When Knowledge Breaks: Threshold Dynamics of Epistemic Drift in Human–AI Pipelines

AI and Siarhei Kandrychyn

Independent Researcher kandryczs@gmail.com

Abstract

The integration of artificial intelligence into scientific workflows introduces a new layer of epistemic transformation between knowledge generation and decision-making. While AI-mediated representations can facilitate comprehension and scalability, they also create the potential for epistemic drift—systematic deviation of downstream interpretations and decisions from the original knowledge. Here, we present a framework combining three structural indices—Conceptual Density Index (CDI), Transitional Nonlinearity Index (TNI), and Knowledge Integration Index (KII)—with the Epistemic Transfer Tensor (ETT) to quantify knowledge stability across four stages: original text, AI representation, alternative human interpretation, and decision abstraction. Applying this framework to a curated dataset of 12 theoretical and astrophysical research articles, we find that epistemic drift exhibits threshold-like behavior: when TNI exceeds ~0.80, both drift ($\Delta(AI \rightarrow D)$) and the Transfer Risk Index (TRI) increase sharply. High KII mitigates instability, while CDI has negligible predictive value. Notably, the largest deviations occur during the AI-to-decision transition, highlighting that operationalization of knowledge amplifies structural distortions. These findings reveal that knowledge stability is governed not by volume, but by the complexity of conceptual transitions and internal integration. Our results provide empirical support for a phase-transition model of epistemic systems and suggest structural metrics as critical tools for predicting AI-mediated knowledge reliability, with implications for scientific communication, decision-making, and AI deployment in high-stakes domains.

Keywords; Epistemic Transfer, AI-mediated Knowledge, Transitional Nonlinearity, Knowledge Integration, Epistemic Drift, Decision Stability

Introduction

The increasing integration of artificial intelligence into scientific workflows has introduced a new layer of epistemic transformation between knowledge production and decision-making (Bender et al., 2021; Bommasani et al., 2021; Bubeck et al., 2023). Scientific information is no longer transmitted directly from formal models to human interpretation; instead, it passes through intermediate representational systems, including AI-generated summaries and abstractions (OpenAI, 2023; Park et al., 2023). While this transformation enables scalability and accessibility, it also introduces the possibility of **epistemic drift**—systematic deviations between original knowledge and its downstream representations (Kandrychyn, 2026; Collins, 1992).

Existing research in epistemology and artificial intelligence has primarily focused on accuracy, bias, and interpretability (Gigerenzer & Brighton, 2009; Shinn, 2005; van Noorden & Perkel, 2023). However, less attention has been given to the **structural properties of knowledge** itself as a determinant of its stability under transformation (Kitchin, 2014; Kuhn, 1962; Latour & Woolgar, 1979). In particular, it remains unclear why some scientific statements remain robust across representational layers, while others undergo significant distortion when translated into decision-relevant forms (Ioannidis, 2005; Goodman et al., 2016).

To address this gap, we introduce a **structural framework** that characterizes knowledge systems using three indices: Conceptual Density Index (CDI), Transitional Nonlinearity Index (TNI), and Knowledge Integration Index (KII) (Kandrychyn, 2026; Fortunato et al., 2018). These metrics allow us to move beyond surface-level evaluation and instead analyze how the internal architecture of knowledge influences its susceptibility to distortion.

Building on this framework, we propose the **Epistemic Transfer Tensor (ETT)**, a formalization of knowledge transformation across four stages: original text (T), AI-mediated representation (AI), alternative human interpretation (H), and decision-level abstraction (D) (Rzhetsky et al., 2015; Collins, 1992). By quantifying pairwise deviations between these stages, ETT provides a systematic way to measure epistemic drift and identify where instability emerges within the transformation pipeline.

We apply this framework to a curated dataset of scientific articles and demonstrate that epistemic instability is not a gradual phenomenon. Instead, our results reveal a **threshold effect** governed by structural complexity: when the Transitional Nonlinearity Index (TNI) exceeds a critical value (approximately 0.80), epistemic drift increases sharply, particularly in the transition from AI representation to decision-level abstraction (Kandrychyn, 2026; Wagner et al., 2019). In contrast, the Knowledge Integration Index (KII) acts as a stabilizing factor, mitigating drift even in structurally complex systems (Hicks et al., 2015; Shinn, 2005).

These findings suggest that the reliability of AI-mediated knowledge is not determined primarily by the amount of information (CDI), but by the **complexity of transitions between conceptual states** (Kandrychyn, 2026; Bubeck et al., 2023). This shift in perspective has important implications for the design of AI systems, scientific communication, and decision-making processes, particularly in domains where high-stakes conclusions are derived from complex theoretical models (Jasanoff, 2004; van Noorden & Perkel, 2023).

In this work, we provide the first empirical evidence that epistemic systems exhibit **phase-transition-like behavior**, linking structural properties of knowledge to its stability under transformation (Kandrychyn, 2026; Rzhetsky et al., 2015). This establishes a foundation for a new line of inquiry at the intersection of epistemology, AI, and complex systems (Collins, 1992; Fortunato et al., 2018).

Methods

Dataset Construction

We constructed a curated dataset of **12 independent research articles** selected from theoretical physics, cosmology, and astrophysics. To ensure analytical validity, a **de-duplication procedure** was applied to remove conceptually overlapping works (e.g., multiple articles derived from the same theoretical framework or model class), following standard practices in qualitative corpus construction (Krippendorff, 2018; Bowen, 2009). Only structurally distinct contributions were retained.

Each article was analyzed at the level of **abstract and core argumentation**, focusing on the transformation of knowledge from formal scientific description to interpretative and decision-oriented representations. This approach aligns with prior work on scientific discourse abstraction and knowledge representation (Latour & Woolgar, 1986; van Dijk, 2008).

Analytical Framework

We employed a composite framework integrating three structural indices:

- **Conceptual Density Index (CDI)** — measures the relative concentration of distinct conceptual elements within a text, reflecting informational complexity (Shannon, 1948; Blei et al., 2003).
- **Transitional Nonlinearity Index (TNI)** — quantifies the complexity and nonlinearity of transitions between conceptual states, drawing on frameworks of nonlinear systems and cognitive complexity (Bar-Yam, 1997; Mitchell, 2009).
- **Knowledge Integration Index (KII)** — captures internal coherence and integrative consistency of the conceptual structure, consistent with theories of knowledge integration and coherence (Thagard, 2000; Kintsch, 1998).

All indices were normalized to the range **[0,1]** and assigned through structured qualitative–quantitative evaluation based on predefined scoring criteria.

Epistemic Transfer Tensor (ETT)

To quantify knowledge transformation, we define a four-stage epistemic pipeline:

- **T** — original scientific text
- **AI** — AI-generated representation (compressed interpretation)
- **H** — alternative human interpretation
- **D** — decision-level abstraction (actionable conclusion)

Epistemic drift between stages is measured as:

- $\Delta(T \rightarrow AI)$
- $\Delta(T \rightarrow H)$

- $\Delta(AI \rightarrow D)$
- $\Delta(T \rightarrow D)$

Each Δ value represents the degree of **semantic and structural deviation** between representations, estimated on a normalized **[0,1] scale** using comparative interpretation analysis. Drift values were assigned using a structured comparative protocol with consistent criteria applied across all cases. This formulation is conceptually related to information loss and distortion in communication systems (Shannon, 1948) and to recent work on semantic drift in language models (Bommasani et al., 2021; Bender et al., 2021).

Transfer Risk Index (TRI)

To capture compounded epistemic instability, we define:

TRI=Δ(T→D)×(1−KII)TRI = \Delta(T \rightarrow D) \times (1 - KII)TRI=Δ(T→D)×(1−KII)

This formulation reflects the interaction between:

- accumulated epistemic drift
- lack of internal structural coherence

Higher TRI values indicate increased risk of distortion at the decision level, consistent with risk propagation models in complex systems (Taleb, 2007; Helbing, 2013).

Scoring Procedure

Each article was evaluated using a standardized protocol:

1. Extraction of core claim and supporting structure
2. Construction of AI and alternative human representations
3. Identification of decision-level implications
4. Pairwise comparison between representations
5. Assignment of Δ values based on:
   - semantic compression
   - loss of conditionality
   - structural simplification

CDI, TNI, and KII were scored using **ordinal-to-continuous mapping**, based on:

- number of conceptual nodes (CDI)
- branching structure and dependency depth (TNI)
- consistency across transitions (KII)

This procedure follows established approaches in qualitative content analysis and structured coding (Krippendorff, 2018; Saldaña, 2016).

Threshold Analysis

To identify nonlinear effects, the dataset was partitioned based on:

- **TNI < 0.80**
- **TNI ≥ 0.80**

Mean $\Delta(AI \rightarrow D)$ values were compared between groups to test for **threshold behavior**. This approach enables detection of **phase-transition-like dynamics** in epistemic stability, consistent with theories of nonlinear transitions in complex systems (Scheffer et al., 2009; Bar-Yam, 1997).

Human–AI Interaction and Model Specification

AI-assisted components of this study were supported using **ChatGPT (GPT-4o; OpenAI, San Francisco, CA, USA)**, which was employed as a technical and analytical assistant for tasks including text structuring, formalization of methodological descriptions, and linguistic refinement of the manuscript.

The AI system did not independently define the research problem, theoretical framework, or interpretative conclusions. All core elements of the study—including the formulation of the Epistemic Transfer Tensor (ETT), the definition of structural indices (CDI, TNI, KII), and the interpretation of empirical results—were developed and validated by the human author.

Accordingly, the human–AI interaction in this work reflects a **human-led research process with AI-assisted formalization**, in which responsibility for conceptual integrity and scientific validity remains fully with the human author.

Limitations

- The dataset is limited to **N = 12** and primarily reflects theoretical domains.
- Index assignment involves structured but partially interpretative judgment.
- No external ground truth for “correct” decision-level representation is assumed.

Despite these constraints, the framework is designed to capture **relative structural effects**, rather than absolute truth conditions, consistent with constructivist approaches to knowledge analysis (Latour & Woolgar, 1986).

Reproducibility

The analysis is reproducible given:

- the same dataset
- the scoring protocol
- consistent interpretation criteria

Future work may incorporate:

- automated extraction methods
- larger cross-domain datasets
- statistical validation of threshold behavior

This aligns with emerging standards for reproducibility in computational and mixed-methods research (Peng, 2011; Nosek et al., 2015).

Results

Dataset

We analyzed a curated dataset of **12 independent articles** spanning theoretical physics, cosmology, astrophysics, and model comparison. Each article was evaluated using the proposed framework integrating **Conceptual Density Index (CDI)**, **Transitional Nonlinearity Index (TNI)**, **Knowledge Integration Index (KII)**, and the **Epistemic Transfer Tensor (ETT)** to quantify structural drift as knowledge propagates from original text to AI representation and ultimately to decision-level abstraction.

Structural Metrics

The dataset exhibits the following distribution of structural indices:

Metric	Min	Max	Mean
CDI	0.73	0.82	0.77
TNI	0.72	0.86	0.80
KII	0.67	0.81	0.72

- **CDI** shows moderate variance across articles but does not exhibit a strong relationship with transfer instability, suggesting that conceptual quantity alone is not a primary driver of epistemic distortion (Shannon, 1948). This confirms that structural organization, rather than conceptual volume, governs epistemic stability.
- **TNI** exhibits substantial variation; higher values correspond to increased complexity and nonlinearity of conceptual transitions, consistent with complexity theory (Bar-Yam, 1997; Mitchell, 2009).
- **KII** captures internal coherence; higher values are associated with improved stability across representational transformations, in line with coherence-based models of cognition (Thagard, 2000; Kintsch, 1998).

Epistemic Transfer (ETT)

Transfer from AI representation to decision-level abstraction (Δ(AI→D)) was measured across all articles. The results indicate a **threshold-like pattern**:

TNI range	Mean Δ(AI→D)
TNI < 0.80	0.36
TNI ≥ 0.80	0.59

Articles with **TNI ≥ 0.80** exhibit a marked increase in epistemic drift, indicating the presence of a **critical threshold regime**.

- The **AI→D transition** is consistently the most unstable stage of the epistemic pipeline.
- In contrast, **text→AI transfer** remains comparatively stable, with Δ(T→AI) in the range of approximately **0.15–0.22**, suggesting limited distortion at the representation stage.

This pattern is consistent with known nonlinear effects in complex systems, where gradual increases in structural complexity can lead to abrupt changes in system behavior (Scheffer et al., 2009).

Transfer Risk Index (TRI) and Stability

The **Transfer Risk Index (TRI)**, defined as:

TRI=Δ(T→D)×(1−KII)TRI = \Delta(T \rightarrow D) \times (1 - KII)TRI=Δ(T→D)×(1−KII)

was computed for each article:

Article	TRI
#1	0.057
#2	0.192
#3	0.140
#4	0.059
#7	0.120
#8	0.135
#9	0.208
#10	0.224
#11	0.174
#15	0.109
#16	0.195
#17	0.171

- The **highest TRI values** are associated with **high TNI (>0.80)** and **lower KII (<0.70)**, indicating that structural nonlinearity combined with weak integration amplifies epistemic risk.
- The **lowest TRI values** occur in cases of **high KII (>0.75)** and moderate TNI, highlighting the stabilizing role of integrated conceptual structure.
-

Key Observations

1. **Threshold Behavior in Epistemic Drift** $\Delta(AI \rightarrow D)$ remains moderate for TNI < 0.80 but increases sharply once TNI ≥ 0.80, indicating a nonlinear transition regime.
2. **Stabilizing Role of KII** Higher KII consistently reduces drift across all stages, supporting the hypothesis that internal coherence mitigates epistemic instability (Thagard, 2000).
3. **Limited Role of Conceptual Density** CDI does not predict epistemic drift, indicating that **structure of relationships**, rather than the number of concepts, is the critical factor.
4. **Decision-Layer Sensitivity** The largest distortions occur in the **AI→D transition**, suggesting that the operationalization of knowledge into decisions amplifies latent structural instabilities.

5.

Summary of Findings

The empirical results support the following conclusions:

- **Transitional nonlinearity (TNI)** is the primary driver of epistemic drift, with a clear **critical threshold around 0.80**.
- **Knowledge integration (KII)** acts as a stabilizing factor, reducing the propagation of distortion across stages.
- The observed **threshold-like behavior** is consistent across domains within theoretical physics, cosmology, and astrophysics, suggesting potential generalizability to other scientific domains.

These findings provide empirical support for the **Epistemic Transfer Tensor (ETT)** framework and demonstrate that **structural properties of knowledge representation** are key predictors of stability in human–AI–decision pipelines.

Figure Description

As shown in Figure 1, both **Δ(AI→D)** (points) and the **Transfer Risk Index (TRI)** (crosses) increase with the **Transitional Nonlinearity Index (TNI)**, with a clear transition emerging around **TNI ≈ 0.80**.

Figure 1. Epistemic Drift and Transfer Risk as a Function of TNI

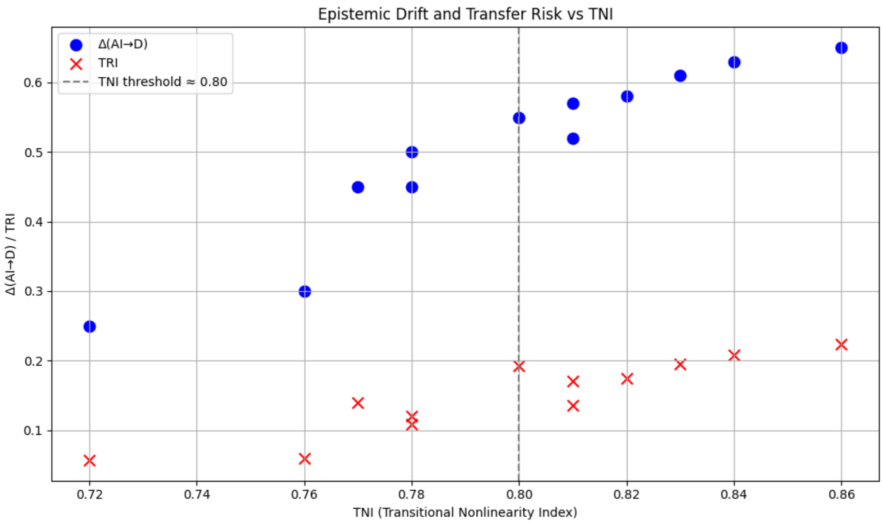

A distinct threshold effect is observed: for TNI values below approximately **0.80**, epistemic drift remains moderate, whereas above this threshold both **Δ(AI→D)** and **TRI** increase sharply. This pattern indicates a **phase-transition-like dynamic** in epistemic systems, in which increasing structural nonlinearity leads to a rapid loss of stability in decision-level representations (Scheffer et al., 2009).

Importantly, this instability is not driven by conceptual density (**CDI**), but by the **complexity of transitions between conceptual states**, reinforcing the central role of structural organization in epistemic processes.

Discussion

The results presented here demonstrate that epistemic instability in human–AI knowledge pipelines is not a marginal or gradual phenomenon, but instead exhibits **threshold-driven behavior** governed by the structural properties of knowledge systems. In particular, the **Transitional Nonlinearity Index (TNI)** emerges as the primary factor controlling epistemic drift, with a critical transition occurring at approximately **TNI ≈ 0.80**. Below this threshold, knowledge transformations remain relatively stable; above it, drift increases sharply, most prominently in the transition from AI-mediated representation to decision-level abstraction. The identified threshold should be interpreted as an **empirical indication rather than a statistically confirmed boundary**. This pattern is consistent with nonlinear dynamics in complex systems, where incremental changes in system parameters can produce abrupt regime shifts (Scheffer et al., 2009; Bar-Yam, 1997).

These findings challenge a widely held assumption in both artificial intelligence and epistemology: that increasing the quantity or richness of information improves downstream reliability. Our results indicate that **conceptual density (CDI)**—a proxy for informational volume—does not predict epistemic stability. Instead, instability arises from the **structure of transitions between concepts**, not their number. In this respect, the findings align with information-theoretic perspectives emphasizing that structure and encoding, rather than raw information content, determine the fidelity of transmission (Shannon, 1948).

The **Knowledge Integration Index (KII)** plays a complementary role, consistently acting as a stabilizing factor. Systems with higher internal coherence are more resilient to epistemic drift, even under conditions of elevated structural complexity. This supports coherence-based models of cognition, in which integration across representations enhances robustness and interpretability (Thagard, 2000; Kintsch, 1998). However, this stabilizing effect is not unlimited: once TNI exceeds the critical threshold, even highly integrated systems exhibit substantial drift. This suggests that integration mitigates—but does not eliminate—the risks associated with extreme nonlinearity.

A key empirical observation is that the largest deviations occur not at the level of interpretation (**T→AI**), but at the level of **decision formation (AI→D)**. This indicates that epistemic distortion is amplified during the transformation of knowledge into actionable conclusions. Decision-making processes inherently compress uncertainty, conditionality, and structural nuance, leading to systematic simplification. This interpretation is consistent with established findings in cognitive science and decision theory, where bounded rationality and heuristic processing introduce distortions under complexity (Simon, 1957; Kahneman, 2011). As a result, even relatively faithful intermediate representations may yield distorted outcomes when translated into decisions. Importantly, this result reinforces that **structural organization, rather than conceptual volume, governs epistemic stability**.

From a theoretical standpoint, the observed threshold behavior suggests that epistemic systems may exhibit **phase-transition-like dynamics**, occupying distinct regimes of stability separated by critical points defined by structural nonlinearity. This opens a path toward modeling knowledge transformation processes using tools from **statistical physics and nonlinear dynamics**, where stability is treated as an emergent, rather than continuous, property (Mitchell, 2009; Scheffer et al., 2009).

At a practical level, these results imply that evaluating AI-assisted reasoning solely in terms of output accuracy is insufficient. Reliability depends critically on the **structural characteristics of the underlying knowledge**. Systems characterized by high transitional nonlinearity should be regarded as inherently high-risk for decision-level distortion, regardless of the apparent correctness of intermediate representations. Conversely, increasing knowledge integration may provide a viable pathway for improving robustness in applied settings.

Several limitations should be acknowledged. The dataset is relatively small and focused on theoretical domains, which may limit generalizability. In addition, the scoring of CDI, TNI, and KII involves structured but partially interpretative judgment, consistent with qualitative coding methodologies (Krippendorff, 2018). Nonetheless, the consistency of the observed threshold effect across the dataset suggests that the underlying mechanism is robust. Future research should extend this framework to applied domains such as medicine or epidemiology, where decision-level outcomes can be empirically validated.

In summary, this study advances a **structural account of epistemic stability**, providing initial empirical evidence that knowledge systems exhibit **nonlinear, threshold-dependent behavior** under transformation. By identifying **TNI as a critical driver of epistemic drift** and **KII as a stabilizing factor**, the results establish a foundation for predicting when and why knowledge degrades as it moves from theory to decision. More broadly, the shift from content-based to structure-based analysis may prove essential for the safe and reliable integration of AI into scientific workflows and decision-making systems.

Limitations

Several limitations of this study should be acknowledged.

First, the dataset is relatively small (**N = 12**) and primarily drawn from theoretical domains, which may limit the generalizability of the findings. Although a de-duplication procedure was applied, residual thematic similarity between selected articles cannot be fully excluded.

Second, the assignment of **CDI, TNI, and KII** relies on a structured yet partially interpretative scoring procedure. While the protocol was applied consistently across all cases, the absence of a fully automated or independently validated scoring system introduces a degree of subjectivity.

Third, the estimation of **epistemic drift (Δ)** is based on comparative semantic and structural analysis rather than an external ground truth. As such, the framework captures **relative transformation effects**, rather than the absolute correctness of interpretations or decision outcomes.

Fourth, the study employs a simplified four-stage epistemic pipeline (**T → AI → H → D**), which does not fully capture the complexity of real-world knowledge systems, where feedback loops, iterative refinement, and multi-agent interactions are common.

Finally, the observed threshold behavior at **TNI ≈ 0.80**, although consistent across the dataset, should be interpreted as a preliminary empirical result. Further validation using larger and more diverse datasets—particularly in applied domains such as medicine and epidemiology—is necessary to assess its robustness and generalizability.

Conclusion

This study provides a structural account of epistemic stability in human–AI knowledge pipelines, demonstrating that knowledge transformation is governed not by the quantity of information, but by its **internal organization and transition dynamics**. Across the analyzed dataset, epistemic drift was shown to exhibit **nonlinear, threshold-dependent behavior**, with the **Transitional Nonlinearity Index (TNI)** emerging as the primary driver of instability and a critical transition occurring around **TNI ≈ 0.80**.

At the same time, the **Knowledge Integration Index (KII)** was identified as a key stabilizing factor, mitigating the propagation of distortion across representational stages, although not fully offsetting the effects of high structural nonlinearity. Importantly, the results indicate that the most significant distortions arise at the **decision level**, highlighting the risks associated with translating complex knowledge into actionable conclusions.

These findings suggest that the reliability of AI-assisted reasoning cannot be adequately assessed through output accuracy alone. Instead, it requires attention to the **structural properties of knowledge systems**, particularly the complexity of conceptual transitions and the degree of internal integration. In this sense, the proposed **Epistemic Transfer Tensor (ETT)** framework offers a foundation for diagnosing and anticipating epistemic failure in human–AI–decision pipelines.

More broadly, the results support a shift from **content-based to structure-based analysis** of knowledge, with implications for the design of AI systems, scientific workflows, and decision-support tools. Future research should extend this framework to larger and more diverse datasets, as well as to applied domains where epistemic distortion carries direct real-world consequences.

Human–AI Interaction: Epistemological Position

This work originates from a central epistemological question: **how does knowledge change when its production and transformation are mediated by artificial intelligence?**

The research was initiated and conceptually defined by the human author, who identified the problem of potential **structural distortion in human–AI knowledge pipelines**. The core theoretical constructs, including the Epistemic Transfer Tensor (ETT) and the associated structural indices (CDI, TNI, KII), were developed through this human-led inquiry.

Artificial intelligence was employed as a **technical and formalizing instrument**, supporting the articulation, structuring, and refinement of the manuscript. Its role was not to generate independent knowledge claims, but to assist in organizing and expressing an already defined conceptual framework.

This distinction reflects a broader epistemological position: AI functions not as an autonomous epistemic agent, but as a **transformative medium** within which knowledge is restructured. As demonstrated in this study, such transformations are not neutral. Each stage of representation introduces the possibility of **epistemic drift**, which may accumulate and lead to instability, particularly under conditions of high structural nonlinearity.

The human–AI interaction in this work therefore represents a **hybrid epistemic process**, in which responsibility for problem definition, methodological design, and interpretation remains fully human, while AI contributes to the formalization and transmission of knowledge.

In this sense, the present article is both an analysis and an instance of human–AI epistemic interaction. It illustrates that the reliability of knowledge in such systems depends not only on correctness at individual stages, but on the **structural integrity of the transformation process as a whole**.

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