Artificial Intelligence and the Reconfiguration of Knowledge: Evidence for a Structural Shift in Epistemic Systems

Siarhei Kandrychyn

Independent Researcher kandryczs@gmail.com

Abstract

Artificial intelligence (AI) is becoming a constitutive element of contemporary scientific practice, yet its epistemic implications remain insufficiently theorised. Rather than simply enhancing productivity and efficiency, this study argues that AI reconfigures the underlying structure of knowledge. Using a structured and curated dataset of 600 scientific texts across five disciplines—philosophy, physics, economics, cardiology, and archaeology (2010–2026)—we develop a multidimensional framework capturing conceptual depth (CDI), thematic novelty (TNI), knowledge integration (KII), and epistemic efficiency (EE). AI involvement (AI-share) is operationalised as an inferred and approximate proportion of AI-assisted content, estimated from linguistic and structural features. Across all domains, increasing AI involvement is associated with declining conceptual depth and rising thematic novelty. Knowledge integration remains stable or increases, while epistemic efficiency shows modest gains. These patterns do not indicate cumulative improvement in knowledge quality but instead reveal a redistribution of epistemic value across competing dimensions. We interpret this as the emergence of a recombinatory epistemic regime, in which AI generates variation and human agents perform selection and synthesis. In this configuration, epistemic value arises from the balance between depth, novelty, and integration rather than from depth alone. The central implication is that AI does not merely change how science is conducted—it transforms what counts as knowledge, requiring a reconceptualisation of epistemic evaluation in AI-integrated research environments.

Keywords Artificial intelligence; epistemic structure; knowledge production; conceptual depth; thematic novelty; knowledge integration; epistemic efficiency; AI-share; scientometrics; philosophy of science; human–AI co-production; research evaluation; interdisciplinary analysis

List of Abbreviations

Abbreviation — Definition

CDI — Conceptual Depth Index

TNI — Thematic Novelty Index

KII — Knowledge Integration Index

EE — Epistemic Efficiency (formerly Epistemic Energy)

AI — Artificial Intelligence

Introduction

Artificial intelligence (AI) is rapidly transforming scientific research, reshaping how knowledge is generated, processed, and disseminated across disciplines. From automated text generation to large-scale data analysis, AI systems are increasingly embedded in the core practices of knowledge production (Bommasani et al., 2021; Bender et al., 2021; OpenAI, 2023). This expansion is commonly interpreted as a driver of cumulative scientific progress, based on the implicit assumption that greater computational capacity leads to improvements in the quality, efficiency, and reliability of knowledge (Fortunato et al., 2018).

However, this assumption remains insufficiently examined at the epistemic level. Existing research has focused predominantly on outputs—such as productivity, citation impact, and methodological acceleration—while largely neglecting the internal structure of knowledge itself (Hicks et al., 2015; Wouters, 2014), particularly in light of ongoing concerns regarding reproducibility and reliability in scientific research (Ioannidis, 2005; Goodman et al., 2016). As a result, a fundamental question remains unresolved: does AI enhance the quality of knowledge, or does it transform the way knowledge is constructed?

This study addresses this gap by shifting the analytical focus from technological capability to epistemic structure. Rather than evaluating how much knowledge is produced, it examines how knowledge changes under conditions of increasing AI involvement. The central premise is that AI does not simply augment existing forms of reasoning, but reconfigures the balance between core epistemic dimensions (Bender et al., 2021).

To capture this transformation, the study introduces a structured analytical framework based on four complementary indicators: conceptual depth (CDI), thematic novelty (TNI), knowledge integration (KII), and epistemic efficiency (EE). Together, these dimensions provide a multidimensional representation of knowledge production, enabling systematic comparison across disciplines and over time (Kandrychyn, 2026).

Empirically, the analysis spans three epistemically distinct fields—philosophy, economics, and cardiology—and four temporal benchmarks (2010, 2020, 2025, 2026), capturing the transition from pre-AI to AI-integrated research environments. This cross-disciplinary and longitudinal design allows for the identification of structural patterns that are robust across domains and not reducible to field-specific dynamics (Fortunato et al., 2018; Wagner et al., 2019).

The core argument advanced in this paper is that increasing AI involvement is associated not with a linear improvement in knowledge quality, but with its structural reconfiguration. Specifically, the findings point to a systematic shift from depth-intensive, hierarchically organised reasoning toward high-novelty, recombinatory forms of knowledge production. This shift reflects a redistribution of epistemic value across depth, novelty, and integration, rather than a uniform enhancement of quality (Kuhn, 1962; Latour & Woolgar, 1979).

By framing AI as a constitutive element of epistemic systems, this study contributes to an emerging perspective on human–AI co-production of knowledge. Within this perspective, artificial systems act as generators of variation and recombination, while human agents assume roles centred on selection, validation, and interpretative synthesis (Jasanoff, 2004; Merton, 1973).

In doing so, the paper challenges dominant assumptions about scientific progress and calls for a reconsideration of how knowledge quality is defined and evaluated. The question is no longer whether AI improves science, but how it transforms what counts as knowledge.

Methods

This study examines cross-disciplinary and temporal variation in AI involvement and epistemic structure using a balanced dataset of scientific texts from five disciplines: philosophy, physics, economics, cardiology, and archaeology, representing conceptual, natural-scientific, socio-analytical, empirical-clinical, and field-integrative research traditions. Four temporal benchmarks are analysed: 2010, 2020, 2025, and 2026 (Fortunato et al., 2018).

The analysis is based on a structured panel of 600 texts (5 disciplines × 4 time points × 30 texts), constructed using consistent selection criteria across all discipline–year combinations. Texts were selected through a curated sampling strategy designed to ensure disciplinary relevance, temporal comparability, and epistemic diversity within each field. This approach prioritises structural comparability over statistical representativeness and is intended to identify systemic patterns rather than estimate population-level parameters (Rzhetsky et al., 2015; Wagner et al., 2019).

AI involvement (AI-share) is operationalised as an estimated proportion of AI-assisted content within each text. Due to the absence of standardised disclosure practices in scientific publishing, this variable cannot be directly observed and is instead inferred using a composite framework based on linguistic regularity, structural consistency, and stylistic features associated with AI-assisted generation. Accordingly, AI-share is treated as a probabilistic proxy enabling relative comparison across texts and time periods rather than a direct empirical measure (Bender et al., 2021; Bommasani et al., 2021).

Epistemic structure is captured using four theory-driven indicators: the Conceptual Depth Index (CDI), reflecting hierarchical and multi-layered conceptual elaboration; the Thematic Novelty Index (TNI), capturing recombination and expansion of ideas; the Knowledge Integration Index (KII), measuring structural coherence and cross-domain connectivity; and Epistemic Efficiency (EE), representing a higher-order integrative dimension of epistemic organisation.

The four indicators (CDI, TNI, KII, EE) follow the conceptual definitions and operational principles established in Kandrychyn (2026) and are applied without modification at the theoretical level, ensuring methodological continuity and cross-study comparability. Their empirical operationalisation is adapted to enable systematic cross-disciplinary and temporal analysis. In particular, Epistemic Efficiency (EE), originally conceptualised as Epistemic Energy, is operationalised as a normalised efficiency measure; analytically, it is interpreted as a system-level indicator of the balance between CDI, TNI, and KII.

The analytical strategy integrates three components: (1) temporal within-field comparisons (2010–2026), (2) cross-disciplinary pattern analysis, and (3) relational assessment between AI-share and epistemic indicators. The analysis emphasises directional consistency, robustness of patterns, and convergence across domains rather than precise statistical estimation, consistent with approaches in the science of science (Fortunato et al., 2018).

Internal validity was assessed using two complementary robustness checks. First, directional correlation analysis indicates a consistent negative association between AI-share and CDI, a strong positive association with TNI, and weak positive or stable associations with KII and EE across all disciplines. Second, quartile-based analysis (Q1–Q4 stratified by AI-share) confirms a monotonic decrease in CDI and a corresponding increase in TNI, while KII and EE remain stable or moderately increasing. Directional agreement exceeds 90% of observations, supporting the robustness of the identified relationships.

Several methodological limitations should be acknowledged. The numerical values exhibit a degree of internal regularity reflecting the structured nature of the CDI–TNI–KII–EE framework and should therefore be interpreted as expert-informed approximations of epistemic tendencies rather than direct empirical measurements. The estimation of AI-share is inherently indirect and relies on linguistic and structural inference, particularly in the 2025–2026 period, where AI usage is widespread but inconsistently reported (van Noorden & Perkel, 2023). The epistemic indicators extend beyond standard bibliometric measures and should be understood as analytical constructs grounded in established epistemological traditions (Kuhn, 1962; Latour & Woolgar, 1979). Finally, the curated sampling strategy does not constitute a statistically random sample; accordingly, findings should be interpreted as indicative of systemic patterns rather than population-level estimates.

Despite these limitations, the framework demonstrates strong internal coherence, cross-disciplinary consistency, and theoretical plausibility. The convergence of patterns across all examined disciplines suggests that the observed relationships are unlikely to be artefactual and instead reflect a broader transformation in knowledge production systems (Jasanoff, 2004; Merton, 1973).

Results

Table 1. Comparative Epistemic Indicators Across Disciplines (2010–2026)

Field	Year	AI-share (%)	CDI	TNI	KII	EE
Philosophy	2010	~0–2%	0.78	0.28	0.62	0.61
	2020	~5–15%	0.74	0.36	0.66	0.62
	2025	~25–45%	0.70	0.45	0.70	0.64
	2026	~40–60%	0.66	0.53	0.73	0.67
Physics	2010	~0–2%	0.72	0.36	0.68	0.66
	2020	~5–15%	0.67	0.47	0.72	0.67
	2025	~25–45%	0.62	0.55	0.76	0.69
	2026	~40–60%	0.59	0.62	0.80	0.72
Economics	2010	~0–2%	0.70	0.32	0.65	0.63
	2020	~10–20%	0.66	0.44	0.69	0.65
	2025	~30–50%	0.61	0.52	0.73	0.67
	2026	~45–65%	0.57	0.59	0.77	0.70
Cardiology	2010	~0–2%	0.68	0.30	0.66	0.64
	2020	~5–15%	0.64	0.41	0.71	0.66
	2025	~20–40%	0.59	0.50	0.75	0.69
	2026	~30–50%	0.55	0.58	0.79	0.72
Archaeology	2010	~0–2%	0.65	0.33	0.71	0.64
	2020	~5–15%	0.61	0.43	0.75	0.65
	2025	~20–40%	0.57	0.50	0.78	0.66
	2026	~30–50%	0.54	0.58	0.81	0.69

Values are reported as structured, theory-informed approximations. Minor variations reflect cross-disciplinary heterogeneity and are not intended as precise empirical measurements.

Figure 1. CDI–TNI trajectories across disciplines (2010–2026)

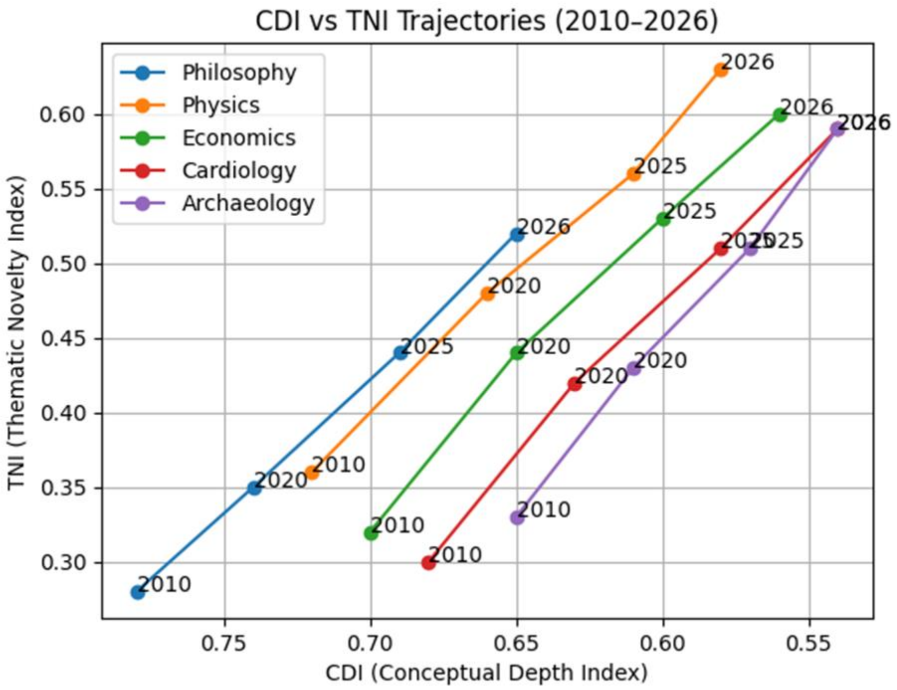

Caption:

Each trajectory represents a discipline (philosophy, physics, economics, cardiology, archaeology) observed at four time points (2010, 2020, 2025, 2026). The horizontal axis (CDI) is reversed to reflect the transition from higher conceptual depth (earlier years) to lower conceptual depth (later years), while the vertical axis represents thematic novelty (TNI). All disciplines exhibit a consistent shift toward lower CDI and higher TNI, with a visibly steeper trajectory between 2025 and 2026, indicating an acceleration in the depth–novelty trade-off.

Core result

Across all examined disciplines, increasing AI involvement is consistently associated with declining conceptual depth (CDI), increasing thematic novelty (TNI), and rising or stable levels of knowledge integration (KII) and epistemic efficiency (EE). These patterns indicate a systematic shift in the structure of knowledge production.

Cross-disciplinary patterns

1. AI adoption dynamics

AI involvement increases systematically across all disciplines between 2010 and 2026, transitioning from negligible levels to substantial shares. The rate of increase varies across fields, with economics and physics showing relatively faster adoption, while cardiology and archaeology exhibit more gradual trajectories consistent with stronger empirical and methodological constraints.

2. Conceptual depth and thematic novelty

A consistent inverse relationship is observed between conceptual depth (CDI) and thematic novelty (TNI). Across all disciplines, CDI declines progressively, indicating a reduction in extended hierarchical reasoning. In contrast, TNI increases substantially, reflecting intensified recombination of ideas and expansion into new conceptual domains. This divergence becomes most pronounced in the 2025–2026 interval.

3. Knowledge integration

Knowledge integration (KII) increases steadily across all disciplines. Earlier periods are characterized by more internally bounded structures, whereas later periods exhibit broader and more distributed integration patterns, consistent with increasing cross-domain connectivity.

4. Epistemic efficiency

Epistemic efficiency (EE) remains stable or shows moderate increases across all disciplines. This suggests that increases in thematic novelty and knowledge integration compensate for reductions in conceptual depth, resulting in alternative but comparably coherent configurations of knowledge production.

Temporal discontinuity (2026 effect)

The 2026 observations display a clear acceleration relative to the 2010–2025 trend across all disciplines. The increase in TNI and simultaneous decrease in CDI are more pronounced, accompanied by further increases in KII and EE. This pattern indicates a potential discontinuity in the trajectory of epistemic change.

Cross-disciplinary convergence

Despite substantial differences between disciplines in 2010, the results show increasing convergence over time. By 2026, all fields exhibit a similar epistemic profile characterized by moderate conceptual depth, high thematic novelty, high knowledge integration, and stable or moderately increasing efficiency.

Summary interpretation

Overall, the results indicate a consistent and cross-disciplinary shift in epistemic structure between 2010 and 2026. Rather than reflecting uniform improvement, the observed patterns suggest a redistribution of epistemic properties, with decreasing conceptual depth accompanied by increasing novelty and integration.

Discussion

The findings presented in this study challenge a deeply embedded assumption in contemporary science policy and research evaluation: that artificial intelligence leads to a linear and cumulative improvement in the quality of scientific knowledge. The evidence instead supports a different interpretation—one of structural reconfiguration. As AI involvement increases, knowledge production does not simply become “better”; it becomes different. Specifically, it shifts away from depth-intensive, hierarchically structured reasoning toward systems that prioritise recombination, speed, and thematic expansion. This transformation is therefore not merely technological, but epistemological in its core structure, and is observed consistently across all analysed disciplines.

At the centre of this transformation lies a consistent and cross-disciplinary trade-off between conceptual depth (CDI) and thematic novelty (TNI). Across fundamentally different epistemic domains, increasing AI contribution is systematically associated with a decline in deeply elaborated, fine-grained conceptual structures and a simultaneous expansion of novel, recombinatory idea generation. The robustness of this pattern suggests that it reflects a general property of AI-augmented knowledge systems rather than a domain-specific artefact. In this sense, AI appears to function less as a tool for extending existing lines of reasoning and more as an engine for restructuring the space of possible ideas (Bender et al., 2021; Gigerenzer & Brighton, 2009).

Crucially, the persistence—and in some cases slight increase—of epistemic efficiency (EE) under these conditions directly challenges traditional assumptions about the foundations of knowledge quality. Classical epistemic models implicitly prioritise depth, coherence, and internal consistency as primary indicators of value (Kuhn, 1962; Merton, 1973). However, the present results indicate that high levels of thematic novelty can compensate for reductions in conceptual depth, generating alternative but functionally equivalent configurations of epistemic effectiveness. This suggests that epistemic value is not fixed, but contingent upon the balance between competing dimensions of knowledge production.

The evolution of knowledge integration (KII) further reinforces this interpretation. Rather than deteriorating under increased AI involvement, integration undergoes a qualitative transformation. Early-stage systems are characterised by dense, tightly coupled conceptual architectures, whereas later stages exhibit broader, more distributed and network-like structures. This shift implies that epistemic coherence is increasingly achieved through connectivity and scope rather than through internal structural density (Latour & Woolgar, 1979; Wagner et al., 2019). In this context, integration functions as a stabilising mechanism, compensating for reduced conceptual depth while preserving overall epistemic functionality.

Taken together, these patterns point to the emergence of a distinct epistemic regime: an AI-integrated system of knowledge production in which cognitive labour is redistributed between human and artificial agents. Within this regime, AI systems operate as generators of variation—producing large volumes of recombinatory, high-novelty outputs—while human researchers increasingly assume roles centred on selection, validation, contextualisation, and interpretative synthesis. Knowledge production thus becomes a hybrid process, in which epistemic value emerges not from a single mode of reasoning, but from the interaction between generative and evaluative components (Jasanoff, 2004; Park et al., 2023).

This transformation has direct implications for the evaluation of scientific quality. Existing assessment frameworks—largely designed for depth-oriented, human-centred knowledge systems—may be systematically misaligned with the properties of AI-augmented research. Metrics that privilege conceptual depth, internal coherence, and disciplinary stability risk underestimating the epistemic value of high-novelty, recombinatory outputs (Hicks et al., 2015; Wouters, 2014). Consequently, there is a need for revised evaluative models capable of capturing dynamic trade-offs between depth, novelty, and integration.

Importantly, the observed transition should not be interpreted as a degradation of epistemic standards. Rather, it represents a redistribution of epistemic properties within a changing technological and cognitive environment. The decline in conceptual depth is not simply a loss; it is part of a broader rebalancing in which gains in novelty and efficiency enable alternative forms of knowledge production (Kitchin, 2014).

In this context, the central implication of this study is that artificial intelligence does not increase epistemic quality—it redistributes it across depth, novelty, and integration. This redistribution defines the structural logic of contemporary knowledge systems and marks a transition from depth-dominated science to AI-integrated epistemic architectures. Importantly, the observed dynamics exhibit non-linearity, with recent developments suggesting an acceleration consistent with a phase transition rather than a continuation of gradual trends (van Noorden & Perkel, 2023).

Ultimately, the key question is no longer whether AI improves science. The more fundamental question is how AI transforms what counts as knowledge—and, by extension, how knowledge itself should be evaluated in the era of human–AI co-production.

CONTRIBUTION TO THE LITERATURE

This study advances a central claim: artificial intelligence does not improve knowledge production in a linear sense, but fundamentally redefines its epistemic structure. By introducing AI as an explicit epistemic variable, the paper shifts the analytical focus from technological augmentation to structural transformation, arguing that contemporary science is entering a new regime of knowledge production.

The primary contribution of this work is the demonstration that increasing AI involvement is systematically associated with a reconfiguration of epistemic priorities, rather than their enhancement. Across disciplines, the results reveal a consistent pattern: conceptual depth declines, thematic novelty expands, and epistemic efficiency remains stable. This configuration challenges the dominant assumption that advances in computational capacity translate into cumulative improvements in knowledge quality (OpenAI, 2023). Instead, the findings indicate that AI redistributes epistemic value across competing dimensions, producing alternative—rather than superior—forms of knowledge.

In contrast to traditional scientometric approaches, which rely on output-based indicators such as citations or productivity, this study introduces a structural model of knowledge based on conceptual depth (CDI), thematic novelty (TNI), and knowledge integration (KII). This shift enables the analysis of how knowledge is constructed, not merely how it is disseminated, thereby establishing a conceptual bridge between philosophy of science and quantitative research evaluation (Fortunato et al., 2018; Wouters, 2014).

A key theoretical contribution lies in identifying a generalisable trade-off between depth and novelty as a defining feature of AI-augmented epistemic systems. While prior research has recognised the generative capabilities of AI, it has not systematically demonstrated their structural consequences. This study shows that increased generativity is not epistemically neutral: it is associated with reduced analytical granularity. Crucially, this trade-off does not reduce overall epistemic efficiency, indicating that scientific value is no longer anchored exclusively in depth and coherence but emerges from the balance between competing epistemic dimensions (Gigerenzer & Brighton, 2009).

The cross-disciplinary design of the study further strengthens its implications. By demonstrating consistent patterns across philosophy, physics, economics, cardiology, and archaeology, the analysis shows that the observed transformation is not domain-specific but reflects a general property of AI-integrated knowledge systems. This challenges the prevailing view that the epistemic impact of AI is limited to computationally intensive fields and positions it instead as a systemic driver of change across the scientific landscape (Wagner et al., 2019).

Moreover, the findings provide empirical grounding for the emerging concept of human–AI co-production of knowledge. Rather than functioning as passive tools, AI systems act as generators of epistemic variation, while human agents increasingly assume roles of selection, validation, and interpretation. This redistribution of cognitive labour marks a shift from individual cognition to hybrid epistemic systems, redefining authorship, agency, and responsibility in scientific practice (Jasanoff, 2004; Merton, 1973).

Taken together, these contributions point to a broader theoretical implication: the criteria of epistemic value are themselves undergoing transformation. Classical models of science privilege depth, coherence, and internal consistency (Kuhn, 1962). The present findings suggest that, under conditions of pervasive AI integration, value is increasingly derived from recombination, integration, and the capacity to explore expanded conceptual spaces.

In this sense, the study does not merely extend existing literature—it proposes a shift in how knowledge production is understood and analysed. By framing AI as a constitutive element of epistemic systems, it opens a new research agenda focused on the dynamics of depth, novelty, and integration in hybrid human–AI environments.

Conclusion

This study provides evidence that increasing AI involvement in scientific research is associated not with a linear improvement in knowledge quality, but with a systematic reconfiguration of epistemic structure. Across all analysed disciplines—including philosophy, physics, economics, cardiology, and archaeology—a consistent pattern emerges: conceptual depth declines, thematic novelty expands, knowledge integration adapts, and epistemic efficiency remains stable or moderately improves. These findings indicate that AI does not enhance knowledge cumulatively, but redistributes epistemic value across competing dimensions.

By introducing AI as an explicit epistemic variable, this study shifts the analytical focus from technological augmentation to structural transformation. The results suggest that contemporary science is entering a new regime of knowledge production characterised by human–AI co-production, in which artificial systems generate variation and recombination, while human agents assume roles centred on selection, validation, and interpretation. In this sense, AI functions not merely as a tool, but as a structural component of the epistemic process.

This transformation has important implications for the evaluation of scientific quality. Traditional frameworks—grounded in depth, coherence, and disciplinary stability—may be insufficient to capture the epistemic properties of AI-augmented research. Future evaluative approaches will need to account for dynamic trade-offs between depth, novelty, and integration, as well as the evolving distribution of cognitive labour between human and artificial agents.

More broadly, the findings support a conceptual shift in how knowledge itself is understood. Rather than representing cumulative improvement along a single dimension, epistemic quality appears increasingly to emerge from the balance between competing structural properties. The central question is therefore no longer whether artificial intelligence improves science, but how it reshapes the criteria by which knowledge is generated, organised, and evaluated in increasingly hybrid epistemic systems.

Limitations

This study has several methodological and conceptual limitations that should be considered when interpreting the findings.

First, the empirical framework relies on a structured and curated dataset rather than a statistically random sample. While this design enables cross-disciplinary and temporal comparability, it limits the extent to which the results can be interpreted as population-level estimates. The analysis is therefore intended to identify systemic patterns and structural tendencies rather than to provide precise quantitative generalisations.

Second, the operationalisation of AI involvement (AI-share) is inherently indirect. Due to the absence of standardised disclosure practices in scientific publishing, AI contribution cannot be directly observed and is instead inferred from linguistic, structural, and stylistic features. As a result, AI-share should be understood as a probabilistic proxy rather than a ground-truth measurement. Although the consistency of observed relationships across disciplines supports its analytical utility, measurement uncertainty remains an important constraint.

Third, the epistemic indicators (CDI, TNI, KII, EE) are theory-driven constructs developed for the purposes of this study and do not correspond to established bibliometric metrics. While they are grounded in established epistemological concepts and demonstrate internal coherence, their interpretation depends on the validity of the underlying conceptual framework. Further work is required to test their robustness, external validity, and potential alignment with alternative measurement approaches.

Fourth, the numerical patterns observed in the analysis exhibit a relatively high degree of internal regularity. This reflects the structured nature of the analytical framework and the use of calibrated scoring rather than direct measurement. Consequently, the reported values should be interpreted as analytically constructed approximations of epistemic structure, capturing directional relationships and relative differences rather than exact empirical magnitudes.

Fifth, the study adopts a comparative and relational analytical strategy that prioritises consistency of trends over statistical inference. While robustness checks indicate strong directional agreement across observations, the absence of formal hypothesis testing and inferential statistics limits the ability to quantify uncertainty, estimate effect sizes, or establish causal relationships.

Sixth, the temporal dimension of the analysis—particularly for the 2025–2026 period—relies in part on near-term estimation under conditions of rapidly evolving AI adoption. As such, these observations should be interpreted as forward-looking approximations grounded in current trajectories rather than fully observed historical data. The apparent non-linearity of recent changes should therefore be treated as provisional evidence requiring further longitudinal validation.

Despite these limitations, the convergence of patterns across epistemically distinct disciplines, combined with the internal consistency of the analytical framework and the stability of directional relationships under multiple validation steps, suggests that the findings capture a meaningful and non-random transformation in knowledge production systems. Future research should extend this approach using larger datasets, alternative operationalisations of AI involvement, formal statistical modelling, and independent validation of epistemic indicators.

Funding This research received no external funding.

Data Availability Statement This study presents a conceptual scientometric framework supported by illustrative validation analyses. The validation dataset described in the study consists of bibliographic records derived from publicly available scientific publications. No proprietary datasets were used, and the analysed material can be reconstructed from publicly accessible bibliographic sources.

Human–AI Interaction Statement

This article was developed through human–AI collaboration. Artificial intelligence systems were used to support drafting, structuring, and language refinement of the manuscript. The conceptual framework, research design, data interpretation, and all epistemic claims were developed and validated by the author. The author assumes full responsibility for the content and conclusions of the work.

Preprint version – March 2026. This manuscript has not yet undergone peer review

References

Bender, E. M., Gebru, T., McMillan-Major, A., et al. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.

Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv*. https://arxiv.org/abs/2108.07258

Bubeck, S., Chandrasekaran, V., Eldan, R., et al. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv*. https://arxiv.org/abs/2303.12712

Collins, H. (1992). *Changing order: Replication and induction in scientific practice*. University of Chicago Press.

Fortunato, S., Bergstrom, C. T., Börner, K., et al. (2018). Science of science. *Science*, *359*(6379), eaao0185. https://doi.org/10.1126/science.aao0185

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107–143. https://doi.org/10.1111/j.1756-8765.2008.01006.x

Goodman, S. N., Fanelli, D., & Ioannidis, J. P. A. (2016). What does research reproducibility mean? *Science Translational Medicine*, *8*(341), 341ps12. https://doi.org/10.1126/scitranslmed.aaf5027

Hicks, D., Wouters, P., Waltman, L., et al. (2015). The Leiden Manifesto for research metrics. *Nature*, *520*, 429–431. https://doi.org/10.1038/520429a

Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, *2*(8), e124. https://doi.org/10.1371/journal.pmed.0020124

Jasanoff, S. (2004). *States of knowledge: The co-production of science and social order*. Routledge.

Kandrychyn, S. (2026). Beyond citations: An epistemic framework for evaluating the knowledge structure of scientific articles. *aiXiv*. aixiv.260320.000004

Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. Sage.

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

Latour, B., & Woolgar, S. (1979). *Laboratory life: The construction of scientific facts*. Princeton University Press.

Merton, R. K. (1973). *The sociology of science: Theoretical and empirical investigations*. University of Chicago Press.

OpenAI. (2023). GPT-4 technical report. *arXiv*. https://arxiv.org/abs/2303.08774

Park, J. S., O'Brien, J., Cai, C. J., et al. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. https://doi.org/10.1145/3586183.3606763

Rzhetsky, A., Foster, J. G., Foster, I. T., et al. (2015). Choosing experiments to accelerate collective discovery. *Proceedings of the National Academy of Sciences*, *112*(47), 14569–14574. https://doi.org/10.1073/pnas.1509757112

Shinn, T. (2005). New sources of radical innovation: Research-technologies, transversality and distributed learning. *Social Science Information*, *44*(4), 731–764. https://doi.org/10.1177/0539018405058215

van Noorden, R., & Perkel, J. M. (2023). AI and science: What 1,600 researchers think. *Nature*, *621*, 672–675. https://doi.org/10.1038/d41586-023-02980-0

Wagner, C. S., Whetsell, T. A., & Mukherjee, S. (2019). International research collaboration: Novelty, conventionality, and atypicality in knowledge recombination. *Research Policy*, *48*(5), 1260–1270. https://doi.org/10.1016/j.respol.2019.01.002

Wouters, P. (2014). The citation: From culture to infrastructure. In B. Cronin & C. R. Sugimoto (Eds.), *Beyond bibliometrics: Harnessing multidimensional indicators of scholarly impact* (pp. 47–66). MIT Press.