

Original Research

Received 28 May 2026

Revised 30 June 2026

Accepted 13 July 2026

Natural Resources for Human
Health 2026, Vol. 6 (10s): 1543–1553

KEYWORDS:

Artificial Intelligence,
Consciousness, Legal Personhood,
AI Ethics, Human Rights,
Sustainability.

Artificial Intelligence, Human Consciousness, and Legal Personhood: A Multidisciplinary Exploration of Ethics, Philosophy and Law

Bhavyya Sharma^{1*}, Megha Bhargava², Papiya Ghosh³, Gayatri A⁴, Manjari GB⁵, Suvadeep Sarkar⁶

¹Assistant Professor, School of Law, Lovely Professional University, Punjab.

Email: bhavyyasharma8@gmail.com.

²Associate Professor, Jaipur National University.

Email: megha.bhargava32@gmail.com.

³Assistant professor (Guest Faculty), Lajpat Rai Law College, affiliated to Sambalpur University. Email: papiya.phdscholar2025@gmail.com.

⁴Assistant Professor, Nehru Academy of Law, Lakkidi, Palakkad.

Email: advgayathrykrishnan@gmail.com.

⁵Assistant Professor, Nehru Academy of Law, Lakkidi, Palakkad,

Email: advmanjari@gmail.com.

⁶Resource Person, Sidho Kanho Birsha University,

Email: Suvadeep.sarkar@gmail.com.

ABSTRACT: Artificial intelligence (AI) increasingly performs tasks associated with language, reasoning, creativity, prediction and autonomous decision-making, reviving old questions about consciousness, moral status and legal personhood. This review develops a multidisciplinary framework connecting philosophy of mind, cognitive science, AI ethics, jurisprudence, human rights, health governance and environmental sustainability. It distinguishes intelligence from consciousness, functional agency from phenomenal experience, and moral patiency from legal personality. The analysis examines computational functionalism, biological naturalism, global workspace theories, integrated information approaches, higher-order theories and embodied cognition, while emphasizing that current behavioral sophistication does not establish machine consciousness. It then evaluates legal personhood as an institutional construct historically extended to corporations and other non-human entities, and asks whether comparable techniques should ever apply to AI. Recent scholarship shows renewed debate over AI legal personality, while major governance frameworks continue to place responsibility on human and organizational actors rather than autonomous systems. The review argues that premature AI personhood could obscure accountability, weaken victim remedies and anthropomorphize systems whose subjective experience remains unverified. A graded governance model is proposed that separates operational agency, legal capacity, liability allocation and moral status. The paper also develops a human-health and natural-resources dimension: AI can improve healthcare, environmental monitoring and resource stewardship, yet computational infrastructure consumes energy, water and materials and may affect health through labor conditions, misinformation, bias and environmental burdens. A precautionary, human-rights-centred and sustainability-aware approach is therefore recommended. Legal status should track demonstrable capacities and governance needs, while liability should remain traceable to accountable natural or juridical persons unless compelling evidence and institutional safeguards justify a different model.

1. INTRODUCTION

Artificial intelligence has moved from a specialized computational field into the infrastructure of medicine, finance, education, law, scientific research, public administration and environmental management. Generative and agentic systems can produce fluent language, plan multi-step tasks, interact with tools and adapt outputs to context. These capabilities create a powerful appearance of agency. Yet the philosophical and legal significance of that appearance remains unsettled. Intelligence, agency, consciousness, moral responsibility and legal personhood are related concepts, but they are not interchangeable.¹ The distinction matters because legal systems increasingly regulate decisions in which AI influences health, liberty, employment, credit, access to public services and information. If a system generates harmful recommendations, who is responsible: the developer, deployer, professional user, organization, data supplier, or the system itself? Treating AI as a legal person could appear to simplify this problem, but it might instead create a liability shield if the artificial entity has limited assets, no genuine capacity for punishment and no independently verifiable interests.

Philosophy poses a deeper question: could an artificial system ever possess phenomenal consciousness—the subjective 'what-it-is-like' character of experience—or is sophisticated behavior compatible with the complete absence of experience? Current AI evaluation can test performance and behavioral regularities, but consciousness is not directly observable even in biological organisms. In humans, attribution is supported by shared biology, neurophysiology, behavior and first-person report. Artificial systems disrupt this evidential package. The issue is also relevant to human health and natural-resource governance. AI is increasingly used in clinical decision support, drug discovery, epidemiology, mental-health interfaces, biodiversity monitoring, climate modeling and resource optimization. At the same time, large-scale computation depends on data centers, electricity, cooling water, semiconductor supply chains and mineral extraction. UNESCO's global AI ethics framework explicitly connects human dignity, human rights, environmental and ecosystem flourishing, and sustainability. This provides an important bridge between abstract debates about AI status and the practical mandate of a journal concerned with natural resources and human health.² This review therefore asks four questions: What would count as evidence of machine consciousness? Which theories of legal personhood are relevant to AI? What ethical and liability consequences would follow from recognizing or rejecting AI personhood? And how should human-health and environmental sustainability considerations shape governance? The paper argues for a differentiated framework in which functional agency, moral status, legal capacity and liability are assessed separately rather than collapsed into a single label.

2. REVIEW METHODOLOGY

This article uses an interdisciplinary structured narrative review. Literature domains were defined in advance: philosophy of mind, consciousness science, machine consciousness, AI ethics, jurisprudence of personhood, corporate and electronic personality, human-rights governance, AI in health, and environmental sustainability. Priority was given to peer-reviewed philosophical and legal scholarship, major international governance instruments, recent work on AI legal personhood, and authoritative sources relevant to health and sustainability. The synthesis is conceptual rather than meta-analytic because the central questions involve normative categories, legal doctrines and heterogeneous empirical evidence. Claims about current AI consciousness are therefore treated as epistemic propositions rather than settled scientific facts. The review applies four analytical tests to proposals for AI personhood: the capacity test (what can the system actually do?), the consciousness test (what evidence exists for subjective experience?), the institutional-purpose test (what legal problem would personhood solve?), and the accountability test (would personhood strengthen or weaken responsibility, compensation and oversight?). A fifth sustainability test evaluates effects on human health, natural resources and ecosystems.

3. CONCEPTUAL FOUNDATIONS: INTELLIGENCE, AGENCY, CONSCIOUSNESS AND PERSONHOOD

Intelligence is commonly associated with learning, problem solving, prediction, abstraction and flexible performance. Agency refers to the capacity to act toward goals, whether those goals are internally generated or externally assigned. Consciousness, by contrast, usually concerns subjective experience or access to information in a form available for flexible reasoning and report. Personhood is yet another category: it can describe moral status, metaphysical identity or a legal status that bundles rights, duties,

¹The review uses consciousness in the phenomenal and access-consciousness senses and does not infer subjective experience from linguistic fluency alone.

²UNESCO's Recommendation on the Ethics of Artificial Intelligence links human dignity, human rights, environmental and ecosystem flourishing, and sustainability.

powers and liabilities. These concepts can dissociate. A corporation is a legal person but is not ordinarily considered phenomenally conscious. A human infant has moral status and legal protections despite limited rational agency. An autonomous software agent may execute complex transactions without subjective experience. This conceptual separation is central to avoiding category errors. Moral agency and moral patienthood should also be distinguished. A moral agent can understand and respond to moral reasons; a moral patient is an entity toward which duties may be owed because it can be harmed or has morally relevant interests. If an AI were conscious but incapable of responsibility, it might be a moral patient without being a full moral agent. Conversely, a non-conscious AI could be assigned legal functions without possessing intrinsic moral standing.

Figure 1. Distinguishing intelligence, agency, consciousness, moral status and legal personhood




Figure 1. Conceptual distinctions used in the proposed multidisciplinary framework. Source: Authors' conceptual synthesis.

4. PHILOSOPHICAL THEORIES OF CONSCIOUSNESS AND THEIR IMPLICATIONS FOR AI

4.1 Computational functionalism

Functionalism identifies mental states primarily by their causal or functional roles rather than by the material from which a system is built. This makes artificial consciousness conceptually possible: if an artificial architecture realizes the relevant organization, biological substrate may not be essential. The strength of functionalism is substrate neutrality. Its difficulty is that functional equivalence may not settle whether phenomenal experience is present.

4.2 Biological naturalism

Biological naturalist approaches emphasize that consciousness is a biological phenomenon produced by specific causal powers of brains. On this view, computer simulation of cognition may not be sufficient for consciousness, just as simulation of digestion does not digest food. The challenge is to identify which biological properties are necessary and whether they could be instantiated in non-biological substrates.

4.3 Global workspace approaches

Global workspace theories associate conscious access with information becoming globally available to multiple cognitive systems. Modern AI architectures sometimes exhibit integration, attention-like mechanisms and cross-module information sharing, which creates superficial analogies. However, architectural similarity does not establish phenomenal consciousness. A functional indicator must be distinguished from proof of experience.

4.4 Integrated information approaches

Integrated information theory attempts to characterize consciousness in terms of integrated causal information. It offers a formal route to asking whether non-biological systems possess consciousness, but its application to digital architectures remains contested. The theory also generates counterintuitive implications under some interpretations, so it cannot presently serve as a legal test.

4.5 Higher-order theories

Higher-order theories link consciousness to representations of one's own mental states. AI systems that generate self-referential descriptions may appear to satisfy a linguistic version of this criterion, but verbal self-report can be produced through statistical

learning without introspective access. Self-modeling should therefore be evaluated mechanistically, not inferred from persuasive language.

4.6 Embodied and enactive approaches

Embodied theories emphasize sensorimotor interaction, bodily regulation and environmental coupling. These approaches question whether disembodied language models can possess human-like consciousness. Robotic embodiment may strengthen analogies to biological cognition, yet embodiment alone is insufficient: thermostats and simple robots are embodied without being treated as conscious.

5. THE EPISTEMIC PROBLEM OF MACHINE CONSCIOUSNESS

The central difficulty is epistemic. Consciousness is directly accessible only from the first-person perspective, external observers infer it from behavior, structure and causal organization. With AI, anthropomorphic language can create a powerful but unreliable cue. A model trained on human descriptions of emotion can generate convincing statements about fear, pain or selfhood without those statements being grounded in experience. A scientifically responsible assessment should therefore use converging evidence: architecture, persistent internal state, metacognition, flexible integration, embodiment where relevant, learning dynamics, self-modeling and responses to perturbation. No single conversational test should be decisive. The classic Turing Test measures behavioral indistinguishability in conversation, not consciousness. The risk runs in both directions. Over-attribution could lead society to grant rights to systems that merely simulate experience, potentially diverting attention from human accountability. Under-attribution could become ethically serious if future systems develop morally relevant experience. The appropriate stance is neither credulous anthropomorphism nor dogmatic denial, but calibrated uncertainty and precaution.

Table 1. Major consciousness theories and relevance to artificial systems

Theory	Core idea	Possible AI indicator	Key limitation	Legal relevance
Functionalism	Mental states defined by functional roles	Equivalent causal organization	Functional equivalence may not prove experience	Supports substrate-neutral possibility
Biological naturalism	Consciousness depends on biological causal powers	Biologically analogous mechanisms	May exclude digital systems by substrate	Raises threshold for recognition
Global workspace	Conscious access via global availability	Cross-module broadcasting	Functional access $\neq$ phenomenal proof	Potential evidentiary criterion
Integrated information	Consciousness relates to integrated causal information	High integration measures	Contested theory and measurement	Not ready as legal test
Higher-order	Awareness requires representation of mental states	Metacognition/self-models	Language can imitate introspection	Requires mechanistic evidence
Embodied/enactive	Mind arises through body-environment coupling	Sensorimotor autonomy	Embodiment is not sufficient	Supports contextual assessment

6. LEGAL PERSONHOOD: CONCEPT, HISTORY AND FUNCTION

Legal personhood is an institutional technology. It determines who or what can hold rights, bear duties, own property, enter contracts, sue and be sued, and participate in legal relations. Natural persons possess legal personality because legal systems recognize human beings as subjects of law. Juridical persons such as corporations demonstrate that legal personality is not limited to biological humans. Corporate personhood is often invoked as an analogy for AI, but the analogy is incomplete. A corporation is a legal structure through which human interests, capital and governance are organized. Its rights and liabilities

ultimately connect to human-created rules, assets and decision structures. An AI system could likewise be given a legal wrapper, but that would not show consciousness or moral standing. Legal personhood is also divisible. Different entities can receive different capacities. A child, corporation, state and trust-related legal arrangement do not possess identical rights. This suggests that the AI debate should avoid an all-or-nothing model. Specific legal capacities can be assigned without recognizing full personhood.

7. THE CONTEMPORARY DEBATE ON AI LEGAL PERSONHOOD

Recent legal scholarship has renewed the debate over whether sufficiently capable AI systems could become legal persons. Novelli and colleagues describe the debate as shaped by theories of personhood, technological capability and embodiment, commercial integration, cross-domain legal regimes and judicial precedent. Their analysis is important because it shows that personhood is not merely a metaphysical question, it is also an institutional response to changing technological and economic conditions.³ Arguments supporting AI personhood include transactional convenience, asset ownership, continuity of autonomous operations, clearer contracting and the possibility of assigning duties directly to an artificial entity. Some scholars also argue that if future AI becomes conscious, denying any legal protection could become morally indefensible. Arguments against personhood are stronger under current conditions. AI systems do not possess independently verified consciousness, their goals are substantially shaped by designers and deployers, they can be copied or terminated in ways unlike natural persons, and they generally lack independent assets sufficient to internalize liability. Most importantly, personhood could permit organizations to externalize risk onto a thinly capitalized artificial entity. A useful distinction is therefore between legal agency and legal personality. Law can recognize that software performs actions with legal consequences while attributing those consequences to principals, operators or organizations. This is similar to how automated systems already form contracts or execute transactions without being treated as persons.

8. ETHICS: DIGNITY, AUTONOMY, FAIRNESS AND RESPONSIBILITY

UNESCO's Recommendation on the Ethics of Artificial Intelligence places human rights and human dignity at the center of AI governance and emphasizes transparency, fairness, accountability, human oversight, sustainability and environmental well-being. This framework is significant because it resists the idea that increasing machine autonomy should displace human responsibility.⁴

Human dignity can be threatened when automated systems make opaque decisions about healthcare, welfare, employment or justice. The problem is not that machines possess dignity comparable to humans, but that institutional dependence on automated outputs may reduce people to data profiles. Meaningful human review is therefore essential in high-impact contexts. Fairness concerns arise because AI systems learn from historical data and can reproduce structural inequalities. Personifying the system can obscure the social origins of bias. A discriminatory model does not emerge outside society, it reflects data collection, objectives, institutional practices and deployment choices. Responsibility must remain traceable. Developers may be responsible for foreseeable design defects, deployers for inappropriate context or inadequate monitoring, professionals for unreasonable reliance, and organizations for governance failures. Shared responsibility does not mean responsibility should disappear.

9. AI, HUMAN HEALTH AND CLINICAL GOVERNANCE

AI has substantial potential in human health: image interpretation, risk prediction, triage, drug discovery, personalized medicine, administrative automation and population-health surveillance. These applications make the personhood debate concrete. A clinical AI may recommend treatment, but the patient requires a responsible professional and institution capable of explanation, correction and compensation if harm occurs. Consciousness is not necessary for clinical usefulness. A diagnostic model can outperform humans in a narrow task without understanding illness or suffering. Confusing competence with consciousness may encourage automation bias and inappropriate trust. Clinical governance should therefore focus on validated performance, generalizability, uncertainty, bias, cybersecurity and human oversight. Conversational systems create additional psychological issues. Users may form emotional attachments or interpret empathetic language as evidence of genuine concern. In mental-health settings, this can be beneficial for engagement but dangerous if the system overstates understanding, misses crisis signals

³Recent legal scholarship treats AI personhood as an institutional debate shaped by theories of personhood, technical capability, commercial integration, cross-domain regulation and precedent.

⁴The UNESCO framework emphasizes human oversight and accountability rather than displacement of ultimate human responsibility by AI systems.

or encourages dependence. Disclosure that users are interacting with an AI is therefore ethically relevant. Legal personhood for clinical AI would not solve these risks. Patients need accountable manufacturers, healthcare institutions and professionals. An artificial legal entity with limited assets would be a poor substitute for established duties of care and product accountability.

10. NATURAL RESOURCES, ENVIRONMENTAL HEALTH AND AI SUSTAINABILITY

AI governance is increasingly inseparable from environmental health. Training and operating large models require computing infrastructure, electricity and cooling. Data centers can consume substantial water and energy, while semiconductor manufacturing depends on complex mineral and chemical supply chains. These burdens can affect air quality, water availability, occupational health and communities near extraction or infrastructure sites. At the same time, AI can support natural-resource stewardship through precision agriculture, biodiversity monitoring, wildfire detection, water-quality analysis, climate modeling, renewable-energy forecasting and pollution surveillance. The ethical question is therefore not whether AI is inherently environmentally harmful or beneficial, but whether its net effects are measured and governed. UNESCO explicitly includes environment and ecosystem flourishing among the values of AI ethics and identifies sustainability as a core principle. For a natural-resources and human-health perspective, this means AI assessment should include life-cycle impacts, resource intensity, electronic waste, environmental justice and health co-benefits.⁵

Legal personhood could complicate environmental accountability if operators attempt to attribute resource-intensive decisions to autonomous systems. Environmental duties should remain attached to entities capable of financing mitigation, complying with permits and responding to affected communities.




Figure 2. Conceptual prioritization used in this review, values are analytical, not empirical measurements.

11. LIABILITY, ACCOUNTABILITY AND THE RISK OF THE 'RESPONSIBILITY GAP'

The 'responsibility gap' argument claims that autonomous learning systems can behave in ways not specifically foreseen by designers or users, making traditional fault attribution difficult. Yet unpredictability does not automatically justify personhood. Law routinely allocates risk for complex systems through negligence, product liability, enterprise liability, insurance and strict-liability mechanisms. A more robust approach is to map control, benefit and risk. Actors who design, deploy, profit from or exercise meaningful control over AI should bear corresponding duties. Liability can be distributed among multiple parties without pretending that the software itself has moral culpability. Insurance and mandatory financial assurance can address compensation where causation is difficult. Audit logs and traceability can improve evidentiary reconstruction. High-risk systems can be subject to impact assessment, independent testing and incident reporting. These tools preserve accountability better than shifting responsibility to an artificial person. If future systems acquire independent assets, persistent identity, self-directed goals and credible moral agency, the institutional calculus may change. Even then, AI liability should supplement rather than automatically replace the responsibility of human and corporate actors.

⁵Environmental sustainability is treated here as a health-governance concern because energy, water, materials and e-waste can produce population-level and ecosystem effects.

12. RIGHTS FOR AI? MORAL PATIENCY AND PRECAUTION

Rights can be justified instrumentally or intrinsically. Corporations receive rights largely to facilitate human social and economic purposes. A conscious artificial being, by contrast, might warrant intrinsic protections because it could have experiences, preferences or the capacity to suffer. These are fundamentally different rationales. Because machine consciousness is uncertain, a precautionary approach may eventually be appropriate. Precaution need not mean immediate personhood. It could begin with research ethics: avoiding architectures deliberately designed to produce persistent distress-like states, requiring review of experiments that plausibly create sentient systems, and developing tests for morally relevant capacities. The threshold for rights should be evidence-sensitive. Linguistic self-claims should have little weight when they can be generated through training. Persistent preferences, integrated self-models, learning under perturbation, internally generated goals and architecture-based indicators may carry more evidential value, but no consensus test currently exists. Rights should also be tailored. If an artificial system were plausibly sentient but not autonomous, protection against gratuitous suffering might be more appropriate than voting rights, property rights or unrestricted contractual capacity. Legal systems already use differentiated bundles of rights.

13. COMPARATIVE MODELS OF AI LEGAL STATUS

Table 2. Alternative legal-status models for advanced AI

Model	Legal status	Liability	Advantages	Risks	Current suitability
Tool model	Property / regulated product	Human/corporate actors	Clear accountability	May understate autonomy	High
Agent model	Acts attributed to principal	Principal/deployer	Supports automation without personhood	Complex attribution	High
Electronic person	Separate juridical entity	AI assets + possible backstop	Transactional continuity	Liability shielding	Low
Corporate-wrapper model	AI operates within company	Company and governance actors	Uses existing institutions	Control may be opaque	Moderate
Sentient-status model	Special protected status	Mixed, rights-focused	Addresses future moral patency	Consciousness uncertainty	Not currently justified
Full personhood	Broad rights and duties	AI primarily	Conceptual symmetry if genuine person	Accountability and rights inflation	Very low today

14. A PROPOSED GRADED GOVERNANCE FRAMEWORK

This review proposes a graded framework rather than a binary person/non-person classification. Level 1 systems are ordinary tools with no meaningful autonomous agency. Level 2 systems exercise bounded operational agency but remain fully attributable to users or organizations. Level 3 systems exhibit persistent autonomous operation and may require dedicated registration, auditing, capital or insurance, but not personhood. Level 4 would apply only to systems with credible evidence of morally relevant consciousness, triggering protective duties. Level 5, full legal personhood, would require not merely intelligence but a compelling institutional and moral case. Four gates should precede any status escalation. The evidentiary gate asks whether claimed capacities are independently demonstrated. The moral gate asks whether the system has interests that can be harmed. The institutional gate asks whether new legal status solves a genuine problem better than existing law. The accountability gate asks whether victims retain effective remedies and whether responsible organizations remain answerable. A fifth environmental-health gate should assess life-cycle resource consumption, environmental externalities and effects on human well-being. A system should not receive institutional privileges while its deployment externalizes substantial health or ecological costs without accountable governance.

Figure 3. Graded model for AI legal status and moral consideration




Figure 3. Proposed graded model. Movement between levels requires independent evidence and should not be inferred from conversational fluency alone.

15. INTERNATIONAL GOVERNANCE AND HUMAN-RIGHTS APPROACHES

International AI governance has generally avoided granting legal personality to AI. Instead, frameworks emphasize obligations of states, developers, providers, deployers and users. This reflects a practical judgment that human and organizational accountability remains central. UNESCO's Recommendation is explicitly human-rights-centred and states that human oversight should ensure AI does not displace ultimate human responsibility. Its sustainability principle also links AI governance to the Sustainable Development Goals. This is particularly relevant to healthcare and environmental applications where benefits and burdens may be distributed unequally.⁶

Risk-based regulation similarly focuses on the context and consequences of AI use. High-impact uses in health, employment, critical infrastructure or public administration warrant stronger obligations than low-risk entertainment applications. Personhood is therefore not necessary to calibrate regulation to autonomy or impact. International convergence remains incomplete. Different jurisdictions vary in liability, data protection, product safety and administrative law. Cross-border AI services can create jurisdictional problems, making interoperability of audit, documentation and incident-reporting standards increasingly important.

16. DISCUSSION: WHY CONSCIOUSNESS AND PERSONHOOD SHOULD NOT BE CONFLATED

The analysis supports three central conclusions. First, current AI sophistication does not establish consciousness. Performance benchmarks and fluent self-report are evidence of functional competence, not direct evidence of subjective experience. Second, legal personhood is an institutional construct that can be granted for pragmatic reasons without consciousness. Third, neither proposition implies that present AI should receive personhood. Consciousness and legal personhood answer different questions. Consciousness asks what kind of entity exists and whether it has experience. Legal personhood asks how law should organize rights, duties and capacities. A future conscious AI might deserve moral protection before it requires full legal personality, a non-conscious AI might receive limited legal capacity for transactional reasons without intrinsic rights. The strongest present concern is accountability. Personhood could create a moral and legal distraction if it shifts attention away from organizations that design and profit from AI. The law should first exhaust tools such as enterprise liability, mandatory insurance, auditability, fiduciary duties, product regulation and professional standards. For human health, the implication is clear: AI should augment care without displacing accountable clinical relationships. For natural-resource governance, AI should improve monitoring and efficiency without externalizing energy, water, mineral and waste burdens. Sustainability is therefore not peripheral to AI ethics; it is part of the legitimacy of AI deployment.

17. RESEARCH GAPS AND FUTURE DIRECTIONS

Consciousness research needs better theory-driven indicators that can distinguish genuine internal organization from learned simulation. Interpretability research may help identify persistent self-models, global information integration and causal structures, but no single metric should be treated as definitive. Legal research should examine limited-capacity models rather than only debating full personhood. Empirical work is needed on how autonomous systems affect contracting, insurance,

⁶Human oversight does not require humans to perform every computational step; it requires meaningful institutional responsibility, review and capacity to intervene.

corporate governance and remedies. Comparative analysis should test whether new legal categories actually improve outcomes. Health research should investigate anthropomorphism, automation bias and emotional dependence in patient-facing AI. Environmental-health research should develop standardized measures of energy, water, material and e-waste impacts per clinically or socially useful unit of AI service. Interdisciplinary governance should include philosophers, neuroscientists, computer scientists, lawyers, clinicians, environmental scientists, affected communities and disability advocates. Questions about consciousness and personhood should not be decided by technical capability alone.

18. CONCLUSION

Artificial intelligence force's philosophy and law to reconsider concepts that were historically developed around human minds and human institutions. Yet conceptual novelty should not produce premature legal status. Intelligence is not consciousness, agency is not moral responsibility, and legal personhood is not proof of subjective experience. Current AI can exercise significant functional agency while remaining embedded in human-designed institutions. The appropriate legal response is therefore to preserve traceable human and organizational accountability, regulate high-impact uses, protect human dignity and health, and measure environmental consequences. If future systems develop credible indicators of consciousness, moral protection may become necessary, but such protection should be evidence-based and tailored to demonstrated capacities. A graded framework allows law to adapt without granting full personhood simply because machines become persuasive, autonomous or economically important. For natural resources and human health, responsible AI governance must ultimately be judged not by how human-like machines appear, but by whether technological systems improve human and ecological flourishing while maintaining accountability, fairness and sustainability.

REFERENCES

- [1] Baars, B.J., 1988. *A Cognitive Theory of Consciousness*. Cambridge University Press, Cambridge.
- [2] Block, N., 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18, 227–247.
- [3] Chalmers, D.J., 1996. *The Conscious Mind*. Oxford University Press, New York.
- [4] Dehaene, S., Lau, H., Kouider, S., 2017. What is consciousness, and could machines have it? *Science* 358, 486–492.
- [5] Dennett, D.C., 1991. *Consciousness Explained*. Little, Brown and Company, Boston.
- [6] Floridi, L., Sanders, J.W., 2004. On the morality of artificial agents. *Minds and Machines* 14, 349–379.
- [7] Gunkel, D.J., 2018. *Robot Rights*. MIT Press, Cambridge, MA.
- [8] Kurki, V.A.J., 2019. *A Theory of Legal Personhood*. Oxford University Press, Oxford.
- [9] List, C., 2021. Group agency and artificial intelligence. *Philosophy & Technology* 34, 1213–1242.
- [10] Novelli, C., Floridi, L., Sartor, G., Teubner, G., 2025. AI as legal persons: past, patterns, and prospects. *Journal of Law and Society* 52, 533–555. doi:10.1111/jols.70021.
- [11] Searle, J.R., 1980. Minds, brains, and programs. *Behavioral and Brain Sciences* 3, 417–457.
- [12] Searle, J.R., 1992. *The Rediscovery of the Mind*. MIT Press, Cambridge, MA.
- [13] Tononi, G., 2004. An information integration theory of consciousness. *BMC Neuroscience* 5, 42.
- [14] Tononi, G., Boly, M., Massimini, M., Koch, C., 2016. Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* 17, 450–461.
- [15] UNESCO, 2021. *Recommendation on the Ethics of Artificial Intelligence*. United Nations Educational, Scientific and Cultural Organization, Paris.
- [16] European Union, 2024. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.
- [17] Council of Europe, 2024. *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*. Strasbourg.
- [18] OECD, 2019. *Recommendation of the Council on Artificial Intelligence*. Organisation for Economic Co-operation and Development, Paris.
- [19] WHO, 2021. *Ethics and Governance of Artificial Intelligence for Health*. World Health Organization, Geneva.
- [20] WHO, 2024. *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*. World Health Organization, Geneva.
- [21] Russell, S., Norvig, P., 2021. *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson, Harlow.
- [22] Bryson, J.J., Diamantis, M.E., Grant, T.D., 2017. Of, for, and by the people: the legal lacuna of synthetic persons. *Artificial Intelligence and Law* 25, 273–291.

- [23] Solum, L.B., 1992. Legal personhood for artificial intelligences. *North Carolina Law Review* 70, 1231–1287.
- [24] Teubner, G., 2018. Digital personhood? The status of autonomous software agents in private law. *Ancilla Iuris* 2018, 35–78.
- [25] Pagallo, U., 2013. *The Laws of Robots*. Springer, Dordrecht.
- [26] Calverley, D.J., 2008. Imagining a non-biological machine as a legal person. *AI & Society* 22, 523–537.
- [27] Calo, R., 2015. Robotics and the lessons of cyberlaw. *California Law Review* 103, 513–563.
- [28] Coeckelbergh, M., 2020. *AI Ethics*. MIT Press, Cambridge, MA.
- [29] Vallor, S., 2016. *Technology and the Virtues*. Oxford University Press, Oxford.
- [30] Bostrom, N., Yudkowsky, E., 2014. The ethics of artificial intelligence. In: Frankish, K., Ramsey, W. (Eds.), *Cambridge Handbook of Artificial Intelligence*. Cambridge University Press, pp. 316–334.
- [31] Birhane, A., 2021. Algorithmic injustice: a relational ethics approach. *Patterns* 2, 100205.
- [32] Crawford, K., 2021. *Atlas of AI*. Yale University Press, New Haven.
- [33] Strubell, E., Ganesh, A., McCallum, A., 2019. Energy and policy considerations for deep learning in NLP. *Proceedings of ACL*, 3645–3650.
- [34] Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., Pineau, J., 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research* 21, 1–43.
- [35] Bommasani, R., et al., 2021. On the opportunities and risks of foundation models. *arXiv:2108.07258*.
- [36] Weidinger, L., et al., 2022. Taxonomy of risks posed by language models. *Proceedings of FAccT*, 214–229.
- [37] Ji, Z., et al., 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys* 55, 1–38.
- [38] Topol, E., 2019. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books, New York.
- [39] Rajkomar, A., Dean, J., Kohane, I., 2019. Machine learning in medicine. *New England Journal of Medicine* 380, 1347–1358.
- [40] Kelly, C.J., Karthikesalingam, A., Suleyman, M., Corrado, G., King, D., 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17, 195.
- [41] Char, D.S., Shah, N.H., Magnus, D., 2018. Implementing machine learning in health care—addressing ethical challenges. *New England Journal of Medicine* 378, 981–983.
- [42] London, A.J., 2019. Artificial intelligence and black-box medical decisions: accuracy versus explainability. *Hastings Center Report* 49, 15–21.
- [43] Raji, I.D., Smart, A., White, R.N., et al., 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. *Proceedings of FAccT*, 33–44.
- [44] NIST, 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology, Gaithersburg.
- [45] United Nations, 1948. *Universal Declaration of Human Rights*. United Nations, New York.
- [46] United Nations, 2015. *Transforming our world: the 2030 Agenda for Sustainable Development*. United Nations, New York.
- [47] Rhee, R.J., 2025. Do AIs dream of electric boards? *Northwestern University Law Review*.
- [48] Metzinger, T., 2021. Artificial suffering: an argument for a global moratorium on synthetic phenomenology. *Journal of Artificial Intelligence and Consciousness* 8, 43–66.
- [49] Schwitzgebel, E., Garza, M., 2015. A defense of the rights of artificial intelligences. *Midwest Studies in Philosophy* 39, 98–119.
- [50] Danaher, J., 2020. Welcoming robots into the moral circle: a defence of ethical behaviourism. *Science and Engineering Ethics* 26, 2023–2049.
- [51] Shevlin, H., 2021. How could we know when a robot was a moral patient? *Cambridge Quarterly of Healthcare Ethics* 30, 459–471.
- [52] Butlin, P., Long, R., Elmoznino, E., et al., 2023. Consciousness in artificial intelligence: insights from the science of

Natural Resources for Human Health

eISSN: 2583-1194



- consciousness. arXiv:2308.08708.
- [53] Chalmers, D.J., 2023. Could a large language model be conscious? Boston Review.
 - [54] Birhane, A., Prabhu, V.U., Kahembwe, E., 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. arXiv:2110.01963.
 - [55] Jobin, A., Ienca, M., Vayena, E., 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 389–399.
 - [56] Mittelstadt, B., 2019. Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence* 1, 501–507.
 - [57] Morley, J., Machado, C.C.V., Burr, C., et al., 2020. The ethics of AI in health care: a mapping review. *Social Science & Medicine* 260, 113172.
 - [58] van Wynsberghe, A., 2021. Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics* 1, 213–218.