Version Responsibility in AI-Era Knowledge Production and the Reconstruction of Authorship

Jiacheng Wu*[1] and Zuoxiu Zheng[2]

[1,2]College of Computer Science and Technology,Zhejiang University,Hangzhou,China

Abstract: The widespread application of generative artificial intelligence is profoundly reconstructing the paradigm of knowledge production, shifting it from the traditional model of "a single human author and a single original text" toward a model of human–machine collaboration. In this process, traditional authorship is substantively deconstructed because its internal unity of intentionality and responsibility is interrupted, which in turn produces black-boxed knowledge production and accountability dilemmas. To address this problem, this paper introduces the concept of version responsibility. The concept suspends the metaphysical debate over AI originality and redirects governance toward the question of who bears ultimate social and ethical responsibility for a specific public version of knowledge. Drawing on the RID cognitive dynamics model in *Knowing and Saying: An Ontological Investigation of Human Cognition* , the paper explains version responsibility through three dimensions: factual verification, transparent disclosure, and consequence accountability. It argues that the philosophical essence of version responsibility is to defend the priority of human beings in the dimension of problem pressure, or the D-dimension. On this basis, the paper proposes an institutional principle of functional outsourcing and responsibility centralization, and develops governance frameworks for high-risk fields such as academic publishing, medical decision support, and judicial sentencing. The goal is to rebuild a truthful and reliable public knowledge ecology while still releasing the cognitive dividends of AI.

Keywords: Generative AI; Knowledge production; Version responsibility; Authorship; Human–machine collaboration; Technology governance

1 Introduction: The Impact and Deconstruction of Traditional Knowledge Production by Generative AI

With the rapid development of large language models and generative artificial intelligence, human knowledge production is undergoing an unprecedented paradigmatic transition. From early word-processing software to systems such as ChatGPT, Claude, and other generative models, AI has moved through several technical stages, including symbolic approaches, connectionist approaches, deep learning, and transformer-based pretraining. The breakthrough of pretrained language models has made it possible for machines not merely to process the surface grammar of natural language but also to simulate deeper patterns of semantic generation. In this historical technological transition, AI has changed from a passive auxiliary tool, such as a spelling checker or formatting assistant, into an active generative partner that participates in the construction of concepts, arguments, and text.

* Corresponding author: 1320389424@qq.com

This role change has penetrated not only routine administrative writing but also academic research, journalism, literary creation, and other advanced cognitive activities. Traditional knowledge production relied heavily on human intellectual labor. The author was not only the creator of the text but also the sole bearer of responsibility for its content. Generative AI disrupts this linear causal chain and pushes knowledge production into a new stage of human–AI collaboration. In this new paradigm, AI is no longer merely a tool outside the cognitive process. It can participate in conceptual generation, logical inference, and the drafting of written expression. This deep technological intervention directly challenges the recognition of authorship and the allocation of intellectual-property rights.

Zhao argues that artificial intelligence is redrawing the epistemological boundary of human beings: machines do not merely surpass human beings in computational capacity, but also challenge traditional epistemology through new modes of knowledge generation [24]. Liu's account of embodied cognition similarly emphasizes that AI cognition detached from bodily and social practice is a form of hollow cognition

that lacks intentionality [14]. This fundamental difference in cognitive mechanism creates an unprecedented theoretical dilemma for human–machine collaborative knowledge production.

Current academic and legal debates continue to dispute the copyright status of AI-generated content. Some scholars argue for limited authorial status by analogy to constructive authorship, whereas mainstream legal practice in many jurisdictions remains committed to human-centered authorship and refuses to recognize AI as an author. Yet the crisis of authorship did not begin with AI. Foucault's question "What is an author?" treated the author as a discursive function that limits and organizes meaning [12]. Barthes's declaration of the death of the author went further, locating textual meaning in reading rather than in authorial intention [1]. In the AI era, however, the death of the author is no longer only a metaphor. It has become a technical reality.

Merely denying AI authorial status at the legal level cannot solve the ethical and responsibility dilemmas of knowledge production. When an academic paper, medical diagnostic report, or legal brief is jointly produced by humans and AI, how should responsibility be assigned if it contains hallucinated facts, fabricated data, infringing material, or harmful advice? The traditional coupling of author and responsibility is strained in the complex network of human–machine collaboration. The process is also increasingly dynamic and fragmented. Users interact with AI through multiple prompts, and texts are produced through repeated revision and iteration. Human intention and algorithmic output become entangled, which obscures the boundary of originality and makes responsibility difficult to trace.

For this reason, this paper moves beyond the conventional debate about intellectual-property ownership and introduces the concept of version responsibility from the standpoint of technology policy and knowledge governance. In AI-era knowledge production, the central problem is no longer simply who created the knowledge, but who is responsible for the specific version of knowledge that enters the public sphere. The following sections use the concepts of cognitive compensation and the RID cognitive dynamics model in *Knowing and Saying* to analyze the deconstruction of traditional authorship, clarify the theoretical meaning of version responsibility, and propose an institutional framework for responsible human–machine collaborative knowledge production.

2 The Deconstruction of Traditional Authorship and the Cognitive Dilemma of Human–Machine Collaboration

2.1 The Internal Definition of Traditional Authorship: The Unity of Intentionality and Responsibility

In the traditional knowledge system, the author is not only a legal category but also a philosophical and cognitive one. From the perspective of cognitive dynamics, traditional authorship rests on two pillars: intentionality and responsibility accountability. These two pillars together give authorship its internal definition and allow the human subject to occupy the central position in knowledge production.

First, the author is the motivational source of knowledge generation. According to the RID model in *Knowing and Saying*, genuine cognition begins from problem pressure, namely the D-dimension of drive or demand [23]. A human author experiences cognitive imbalance and the motivation to seek a solution when facing concrete problems in survival, social interaction, or academic inquiry. This D-driven intentionality is the first impetus of knowledge production. Chalmers argues that subjective experience and intentionality are irreducible features of consciousness [4]. Searle's Chinese Room argument similarly shows that syntactic manipulation alone does not amount to semantic understanding [16]. Clark and Chalmers further explain that tools can extend cognition, but the subject of cognition remains anchored in human agency [6]. Whether in scientific inquiry or literary creation, a deeper sense of problem and value concern stands behind the text.

Second, the author transforms internal cognitive schemata into public symbolic text through structural generation and rule-governed expression. In the RID vocabulary, this means that the author moves from the D-dimension through the I-dimension of information or structure and the R-dimension of rule or representation. During this process, the author bears the burden of proof and social responsibility for concepts and propositions. Traditional authorship is therefore a closed loop of drive, structure, and expression. Intentionality and responsibility coincide in the human subject.

2.2 How Generative AI Separates Intentionality from Responsibility

Generative AI breaks the traditional chain between intentionality and responsibility. A large language model is essentially a probabilistic prediction machine trained on massive corpora. It possesses strong rule-like expressive capacity and simulated structural-generation capacity, but it lacks the D-dimension. It has no problem pressure. It does not ask questions out of curiosity, need, social concern, or existential disequilibrium. Sun's ontological critique of generative AI emphasizes that such systems probabilistically fit prior human knowledge but lack life experience and genuine concern [17]. Bender and colleagues describe large language models as stochastic parrots that assemble language patterns without understanding the world [2]. Dreyfus's critique of artificial reason remains relevant: without embodiment and situated understanding, computational systems cannot acquire human rationality in the full sense [7].

The absence of the D-dimension means that AI functions as a symbol manipulator rather than a bearer of meaning. It can generate academic-looking text, but it does not understand the causal relations behind that text and cannot assume ethical responsibility for the consequences. Its output depends on human prompts and institutional use. This produces a cognitive asymmetry: powerful capacities of expression and structure

are outsourced to a machine, while the dimension of problem pressure and responsibility remains with human beings.

The traditional concept of the author cannot fully capture this separation of capacity and responsibility. In traditional authorship, capacity and responsibility are broadly correlated: the more capable the author is, the greater the value and responsibility of the work. In human–machine collaboration, AI displays textual-generation capacity beyond ordinary human ability but bears no responsibility, whereas human users may have weaker generative capacity but remain responsible for the final text. This asymmetry deconstructs authorship as a unified category.

2.3 Black Boxes and Accountability Dilemmas in Knowledge Production

The direct consequence of this deconstruction is black-boxed knowledge production and the resulting difficulty of accountability. In traditional production, the path from literature reading and data collection to reasoning and writing is relatively traceable. Even when errors appear, the production process can be reconstructed. In human–machine collaboration, however, a large interpretive gap separates human prompts from AI outputs.

Deep neural networks generate content through parameterized transformations that are not readable at the level of individual concepts and decisions. Even model designers often cannot reconstruct why a specific sentence, citation, or concept was generated. This lack of explainability may be tolerable in casual conversation, but it is unacceptable in serious knowledge production. If AI-generated content is used in high-risk contexts such as medical decision support, automated driving, or legal reasoning, serious errors can be difficult to trace and correct.

Human users may claim that they only used a tool and could not foresee hallucination or bias. Developers may claim that they merely provided a general-purpose technology and are not responsible for concrete downstream consequences. This suspension and diffusion of responsibility is the social-governance expression of the failure of traditional authorship. Duan's work on ethical regulation and value alignment stresses that ethical governance must be embedded throughout the process of AI application [8]. Chen likewise emphasizes that AI ethics requires clarifying human responsibility during technological evolution and preventing social risk caused by technological alienation [5].

3 Theoretical Construction of Version Responsibility

3.1 From Originality to Responsibility

The crisis of authorship cannot be resolved by minor adjustments within the framework of originality and intellectual-property ownership. Originality emphasizes the creation of something from nothing. In the AI era, however, much knowledge is recombined and optimized from existing data, and pure originality becomes increasingly difficult to define. A shift from an originality paradigm to a responsibility paradigm is therefore required.

Floridi's theory of the logic of information treats philosophy as conceptual design: when the infosphere changes fundamentally, new conceptual frameworks are needed to address moral and cognitive challenges [10]. Version responsibility is such a conceptual design. It means that in a network of human–machine collaborative knowledge production, regardless of whether a particular element is handwritten, AI-generated, or hybrid, the subject who releases a specific public version, such as a published article, diagnostic report, legal document, or software release, must assume final responsibility for that version's truthfulness, legality, and ethical consequences.

This concept suspends the endless question of who the substantive creator is and focuses instead on who serves as the gatekeeper of public knowledge. Once knowledge is published, it is no longer merely private material. It enters a public space and can influence policy, medicine, law, education, and social trust. The publisher of the version must therefore guarantee not only quality but also public accountability.

3.2 Three Dimensions of Version Responsibility

Version responsibility contains three mutually supporting dimensions.

The first is verification responsibility. Because AI can produce hallucinations, conceptual drift, fabricated sources, and false data, directly transforming AI output into public knowledge is dangerous. Ji and colleagues show that hallucination is a common structural problem in natural language generation: models can generate fluent content that is not grounded in facts [13]. The version-responsible subject must therefore check factual statements, data citations, and logical steps before publication. In academic publishing, this means verifying that AI-assisted passages do not contain fabricated references or false evidence. In journalism, it means independently checking AI-generated background information.

The second is disclosure responsibility. Recipients of knowledge, including readers, patients, judges, and citizens, have a right to know the extent to which AI participated in generating the information they receive. Disclosure should include the model or tool used, the stage of use, and where relevant the nature of prompting and revision. Transparency is not a ritual. It is the procedural basis of trust.

The third is consequence accountability. When a specific version causes academic misconduct, plagiarism, defamation, infringement, malpractice, or other harm, the version-responsible subject cannot avoid responsibility by claiming that the mistake was generated by AI. The act of stabilizing AI-generated content into a final version and releasing it to the public is an act of endorsement. Endorsement entails accountability.

3.3 The Philosophical Foundation: Defending Human Priority in the D-Dimension

Philosophically, version responsibility defends a central insight of the RID model: in human–machine collaboration,

the priority of human beings in the D-dimension must be maintained. AI can greatly extend human capacity in the I-dimension and R-dimension. It can process information, generate structures, and produce language at a scale previously unavailable. Yet this compensation must remain under the constraint of human problem pressure. Machines do not possess values, moral conscience, or a pursuit of truth in the human sense.

To assume version responsibility is to anchor AI's structural and expressive outputs in human life interests, social credibility, and moral judgment. Only when a human being or human institution is willing to be responsible for a version does algorithmically generated language become legitimate public knowledge. Version responsibility is therefore not a technical rule alone. It is a defense of human subjectivity in the knowledge ecology.

4 Institutional Design for Human–Machine Collaborative Knowledge Production

4.1 Functional Outsourcing and Responsibility Centralization

The core institutional principle can be summarized as functional outsourcing and responsibility centralization. Repetitive or mechanical tasks, including retrieval, preliminary drafting, language polishing, summarization, or coding assistance, may be outsourced to AI. Value-laden tasks, including fact checking, final judgment, approval, and publication, must remain centralized in human subjects.

This principle is particularly important in academic publishing. O'Connor's discussion of open AI platforms in nursing education and Thorp's editorial on ChatGPT both converge on a clear norm: AI may be useful, but it is not an author and should not be treated as one [15, 18]. The authority of publication must not be confused with the capacity of generation.

4.2 Academic Publishing

In academic publishing, version responsibility requires a structured AI-use disclosure system, automated and manual screening for AI-generated content, and clear standards for misconduct. Authors should record whether AI was used for literature review, language editing, code generation, data analysis, or substantive drafting. Journals and universities should establish a tiered disclosure mechanism: minor language polishing may be disclosed briefly, whereas AI participation in literature review, analysis, or drafting requires detailed explanation.

If hallucinated references or unverified claims are published, the version holder should bear responsibility. The purpose is not to ban AI assistance, but to prevent the authority of publication from being mistaken for the authority of generation. Recent surveys also show that AI use in research is becoming common before institutional norms have fully stabilized. A Nature survey of more than 1,600 researchers reported broad expectations that AI tools will become increasingly central to science [20]. This gap between actual use and mature regulation strengthens the case for version responsibility.

4.3 Medical and Judicial High-Risk Use

In medical decision support, AI should never become the final decision maker. Physicians must remain the version-responsible subjects of diagnoses, treatment plans, and patient communication. Hospitals should add human–machine comparison nodes to clinical workflows, requiring physicians to conduct independent clinical and pathological analysis before adopting AI suggestions. Records of human verification should be retained in the medical file. Such a system protects patients and also clarifies physicians' professional responsibility.

In judicial contexts, AI-assisted legal documents or sentencing references require a double safeguard of algorithmic-transparency review and human verification. Judges and lawyers who submit AI-assisted legal materials must independently verify every statute, precedent, quotation, and citation. The *Mata v. Avianca* sanctions decision illustrates this point: attorneys submitted legal materials containing fabricated cases generated by ChatGPT and were sanctioned by the court [19]. The lesson is not merely that AI can hallucinate. It is that the legal professional who submits a final document remains responsible for its verification.

4.4 Developer Liability and Shared Governance

Emphasizing user version responsibility does not absolve developers. In a complete knowledge-governance framework, responsibility must be distributed along the whole chain. If a model contains severe design defects, biased training data, insufficient risk warnings, or no reasonable verification support, developers may bear product-related or regulatory responsibility. Bender's warning about stochastic parrots and Floridi and Chiriatti's analysis of GPT-3 both show why model output cannot be treated as self-sufficient truth [2, 11].

The European Union's Artificial Intelligence Act offers a useful legislative reference by imposing risk-based requirements on high-risk AI systems, including transparency, data governance, and human oversight [9]. Future AI legislation should similarly define the responsibilities of developers, providers, deployers, and users across different scenarios. The goal is not to transfer all risk to end users, but to build a layered responsibility chain in which the final version holder is accountable for publication while developers remain responsible for design and foreseeable risk.

4.5 Empirical Evaluation of Governance Effectiveness

Institutional design must be accompanied by empirical evaluation. Version responsibility requires operational indicators, including the rate and quality of AI-use declarations, the effectiveness of academic-integrity screening for fabricated or AI-generated content, and the maturity of dispute-handling mechanisms for improper AI use. These indicators should be tracked and periodically revised. Without data-based institutional evaluation, responsibility centralization may become only a slogan.

5 Version Responsibility in Different Modes of Knowledge Production

5.1 AI-Assisted Polishing

In the AI-assisted polishing mode, the human author has already completed the main argument, data analysis, and initial drafting. AI is used only for grammar, wording, style, or translation. This is currently the most widely accepted form of human–AI collaboration in academic and professional contexts.

In this mode, the core responsibility is disclosure. Because AI has not substantively altered the D-dimension or I-dimension of the work, the fact-checking burden is relatively light. Nonetheless, academic integrity requires authors to disclose that AI was used for language polishing or translation. Disclosure respects readers' right to know and prevents misunderstandings caused by sudden stylistic changes.

5.2 AI Content Generation

In the AI content-generation mode, human authors provide central ideas, research frameworks, or detailed prompts, while AI generates substantial passages, literature summaries, or code. This mode significantly increases efficiency but also increases risk.

Here, version responsibility shifts toward verification and consequence accountability. AI may mix fabricated information with real information. The human author must act as a gatekeeper, checking each claim, data point, and citation. If an author publishes erroneous AI-generated content through negligence, the author must bear the consequences, including retraction, reputational harm, or legal liability.

5.3 AI-Led Discovery

AI-led discovery is the frontier of knowledge production. Examples include systems that predict protein structures, mine scientific literature, or identify hidden correlations in large datasets. In this mode, AI is not merely an expressive tool. It participates in structural generation and may even simulate aspects of problem exploration.

Because AI's reasoning process is highly black-boxed, version responsibility here concerns scientific validation and methodological transparency. Researchers cannot publish a result as scientific truth merely because a model produced it. They must design independent experiments, use multiple verification methods, and disclose model architecture, training data, and validation procedures so that the scientific community can reproduce and examine the result. In this mode, version responsibility is the human guarantee of the reliability of technological tools.

6 Challenges and Responses

6.1 Technical Challenge: Detecting AI-Generated Content

One precondition for version responsibility is the capacity to identify AI-generated components in knowledge products. Yet as generative models improve, AI text increasingly resembles human writing. Statistical detectors based on perplexity or burstiness are fragile, and human revision can further conceal AI traces.

Two responses are necessary. First, detection methods should move beyond surface statistics toward semantic, evidential, and logical analysis. Second, digital watermarking or provenance metadata should be embedded in model outputs where feasible, allowing traceability from the source. This requires cooperation among developers, standardization bodies, publishers, and regulators.

6.2 Institutional Challenge: Legal Lag

Existing copyright and liability systems were built around human authorship. They do not map cleanly onto human–machine collaboration. If a human is the version-responsible subject, should that person receive full copyright in a heavily AI-assisted work? If AI contributes most of the expression, is full human ownership fair? Current law has not reached a stable answer.

Legislative revision and judicial clarification are therefore needed. New legal categories, neighboring rights, or special protections for human–machine collaborative products may be required. More importantly, the law must clarify the relationship between rights and responsibilities: those who release a version to the public must assume responsibility for it.

6.3 Ethical Challenge: Cognitive Laziness and Responsibility Avoidance

AI collaboration can increase efficiency, but it can also produce cognitive laziness. When AI can draft literature reviews and summarize data, researchers may gradually lose habits of independent thought and deep inquiry. In organizations, responsibility may diffuse further as team members blame the model, the workflow, or one another.

The response is AI ethics education, digital literacy training, and explicit internal responsibility matrices. Professionals such as researchers, physicians, lawyers, and editors should be trained to use AI critically. Organizations should specify who is the final publisher and accountable subject in every human–AI collaborative project.

6.4 Case Analysis: *Mata v. Avianca*

The *Mata v. Avianca* case provides a vivid example of version responsibility. Attorneys submitted a filing containing fictitious cases generated by ChatGPT and were sanctioned by a federal court [19]. The lawyers attempted to explain that they had relied on the tool without understanding its tendency to fabricate legal materials. The court's reasoning made clear that technological novelty does not remove the lawyer's duty to verify legal authorities before filing them.

This case is an early judicial rehearsal of version responsibility. It shows that in a complex human–machine network, responsibility can still be fixed at the version node: the human subject who submits the final public document must be accountable for it.

7 Future Outlook: Building a Human–Machine Coexistent Knowledge Ecology

7.1 Upgrading and Standardizing Cognitive Tools

The purpose of version responsibility is not to restrict AI development, but to anchor it ethically and institutionally. Future AI systems should evolve from black-box generators into white-box or grey-box reasoning assistants that provide evidence sources and reasoning traces. This would reduce the cost of verification and improve reliability. Domain-specific models should also be developed and certified for medicine, law, education, and other high-risk fields.

Weizenbaum warned that tasks requiring human wisdom and empathy should not be handed over uncritically to computers [21]. Wiener similarly insisted that technology must serve human values [22]. Bostrom's concern about superintelligence reminds us that governance must remain ahead of capability [3]. These warnings remain relevant in the AI era.

7.2 Cultivating Digital Intellectuals

In the AI era, the figure of the intellectual will change. Future digital intellectuals will need not only professional knowledge but also the capacity to collaborate with AI responsibly. They should possess three capacities: a sharp sense of problems in the D-dimension, the structural ability to guide AI in the I-dimension, and rigorous critical thinking for verification and accountability. These capacities make version responsibility practical rather than merely rhetorical.

7.3 Rebuilding Trust in Public Knowledge

Trust is the foundation of public knowledge. In an age of abundant AI-generated content, trust cannot rest only on institutional prestige. It must also rest on transparency and traceability. Future academic outputs may need to include AI-use statements, version logs, or provenance records. The purpose is not surveillance, but accountable knowledge production. A healthy knowledge ecology is one in which AI performs functions, human beings retain responsibility, and public versions remain answerable to truth, legality, and ethical consequence.

8 Conclusion

The rise of generative AI marks the beginning of a new era of human–machine collaborative knowledge production. In this historical process, traditional authorship is deconstructed because intentionality and responsibility no longer coincide naturally in a single human subject. The response should be neither technological rejection nor technological determinism.

The theory of version responsibility offers a pragmatic and philosophically grounded path for reconstructing AI-era knowledge governance. It redirects attention from whether AI has creativity to who is responsible for the social consequences of knowledge. By establishing verification responsibility, disclosure responsibility, and consequence accountability, and by applying the principle of functional outsourcing and responsibility centralization in academic publishing, medicine, law, and other fields, society can use AI's powerful structural and expressive capacities while preserving human priority in problem pressure and value judgment.

Human–machine collaboration should not mean the withdrawal of human subjectivity. It should mean the upgrading of human cognitive responsibility. Only when every publisher of knowledge assumes version responsibility can we rebuild a truthful, reliable, and vibrant public knowledge ecology in the flood of AI-generated information.

References

[1] R. Barthes, "The death of the author," in *Image, Music, Text*. New York: Hill and Wang, 1977, pp. 142–148.

[2] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 610–623. [Online]. Available: https://doi.org/10.1145/3442188.3445922

[3] N. Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.

[4] D. J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press, 1996.

[5] X. Chen, "The ethical foundations and technological evolution of artificial intelligence," *China Soft Science*, no. 8, pp. 55–63, 2022, chinese-language source; translated title retained from the source manuscript; exact public primary metadata was not accessible in this audit.

[6] A. Clark and D. J. Chalmers, "The extended mind," *Analysis*, vol. 58, no. 1, pp. 7–19, 1998. [Online]. Available: https://doi.org/10.1093/analys/58.1.7

[7] H. L. Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press, 1992.

[8] W. Duan, "Ethical regulation and value alignment in artificial intelligence," *Studies in Philosophy of Science and Technology*, vol. 41, no. 2, pp. 33–41, 2024, chinese-language source; translated title retained from the source manuscript; exact public primary metadata was not accessible in this audit.

[9] European Parliament and Council of the European Union, "Regulation (eu) 2024/1689 laying down harmonised rules on artificial intelligence (artificial intelligence act)," 2024, official Journal of the European Union. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[10] L. Floridi, *The Logic of Information: A Theory of Philosophy as Conceptual Design*. Oxford: Oxford University Press, 2019.

[11] L. Floridi and M. Chiriatti, "Gpt-3: Its nature, scope, limits, and consequences," *Minds and Machines*, vol. 30, no. 4, pp. 681–694, 2020. [Online]. Available: https://doi.org/10.1007/s11023-020-09548-1

[12] M. Foucault, "What is an author?" in *The Foucault Reader*, P. Rabinow, Ed. New York: Pantheon Books, 1984, pp. 101–120.

[13] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung, "Survey of hallucination in natural language generation," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, 2023. [Online]. Available: https://doi.org/10.1145/3571730

[14] X. Liu, *The Philosophical Foundations of Cognitive Science and Embodied Cognition*. Beijing: Science Press, 2021, chinese-language monograph; translated title retained from the source manuscript.

[15] S. O'Connor, "Open artificial intelligence platforms in nursing education: Tools for academic progress or abuse?" *Nurse Education in Practice*, vol. 66, p. 103537, 2023. [Online]. Available: https://doi.org/10.1016/j.nepr.2022.103537

[16] J. R. Searle, "Minds, brains, and programs," *Behavioral and Brain Sciences*, vol. 3, no. 3, pp. 417–424, 1980. [Online]. Available: https://doi.org/10.1017/S0140525X00005756

[17] W. Sun, "An ontological critique of generative artificial intelligence," *Studies in Dialectics of Nature*, vol. 39, no. 11, pp. 3–10, 2023, chinese-language source; translated title retained from the source manuscript; exact public primary metadata was not accessible in this audit.

[18] H. H. Thorp, "Chatgpt is fun, but not an author," *Science*, vol. 379, no. 6630, p. 313, 2023. [Online]. Available: https://doi.org/10.1126/science.adg7879

[19] United States District Court for the Southern District of New York, "Mata v. avianca, inc., 678 f. supp. 3d 443," 2023, sanctions decision involving ChatGPT-generated fictitious legal citations. [Online]. Available: https://www.courtlistener.com/docket/63107798/mata-v-avianca-inc/

[20] R. Van Noorden and J. M. Perkel, "Ai and science: What 1,600 researchers think," *Nature*, 2023. [Online]. Available: https://doi.org/10.1038/d41586-023-02980-0

[21] J. Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: W. H. Freeman, 1976.

[22] N. Wiener, *The Human Use of Human Beings: Cybernetics and Society*. Boston: Houghton Mifflin, 1954.

[23] X. Zhang, "Knowing and saying: An ontological investigation of human cognition," PSSXiv (Philosophy and Social Sciences Preprint Server), may 2026, pSSXiv:202605.04152V1; CSTR:32012.36.PSSXiv.202605.04152. [Online]. Available: https://zsyyb.cn/abs/202605.04152

[24] T. Zhao, "The epistemological boundary of artificial intelligence," *Philosophical Research*, no. 5, pp. 15–26, 2023, chinese-language source; translated title retained from the source manuscript; exact public primary metadata was not accessible in this audit.