org/abs/2302.06590

Pew Research Center. (2025, February 25). *Workers' views of AI use in the workplace*. https://www.pewresearch.org/social-trends/2025/02/25/workers-views-of-ai-use-in-the-workplace/

Robinson, B. (2024, July 23). 77% of employees report AI has increased workloads and hampered productivity, study finds. *Forbes*. https://www.forbes.com/sites/bryanrobinson/2024/07/23/employees-report-ai-increased-workload/

Rowley, J. (2007). The wisdom hierarchy: Representations of the DIKW hierarchy. *Journal of Information Science, 33*(2), 163–180.

Shin, J., Hedderich, M. A., Rey, B. J., Lucero, A., & Oulasvirta, A. (2024). Understanding human–AI workflows for generating personas. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)* (pp. 1–25). ACM. https://doi.org/10.1145/3643834.3660729

Wang, Z. Z., Shao, Y., Shaikh, O., Fried, D., Neubig, G., & Yang, D. (2025). How do AI agents do human work? Comparing AI and human workflows across diverse occupations. *arXiv*. https://arxiv.org/abs/2510.22780