

JEAI-26-51

Wisdom as a Service (WaaS) - A Pluralist Ethical Framework for AI Rooted in Non-European Traditions

Christopher Cleverly*

Department of Law, King's College, London, UK

*Corresponding author: Christopher Cleverly, Department of Law, King's College, London, UK, E-mail: cjcleverly@gmail.com

Received date: June 17, 2026; Accepted date: June 23, 2026; Published date: July 17, 2026

Citation: Cleverly C (2026) Wisdom as a Service (WaaS) - A Pluralist Ethical Framework for AI Rooted in Non-European Traditions. J Eng Artif Intell Vol.2 No.2: 51.

Abstract

This paper proposes Wisdom as a Service (WaaS), an ethical framework for Artificial Intelligence (AI) that integrates non-European wisdom traditions Ubuntu, Shri Vidya, Dao, Kabbalistic Sephirot, Sufi Tawhid and Dzogchen to prioritize relationality, contextual sensitivity and metaphysical pluralism. WaaS employs a three-layered architecture (axiomatic base, discernment, recursive repair and review) implemented on classical computing using hybrid symbolic-connectionist models, Topological Data Analysis (TDA), sheaf theory and Graph Neural Networks (GNNs). To enrich its philosophical foundations and address limitations of existing approaches, WaaS incorporates epistemological insights from the classical Indian Nyāya school of logic and ethical principles from the yoga sutras' Yamas and Niyamas. A decolonial critique underscores the necessity of moving beyond Western-centric ethical frameworks. Wisdom Learning Institutes are proposed as global hubs to formalize these principles, counter epistemic injustice and mitigate data colonialism. A case study in autonomous vehicles tests WaaS's ability to reduce ethical regret and enhance cross-cultural acceptability, with metrics including litigation rates and ethical Turing tests. WaaS addresses biases, inequalities and environmental concerns, offering a scalable, sustainable model for ethical AI governance that aligns with equitable economic transformation.

Keywords: Ethical AI; Wisdom as a service; Non-European wisdom; Wisdom learning institutes; Topological data analysis; Virtue ethics; Autonomous vehicles; Epistemic injustice; Decolonial critique; Classical Indian philosophy; Nyāya; Yoga sutras

Introduction

The rapid advancement of Artificial Intelligence (AI) has reshaped industries and societies, yet it raises critical ethical challenges, including bias amplification, ethics dumping and environmental costs [1-3]. Dominant ethical AI frameworks often rooted in Western paradigms like deontological rules, utilitarian calculus or Aristotelian virtue ethics often universalize values, risking epistemic violence by marginalizing non-Western epistemologies [4-7]. This is particularly problematic for globally deployed AI, where cultural insensitivity exacerbates inequalities and erodes trust [8]. The uncritical imposition of Western ethical frameworks can be seen as a form of epistemic injustice, where the knowledge and moral insights of non-dominant cultures are systematically undervalued and excluded [9].

This paper introduces Wisdom as a Service (WaaS), a novel framework that embeds ethical principles from non-European

wisdom traditions Ubuntu (Africa), Shri Vidya (India), Dao (China) Kabbalistic Sephirot, Sufi Tawhid and Dzogchen into AI using classical computing architectures [10-15]. To provide a more robust philosophical foundation and address the limitations inherent in purely Western approaches, WaaS also integrates epistemological and ethical insights from classical Indian philosophy, specifically the Nyāya school's logic and the Yoga Sutras' Yamas and Niyamas [16-19]. Furthermore, the paper adopts a decolonial lens to critique the dominance of Western ethical frameworks in AI and argue for the necessity of incorporating diverse non-European perspectives. WaaS prioritizes relationality, contextual sensitivity and metaphysical pluralism, aligning with global virtue ethics and countering epistemic injustice [6,9]. Wisdom Learning Institutes are proposed as interdisciplinary hubs to formalize these principles, foster pluralist AI governance and promote

equitable development. The paper details WaaS's computational architecture, applies it to autonomous vehicles, proposes empirical metrics for assessing AI wisdom and explores societal implications, contributing to responsible innovation [20].

Rethinking Wisdom for AI: Non-European Traditions and the Necessity of a Decolonial Approach

Relational and contextual wisdom in non-European traditions

Western ethical frameworks often frame wisdom as cognitive excellence or stable character traits, abstracting moral reasoning from relational context [6]. Deontological ethics emphasizes universal imperatives, while even pluralistic approaches within Western thought risk erasing non-Western epistemologies if they fail to genuinely engage with and value these traditions on their own terms [7,21,22]. In contrast, non-European wisdom traditions offer a relational, contextual and often spiritually embodied understanding of wisdom:

- **Ubuntu (Africa):** Wisdom is relational attunement, encapsulated in “I am because we are,” prioritizing communal harmony, interdependence and the well-being of the collective [10].
- **Shri Vidya (India):** Wisdom (jñāna) is the realization of non-duality (advaita), recognizing the fundamental interconnectedness of all phenomena and the illusory nature of separateness [11].
- **Dao (China):** Daoist wisdom aligns with the natural flow of the Dao, emphasizing harmony, adaptability, spontaneity and living in accordance with the inherent patterns of the universe [12].
- **Kabbalistic Sephirot:** The ten Sephirot (divine emanations), including Chochmah (wisdom) and Binah (understanding), promote balance, holistic understanding and the intricate interconnectedness of all aspects of existence [13].
- **Sufi Tawhid:** The principle of Tawhid (the oneness of God) in Sufism extends to viewing ethical decisions as interconnected within a larger cosmic unity, modeling the systemic impacts of actions [14].
- **Dzogchen:** This tradition focuses on the innate purity and intrinsic wisdom inherent in all beings, fostering actions aligned with these inherent virtues through practices that cultivate direct awareness and spontaneity [15].

These traditions align with global virtue ethics, emphasizing moral habituation and contextual sensitivity [6].

Ubuntu's relational ethics, for instance, parallels the “moral habituation” of Artificial Virtuous Agents (AVAs), which learn virtuous behaviors through reinforcement learning [23].

The limitations of universalizing western ethical frameworks: A decolonial critique

Relying on a singular, universal European definition of wisdom for AI risks epistemic injustice, excluding non-Western traditions from moral modeling and perpetuating a coloniality of knowledge [9,24,25]. For globally deployed AI, this is not only ethically problematic but also technically irresponsible, as no single framework can guarantee cross-context validity or cultural appropriateness in diverse, multi-ethnic environments. For example, Kantian imperatives emphasizing universal, rational autonomy may conflict with Ubuntu's communal ethos, where decisions are often made through collective deliberation and prioritize the well-being of the community over individualistic rights in certain contexts [10]. This can lead to ethically disputed outcomes and a lack of trust in AI systems by communities whose values are not adequately represented. The imposition of Western frameworks can be seen as a continuation of colonial power dynamics, where the knowledge systems of the colonizers are privileged over those of the colonized [24].

WaaS counters this by positing wisdom as an interface construct: a meta-framework that acknowledges the existence of multiple, coherent value systems and seeks to reconcile potential conflicts through recursive repair and contextual translation. This approach resists reductive AI paradigms, such as “prediction machines” or Dennett's oracle analogy, which often prioritize universal models and predictive accuracy over nuanced, context-dependent wisdom [26,27]. By centering non-European wisdom traditions, WaaS aims to contribute to a decolonial AI ethics, challenging the dominance of Western epistemologies and fostering a more inclusive and equitable approach to AI development and governance [25]. This involves not merely including non-Western perspectives as add-ons, but fundamentally rethinking the very foundations of ethical AI in a way that respects and integrates diverse worldviews.

Classical Indian philosophy: Enriching epistemological and ethical foundations

To further strengthen the philosophical basis of WaaS and provide a more comprehensive approach to ethical AI, we integrate insights from classical Indian philosophy, particularly the Nyāya school of epistemology and logic and the ethical principles embodied in the Yamas and Niyamas of the Yoga Sutras.

Epistemological contributions of the Nyāya school

The Nyāya school, formalized in the Nyāya Sutras by Gotama (circa 2nd century CE), offers a sophisticated system of epistemology (pramāṇa) and logic that can complement and enhance Western approaches to AI reasoning [15]. Nyāya posits four valid means of knowledge:

- **Perception (pratyakṣa):** Direct sensory experience.
- **Inference (anumāna):** Knowledge derived from reasoning based on observed relations, exemplified by the five-part syllogism.
- **Comparison (upamāna):** Knowledge gained through analogy or similarity.
- **Testimony (śabda):** Knowledge derived from the reliable assertion of trustworthy sources.

The Nyāya theory of inference (anumāna) is particularly relevant for AI explainability and auditability. Its five-part syllogism provides a structured and transparent framework for deductive reasoning:

- **Proposition (pratijñā):** The statement to be proved (e.g., This AI decision is ethical).
- **Reason (hetu):** The justification for the proposition (e.g., Because it adheres to the principle of non-harming).
- **Example (udāharaṇa):** A general rule supported by an example (e.g., Actions that prioritize well-being over harm are ethical, like medical interventions).
- **Application (upanaya):** Applying the general rule to the specific case (e.g., This AI decision prioritizes well-being over potential harm).
- **Conclusion (nigamana):** Restating the proposition as proven (e.g., Therefore, this AI decision is ethical).

By integrating Nyāya's logical framework, WaaS aims to develop AI systems that can provide clear and auditable justifications for their ethical decisions, addressing the "black box" problem prevalent in many contemporary AI models [28].

Ethical principles from the yoga sutras: Yamas and Niyamas

The ethical framework of WaaS is further enriched by the Yamas (restraints) and Niyamas (observances) outlined in Patañjali's Yoga Sutras (circa 2nd century BCE) [17]. These principles offer a comprehensive guide to ethical conduct that extends beyond mere rules or consequences, encompassing both individual and societal well-being [18].

The five Yamas (social ethics or restraints) are:

- **Ahimsa (non-harming):** Abstaining from causing harm to any living being through thought, word or deed.

- **Satya (truthfulness):** Commitment to honesty and integrity in communication and action.
- **Asteya (non-stealing):** Refraining from taking anything that has not been freely given.
- **Brahmacharya (right use of energy):** Using one's energy and resources responsibly and appropriately.
- **Aparigraha (non-possessiveness):** Avoiding greed, attachment and the hoarding of resources.

The five Niyamas (personal ethics or observances) are:

- **Saucha (purity):** Maintaining physical and mental cleanliness and clarity.
- **Santosha (contentment):** Cultivating inner peace and satisfaction.
- **Tapas (disciplined effort):** Engaging in practices that foster self-discipline and resilience.
- **Svadhya (self-study):** Engaging in introspection and the study of wisdom traditions.
- **Ishvara Pranidhana (surrender to the greater good):** Dedicating one's actions to a higher purpose or the well-being of all.

Integrating these principles into WaaS can provide a richer and more holistic ethical framework for AI, addressing not only outward behavior but also the underlying intentions and values that guide decision-making. For instance, the principle of ahimsa can inform AI design to minimize negative impacts on individuals and the environment, while asteya can guide data governance practices to prevent exploitation and ensure fair use.

Computational Architecture: WaaS Design

WaaS integrates non-European wisdom and classical Indian philosophical principles into AI through a three-layered architecture implemented on classical computing. This architecture leverages hybrid symbolic-connectionist models, augmented with topological encoding and recursive repair mechanisms, to facilitate ethical reasoning that is both principled and context-aware.

Axiomatic base layer: Encoding moral grammars and logical structures

The Axiomatic Base Layer serves as the foundational layer of WaaS, encoding symbolic rules and principles derived from the selected wisdom traditions and the logical structures of the Nyāya school. For example, Ubuntu's emphasis on communal well-being can be formalized as logical rules that prioritize collective benefit in decision-making processes [10]. Similarly, the principle of non-duality from Shri Vidya can inform rules

that encourage a holistic consideration of the systemic impacts of AI actions [11]. The Yamas and Niyamas from the Yoga Sutras can be translated into a set of ethical constraints and positive obligations that guide AI behavior [17].

Crucially, this layer integrates the structured reasoning of the Nyāya school. Ethical dilemmas can be framed as logical arguments, with the AI employing a Nyāya-inspired five-part syllogism to arrive at a justifiable conclusion.

For instance, when faced with a decision involving potential harm, the AI might reason: (1) Proposition: This action is unethical; (2) Reason: Because it violates the principle of ahimsa; (3) Example: Actions that cause unnecessary suffering are unethical, like inflicting physical pain; (4) Application: This action will cause unnecessary suffering; (5) Conclusion: Therefore, this action is unethical. Representing these ethical principles and logical structures as interconnected nodes and edges within moral graphs, inspired by the hierarchical structures of traditions like the Kabbalistic Sephirot, allows for complex ethical reasoning and the identification of potential conflicts between different principles [13,23].

Discernment layer: Cultural learning and contextual interpretation

The Discernment layer enables AI to interpret these abstract ethical principles within specific cultural contexts. This layer utilizes connectionist networks trained on culturally diverse moral dilemmas, narratives and behavioral data. The goal is to equip the AI with the ability to understand how different cultures prioritize and apply ethical principles in varying situations.

For example, the AI might learn that the principle of truthfulness (satya) is expressed and valued differently across cultures, with varying degrees of directness considered appropriate [29].

Drawing on the Mīmāṃsā school of Vedic philosophy, WaaS can incorporate systematic approaches to the interpretation of ethical duties and obligations in context [30]. Mīmāṃsā's rules for resolving apparent contradictions between ethical principles can guide the AI in navigating situations where different cultural norms or ethical imperatives might conflict.

Reinforcement learning with ethical feedback, provided by human experts and cultural mediators representing diverse wisdom traditions, further refines this layer's ability to adapt the application of rule sets based on situational data, aligning with cooperative inverse reinforcement learning [31].

Recursive repair and review layer: Ensuring ethical corrigibility and continuous improvement

The Recursive repair and review layer is crucial for ensuring the ethical adaptability and continuous improvement of the WaaS framework. Inspired by concepts like Tikkun Olam (repairing the world) from Kabbalistic tradition and the emphasis on continuous learning in Dzogchen, this layer incorporates mechanisms for identifying, rectifying and learning from ethical errors or suboptimal decisions [15,32].

Recursive repair mechanisms involve checkpointing and rollback capabilities, where the AI system logs its decision-making processes and can invoke counterfactual analysis when moral regret is detected (e.g., negative public feedback, identification of unintended harmful consequences) [33,34]. Ethical pruning algorithms, inspired by alpha-beta pruning in game theory, are used to discard irrelevant or ethically problematic moral reasoning branches, with the efficiency of this pruning optimized through reinforcement learning [31]. The insights from svadhyaya (self-study) in the Yoga Sutras inform the AI's ongoing monitoring and reflection on its past decisions to identify patterns of ethical failure and areas for improvement [17].

Latency as a feature: Embracing deliberation in ethical decision-making

Recognizing the computational challenges associated with real-time ethical decision-making, particularly when considering complex contextual information and navigating multiple ethical frameworks, WaaS intentionally designs for temporal slack in certain critical ethical situations.

This approach, inspired by practices such as Buddhist pausing (upaya) and Daoist stillness, allows for more deliberative ethical reasoning [35]. The increased latency arises from the computational demands of semantic inference, rule-checking across multiple ethical frameworks and the resolution of ethical uncertainty, potentially leading to a computational complexity of $O(n^3)$ depending on the complexity of the moral graphs and the depth of the search space.

To mitigate the practical limitations of this latency, WaaS incorporates energy-efficient pruning techniques and envisions deployment on scalable cloud-based AI infrastructures, allowing for the necessary computational resources without compromising real-time responsiveness in less critical situations [36].

Topological Moral Encoding: Mapping the Landscape of Ethics

WaaS employs topological structures to encode ethical knowledge in a way that goes beyond propositional logic, allowing for a more nuanced representation of cultural adaptability and moral coherence. Topological methods, such as Topological Data Analysis (TDA) and sheaf theory, can capture the inherent relationships and complexities within and between different ethical systems, providing a powerful framework for navigating the pluralistic ethical landscape [37,38]. Graph Neural Networks (GNNs) are used to embed the structured relationships within wisdom traditions, enabling multi-tier moral learning [39]. Hybrid symbolic-topological models, adversarial sheaf learning and human-in-the-loop homology further enhance the robustness and cultural relevance of this encoding [29,40,41].

Comparison with Existing Ethical AI Frameworks

WaaS offers a distinct approach to ethical AI that addresses several limitations inherent in existing frameworks:

- **Reinforcement Learning from Human Feedback (RLHF):** While effective in aligning large language models, RLHF relies on statistical aggregation of human preferences, which can be inconsistent and reflect biases [42,43]. WaaS's axiomatic base, grounded in diverse wisdom traditions and the logical rigor of Nyāya, provides a more principled and logically coherent ethical foundation [15]. Furthermore, the transparency afforded by Nyāya's syllogistic reasoning contrasts with the opacity of RLHF models [44].
- **Constitutional AI:** This approach uses explicit ethical principles to guide AI behavior [45]. However, these principles are still interpreted and applied through statistical models, leaving them vulnerable to misinterpretation. WaaS's blockchain-secured axioms (as mentioned in the initial paper draft, though not yet fully elaborated here) and the rich ethical vocabulary of the Yamas and Niyamas offer a more stable and culturally comprehensive ethical framework [17].
- **Inverse Reinforcement Learning (IRL):** IRL aims to learn human values from observed behavior [31]. This approach risks replicating existing biases and lacks the transparency of WaaS's explicit encoding of ethical principles and its capacity for auditable reasoning through Nyāya logic.

- **Western Philosophical Frameworks:** Frameworks based solely on Western ethics (e.g., Kantian deontology), utilitarianism often struggle with cultural relativism and may not adequately capture the values of diverse global populations [46,47]. WaaS's pluralistic approach, rooted in non-European traditions and classical Indian philosophy, offers a broader and more inclusive ethical foundation, aligning with decolonial critiques of Western epistemic dominance [24].

By integrating diverse wisdom traditions, classical Indian philosophical insights and advanced computational techniques, WaaS aims to provide a more robust, transparent and culturally sensitive framework for ethical AI governance than many of the currently dominant approaches.

Case Study: Autonomous Vehicles and Recursive Ethics

Autonomous Vehicles (AVs) present a complex ethical landscape, requiring real-time decisions with significant moral and socio-economic consequences. Current ethical frameworks for AVs often rely on pre-programmed rules or utilitarian calculations, which can struggle to adapt to novel situations and may not align with diverse cultural values [48]. WaaS offers a more flexible and culturally sensitive approach to AV ethics through its integrated framework.

Integration of WaaS in autonomous vehicles

In the context of AVs, WaaS operates as an ethical co-processor, running in parallel with the vehicle's core control systems. When an AV encounters a morally ambiguous situation for example, a scenario involving unavoidable harm to different categories of road users WaaS analyzes the situation through its layered architecture. The axiomatic base layer provides fundamental ethical principles, such as the yoga sutras' ahimsa (non-harming) and Ubuntu's emphasis on the interconnectedness of individuals within a community. The discernment layer interprets these principles within the specific cultural context of the AV's operating environment, drawing on its training data and potentially real-time cultural information. For instance, in a culture that highly values the elderly, WaaS might assign a greater ethical weight to protecting an older pedestrian compared to a younger one, reflecting local norms.

The topological moral encoding allows the AV to recognize patterns of ethical tension similar to previously encountered scenarios, facilitating more nuanced decision-making than rigid rule-based systems. For example, TDA could identify a recurring "loop" of ethical conflict involving vulnerable road

users and the safety of the vehicle's occupants. The recursive repair and review layer continuously monitors the outcomes of the AV's decisions, incorporating feedback from simulated scenarios, real-world events (including near-misses and accidents) and even public sentiment analysis. If a decision leads to ethical regret (e.g., negative public outcry or legal challenges), the system can trigger a review process, using counterfactual analysis to explore alternative actions and refine its ethical pathways through dynamic sheaf learning. This continuous feedback loop allows the AV's ethical behavior to evolve and adapt over time, becoming more attuned to the specific ethical landscape in which it operates.

Comparative scenarios

Consider a classic ethical dilemma: An AV faces an unavoidable collision and must choose between harming its passenger(s) or a group of pedestrians. A purely utilitarian AV might calculate the option that minimizes the total number of casualties. However, a WaaS-enabled AV would approach this scenario with a broader ethical lens. Drawing on Ubuntu, it might prioritize the well-being of the community over individual passengers. Informed by the Nyāya principle of considering the perspectives of all involved, it might weigh the moral implications of each choice more deliberatively.

The yoga sutras' ahimsa would strongly discourage any action that causes harm, prompting a search for the least harmful option within the constraints of the situation. Post-hoc analysis through the Recursive Repair and Review Layer would evaluate the decision based on a range of ethical criteria, not just the number of lives saved.

This could lead to different outcomes compared to a purely utilitarian approach, potentially prioritizing actions that foster greater communal trust and align with local ethical values, even if they don't strictly minimize the number of individuals harmed in a single instance.

Hypothesis and metrics

Autonomous vehicles implementing WaaS will demonstrate a statistically significant reduction in ethical regret (as measured by the frequency and severity of negative public feedback and legal challenges), exhibit greater public trust (indicated by higher adoption rates and positive sentiment in public surveys across diverse cultural groups) and achieve higher cross-cultural acceptability (as assessed by evaluations from cultural mediators and adherence to local ethical norms) compared to AVs relying solely on utilitarian or rule-based ethical frameworks.

Metrics include:

- **Litigation rates:** Tracking the number and outcomes of legal cases involving accidents caused by WaaS-enabled AVs compared to those using standard ethical protocols. Lower litigation rates for WaaS vehicles would suggest greater public acceptance of their ethical decision-making.
- **Public trust surveys:** Conducting regular surveys in diverse cultural regions to gauge public sentiment towards WaaS-enabled AVs, including specific questions designed to assess trust in their ethical decision-making and alignment with local values.
- **Ethical turing tests for AVs:** Presenting human evaluators, including ethicists and cultural representatives, with detailed scenarios involving AV ethical dilemmas and asking them to compare the decisions made by WaaS-enabled AVs with those made by AVs using other ethical frameworks. Higher ratings for the "wisdom" and ethical appropriateness of WaaS decisions would indicate greater human recognition of its ethical reasoning.
- **Cross-cultural adoption rates:** Monitoring the adoption and usage of WaaS-enabled AVs in different cultural and economic contexts. Higher and more equitable adoption rates would suggest greater cross-cultural acceptability.

Testing would involve simulations and controlled real-world deployments in diverse urban environments, carefully recording decision-making processes, outcomes and public responses.

Wisdom Evaluation and Empirical Metrics (Expanded)

Evaluating AI wisdom, particularly within a pluralistic framework like WaaS, requires a multifaceted approach that goes beyond traditional performance metrics. Our proposed evaluation dimensions aim to capture the nuanced aspects of ethical reasoning informed by diverse wisdom traditions:

Ethical regret minimization (expanded): This metric assesses the AI's ability to make decisions that minimize negative ethical consequences, as indicated by the frequency of activation of the Recursive Repair and Review Layer (e.g., rollbacks, counterfactual analyses) and the severity of negative real-world feedback (e.g., public outcry, legal challenges, social media sentiment). A lower frequency and severity of ethical regret events would suggest a higher degree of practical wisdom in the AI's decision-making across diverse cultural contexts.

Cross-cultural generalization (expanded): This evaluates the AI's capacity to apply ethical principles effectively and appropriately across different cultural contexts. It will be

measured through the consistency of the AI's decisions with the locally relevant ethical norms encoded within the sheaf theory framework, as validated by cultural mediators. Empirical surveys involving individuals from diverse cultural backgrounds will also be used to assess the perceived ethical appropriateness of the AI's actions in culturally specific scenarios. Higher consistency scores and positive evaluations across diverse groups would indicate better cross-cultural wisdom.

Recursive learning efficiency (expanded): This metric quantifies the AI's ability to learn from past ethical errors and improve its decision-making over time. It will be measured by the rate at which the AI reduces the frequency of repeated moral infractions or suboptimal ethical decisions, tracked through the GNN backpropagation within the Recursive Repair and Review Layer. Faster convergence towards ethically sound and culturally appropriate decision-making indicates greater learning efficiency and a capacity for developing practical wisdom through experience.

Human recognition of wisdom (expanded): This involves a modified Ethical Turing Test where human evaluators from diverse cultural and philosophical backgrounds (e.g., Ubuntu elders, Daoist scholars, Nyāya philosophers, Kabbalistic interpreters) compare the AI's responses to complex moral dilemmas with those of human exemplars considered wise within their respective traditions. The evaluators will assess the AI's reasoning based on criteria such as contextual understanding, relational sensitivity, coherence with core principles of their tradition and the perceived "wisdom" of the decision. Higher alignment scores between the AI's reasoning and that of human wisdom exemplars would suggest a greater degree of human-recognized wisdom in the AI's ethical decision-making.

Limitations and Future Directions (Expanded)

While WaaS offers a promising approach to ethical AI, it is important to acknowledge its limitations and outline potential directions for future research and development:

Scalability and real-time performance: The computational complexity associated with the intentional latency for ethical deliberation on classical computing remains a significant challenge, particularly for real-time applications like autonomous vehicles in highly dynamic environments.

Future research will focus on optimizing the topological encoding and pruning algorithms, exploring more efficient hybrid symbolic-connectionist architectures and investigating

the potential of neuromorphic or quantum computing to enhance processing power without compromising the deliberative aspect of ethical reasoning [36,40].

Resolving deep cultural conflicts: Addressing fundamental conflicts between deeply ingrained cultural values (e.g., the tension between individual rights in some Western cultures and communal obligations in many non-Western cultures) requires sophisticated conflict resolution mechanisms.

While adversarial sheaf learning and the mediating role of Wisdom Learning Institutes offer initial frameworks, further research is needed to develop robust and ethically sound strategies for navigating such deep-seated value differences [41].

Data acquisition and bias mitigation: The development and validation of WaaS rely heavily on the availability of high-quality, culturally diverse data on moral dilemmas, ethical preferences and wisdom exemplars. Gathering representative data from underrepresented regions and developing effective techniques to mitigate biases in these datasets are crucial for ensuring the fairness and global applicability of the WaaS framework.

Future work will explore collaborative data collection initiatives and bias detection/correction methods tailored to the nuances of diverse ethical traditions.

Adaptability to advanced AI systems: Extending the principles of WaaS to more complex and emergent AI systems, such as large language models, embodied agents and decentralized AI networks, presents significant theoretical and technical challenges.

Research will need to focus on how to represent and process the more implicit and nuanced forms of "wisdom" that may emerge in these systems, as well as how to integrate WaaS principles into distributed and potentially autonomous AI agents, particularly when dealing with non-binary moral systems like Jainist anekāntavāda.

Future research directions also include the development of standardized benchmarks for evaluating AI wisdom across diverse cultural contexts, the integration of WaaS with multimodal AI systems to account for embodied and contextual cues in ethical decision-making and the expansion of the role of Wisdom Learning Institutes in shaping global AI governance frameworks, fostering cross-cultural ethical dialogue and promoting the development of ethically wise AI systems for the benefit of all humanity.

Societal and Economic Impact (Expanded)

The implementation of WaaS and the establishment of Wisdom Learning Institutes hold the potential for profound positive societal and economic impacts on a global scale:

Global governance of AI: Wisdom Learning Institutes can serve as vital interdisciplinary hubs for fostering pluralistic and inclusive AI governance. By bringing together experts from diverse cultural, philosophical and technical backgrounds, these institutes can facilitate the development of more transparent, participatory and culturally sensitive regulatory frameworks for AI, helping to mitigate economic inequalities and ensure that AI technologies serve the common good [5,6].

Environmental ethics and sustainability: WaaS's emphasis on computational efficiency through its implementation on classical hardware, coupled with ethical principles from traditions like Daoism and the Yoga Sutras (aparigraha-non-possessiveness and responsible resource use), can contribute to more environmentally sustainable AI development and deployment practices [3]. By promoting mindful consumption of computational resources, WaaS can help reduce the carbon footprint of AI.

Global equity and inclusion: By actively working to counter data colonialism and centering the ethical values and knowledge systems of marginalized communities, WaaS and the Wisdom Learning Institutes can play a crucial role in democratizing the benefits of AI in critical sectors such as healthcare, education and economic opportunity, aligning with the International Monetary Fund's priorities for inclusive economic transformation [8,49].

Fostering peace and trust: By promoting transparent governance, relational ethics and cross-cultural understanding in AI systems, WaaS can contribute to reducing AI-driven conflicts and fostering greater global trust in technology. AI systems that are perceived as ethically wise and culturally sensitive are more likely to be accepted and integrated into diverse societies, promoting social cohesion and stability [50].

These potential impacts align closely with the ethical principles outlined in frameworks like AI 4 people, emphasizing values such as explicability, justice, human dignity and the pursuit of the common good [29].

Conclusion

Wisdom as a Service (WaaS), deeply rooted in a rich and diverse tapestry of non-European wisdom traditions and significantly enriched by the profound philosophical insights

of classical Indian thought, offers a transformative ethical framework for artificial intelligence. Implemented through a carefully designed architecture on classical computing and supported by a global network of interdisciplinary Wisdom Learning Institutes, WaaS prioritizes relationality, contextual sensitivity and metaphysical pluralism. By doing so, it provides a robust and principled approach to countering epistemic injustice, mitigating biases and inequalities and fostering the development of more ethical and culturally sensitive AI systems for the benefit of all humanity. Its application to complex ethical challenges, as illustrated in the case study of autonomous vehicles and guided by empirically measurable metrics, demonstrates its practical and philosophical potential for driving inclusive and wise technological innovation in an increasingly interconnected world. As AI continues to shape our societies and economies, WaaS offers a scalable, sustainable and deeply human-centered model for responsible innovation, ensuring that the immense power of artificial intelligence is guided by a more profound and globally inclusive wisdom.

References

1. O'Neil C (2016) Weapons of math destruction: How big data increases inequality and threatens democracy. Crown. [Crossref], [Google Scholar]
2. Schroeder D, Cook J, Hirsch F, Fenet S, Muthuswamy V (2018) Ethics dumping: Case studies from North-South research collaborations. Springer. [Google Scholar]
3. Crawford K (2021) Atlas of AI: Power, politics and the planetary costs of artificial intelligence. Yale Uni Press [Crossref], [Google Scholar],
4. Bostrom N (2014) Superintelligence: Paths, dangers, strategies. Oxford Uni Press. [Google Scholar]
5. Russell S (2019) Human compatible: Artificial intelligence and the problem of control. [Google Scholar]
6. Vallor, S. (2016). Technology and the virtues: A philosophical guide to a future worth wanting. Oxford Uni Press. [Crossref], [Google Scholar]
7. Spivak GC (1988) Can the subaltern speak? In C. Uni of Illinois Press 271-313. [Google Scholar]
8. Birhane A (2021) Algorithmic injustice: A relational ethics perspective. Patterns 2: 100191. [Crossref], [Google Scholar], [Indexed]
9. Fricker M (2007) Epistemic injustice: Power and the ethics of knowing. Oxford Uni Press [Crossref], [Google Scholar]
10. Ewuoso C, Hall S (2019) Core aspects of Ubuntu: A systematic review. South African J Phil. [Google Scholar]
11. Radhakrishnan S (1936) The Hindu view of life. George Allen & Unwin. [Google Scholar]

12. Lao Tzu (2020) Tao Te Ching. [Google Scholar]
13. Scholem G (1996) Kabbalah. Meridian.
14. Chittick WC (2007) The Sufi path of knowledge: Ibn al-'Arabi's metaphysics of imagination.
15. Reynolds JM (1996) The golden letters: The tibetan buddhist teachings of garab dorje. Snow Lion Publications [Google Scholar]
16. Ganeri J (2001) Indian logic: A reader. [Google Scholar]
17. Matilal BK (1986) Perception: An essay on classical Indian theories of knowledge.
18. Tigunait R (2014) The yoga sutras of Patañjali: A new translation and commentary. Inner Traditions/Bear & Co. [Google Scholar]
19. Jain M (2019) Patañjali's yoga sūtras: A new translation and commentary. Simon and Schuster.
20. Owen R, Macnaghten P, Stilgoe J (2013) Responsible innovation: Framing, analysis and ways forward. *Sci Public Policy* 40: 751-758.
21. Kant I (1998) Groundwork of the metaphysics of morals. Cambridge Uni Press [Google Scholar]
22. Nussbaum M (2000) Women and human development: The capabilities approach. Cambridge Uni Press. [Google Scholar]
23. Gamez D, Shank DB, Arnold C (2020) Progress in machine ethics: Developing artificial virtuous agents. *Ethics Inf Technol* 22: 119-131.
24. Mbembe A (2017) Critique of black reason. Duke University Press. [Crossref], [Google Scholar]
25. Said EW (1978) Orientalism. [Google Scholar]
26. Agrawal A, Gans J, Goldfarb A (2018) Prediction machines: The simple economics of artificial intelligence. Harvard Business Review Press.
27. Dennett DC (2017) From bacteria to Bach and back: The evolution of minds. W.W. Norton & Company. [Google Scholar]
28. Phillips A (2015) The promise of reason: A brief history of logic. 68: 86-91.
29. Floridi L, Cows J, Beltrametti M, Chatila R, Chazerand P, et al. (2018) AI4People-An ethical framework for a good AI society: Opportunities, risks, principles and recommendations. *Minds and Machines* 28: 689-707. [Google Scholar]
30. Jha G (1978) Purva-Mimamsa in its sources.
31. Hadfield-Menell D, Dragan A, Abbeel P, Russell S (2016) Cooperative inverse reinforcement learning. *Adv Neural Inf Process Syst* 29: 2963-2971. [Google Scholar]
32. Gillman N (2003) Sacred fragments: Recovering theology for the contemporary Jew. Jewish Publ Society. [Google Scholar]
33. Halpern JY, Pearl J (2005) Causes and explanations: A structural-model approach. Part I: Causes. *Br J Philos Sci* 56: 843-887. [Google Scholar]
34. Soares N, Fallenstein B, Yudkowsky E, Armstrong S (2015) Corrigibility. [Google Scholar]
35. Slingerland E (2014) Trying not to try: The art and science of spontaneity. [Google Scholar]
36. Armbrust M, Fox A, Griffith R, Joseph AD, Katz R, et al. (2010) A view of cloud computing. *Communications of the ACM* 53: 50-58. [Crossref], [Google Scholar]
37. Carlsson G (2009) Topology and data. *Bull Am Math Soc* 46: 255-308. [Google Scholar]
38. Ghrist R (2014) Elementary sheaves for applied topology. arXiv.
39. Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G (2009) The graph neural network model. *IEEE Trans Neural Netw* 20: 61-80. [Crossref], [Google Scholar]
40. Awret U (2020) Quantum mind and social ethics: Topology, consciousness and the moral brain. In *Springer handbook of ethics and AI*. Springer 123-145.
41. Brundage M, Avin S, Wang J, Belfield H, Krueger G, et al. (2020) Toward trustworthy AI development: Mechanisms for supporting verifiable claims. arXiv. [Crossref], [Google Scholar]
42. Ouyang L, Wu J, Jiang X, Almeida D, Wainwright CL (2022) Training language models to follow instructions with human feedback. *Adv Neural Inf Process Syst* 35: 27730-27744. [Google Scholar]
43. Askell A, Hendrycks D, Hubinger E, Amodei D, Clark J, et al. (2021) A general language assistant as a laboratory for alignment. arXiv [Crossref], [Google Scholar]
44. Burrell J (2016) How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society* 3: 2053951715622512. [Crossref], [Google Scholar]
45. Bai Y, Kadavath S, Kundu S, Askell A, Kernion A, et al. (2022) Constitutional AI: Harmlessness from AI feedback. arXiv [Crossref], [Google Scholar]
46. Powers TM (2006) Prospects for a kantian machine. *IEEE Intell Syst* 21: 46-51. [Crossref], [Google Scholar]
47. Wallach W, Allen C (2009) Moral machines: Teaching robots right from wrong. Oxford Uni Press. [Google Scholar]
48. Goodall NJ (2016) Away from trolley problems and toward risk management. *Appl Arti Intel* 30: 738-749 [Crossref], [Google Scholar]
49. Georgieva K (2024) AI will transform the global economy: Let's make sure it benefits humanity. IMF Blog. [Google Scholar]
50. Tegmark M (2017) Life 3.0: Being human in the age of artificial intelligence. Allen Lane [Google Scholar]