

International journal of interdisciplinary and multidisciplinary research

ISSN 2456-4567 (O)

Comparative Ethical Perspectives in Major World Religions: A Philosophical and Computational Framework

Pankoj Kumar Dutta

Department of Philosophy
State Aided College Teacher
Sitalkuchi College
Coochbehar, West Bengal, India

Abstract:

The philosophical study of comparative ethics across major world religions offers profound insights into fundamental human values and complex moral reasoning. While there is universal agreement across societies that agents ought to act ethically, there remains profound disagreement as to what precisely constitutes ethical behavior within varying theological frameworks. This paper presents a structured philosophical and computational framework for systematically analyzing the distinct ethical perspectives of Hinduism, Jainism, Buddhism, Christianity, and Islam. By rigorously mapping these ancient traditions into structured representations, we aim to bridge the persistent gap between historical theological doctrines and modern models of moral uncertainty. Ultimately, this approach provides a robust foundation for encoding pluralistic ethical viewpoints into artificial intelligence systems, ensuring culturally sensitive and globally representative moral reasoning.

Introduction

The comparative study of religious ethics provides a necessary window into the diverse moral frameworks that have guided human civilizations for millennia. Hinduism, Jainism, and Buddhism—often categorized under the Dharmic traditions—emphasize concepts such as Karma, Dharma, and Ahimsa

(non-violence), which inherently dictate a holistic and interconnected view of ethical obligations toward all sentient life. Conversely, the Abrahamic traditions of Christianity and Islam place a strong emphasis on divine command theory, justice, and agape (unconditional love) or adl (justice) as revealed through sacred scriptures. Understanding these distinct paradigms is crucial in our highly interconnected world, where moral pluralism continuously shapes global discourse and societal norms.

Our core problem definition revolves around the challenge of formally defining and integrating these disparate religious ethical perspectives into coherent, actionable frameworks for modern socio-technical systems. The scope of this paper encompasses the foundational ethical tenets of five major world religions—Hinduism, Jainism, Buddhism, Christianity, and Islam—and seeks to translate their philosophical principles into structured methodologies. As religious adherence acts as a crucial degree of freedom for human agents within global populations, capturing the nuances of these traditions is essential for accurately modeling societal dynamics and moral decision-making (Ausloos & Petroni, 2008).

However, existing approaches to modeling comparative religious ethics and integrating them into modern frameworks are fundamentally insufficient for several reasons. First, contemporary philosophical and computational approaches to moral uncertainty predominantly rely on Western-centric, secular ethical theories, such as utilitarianism or Kantian deontology, largely ignoring the rich doctrinal nuances of non-Western and religious traditions (Szabo et al., 2023). Second, quantitative studies on religion often treat religious groups merely as statistical entities adhering to distribution laws, without engaging with the profound internal ethical logic that drives the behavior of these populations (Mir, 2011).

To address these critical gaps, this paper introduces a novel intersection of comparative religious philosophy and formal modeling. Our primary paper contributions are as follows:

- We propose a comprehensive philosophical taxonomy that categorizes the core ethical principles of five major world religions, allowing for direct comparative analysis.
- We introduce a structured, hypothetical computational approach designed to safely encode these diverse religious ethical personas into artificial intelligence systems without succumbing to moral fanaticism.

Related Work

Moral Uncertainty and Ethical Aggregation

The philosophical problem of determining how an agent should act when uncertain about which ethical theory is correct has gained significant traction. Foundational proposals suggest aggregating multiple ethical theories to guide behavior, often through mechanisms like Maximising Expected Choiceworthiness (MEC) (Szabo et al., 2023). The core idea is that acknowledging different ethical perspectives requires mathematical aggregation to yield an overall moral evaluation. The strength of this approach lies in its democratic motivation to accommodate pluralistic views; however, its primary weakness is a vulnerability to fanaticism, where an extreme ethical theory might dominate an agent's behavior despite possessing low credence (Szabo et al., 2023). Compared to these theoretical welfare aggregations, our work grounds moral uncertainty directly within the specific theological doctrines of world religions, providing a concrete domain where aggregating diverse religious perspectives is practically necessary.

Encoding Personas and Ethical Views in Artificial Intelligence

A distinct body of literature focuses on how specific human characteristics, values, beliefs, and ethical views are represented within the architecture of large language models (LLMs). Research indicates that distinct personas, including ethical perspectives like moral nihilism or political ideologies, can be localized within the representation space of decoder-only LLMs (Cintas et al., 2025). Furthermore, custom generative agents are actively being developed to present alternate ethical perspectives to software developers, aiding in the creation of ethically compliant solutions (Olson, 2024). While these technological advancements show strength in their practical application for encoding human traits, they often lack a rigorous philosophical foundation regarding traditional religious ethics. Our framework complements these engineering efforts by supplying the necessary philosophical depth required to accurately represent complex Dharmic and Abrahamic moral paradigms.

Quantitative Evolution and Distribution of Religions

A third category of research focuses on the demographic and statistical evolution of world religions over time. Sociophysics approaches analyze the distribution and life time of religions using interacting agent modeling, applying differential equations to the growth and decay of religious adherence (Ausloos & Petroni, 2008). Similarly, empirical investigations have shown that the adherent distribution of major religions—such as Islam, Buddhism, Hinduism, and Judaism—often conforms to mathematical

phenomena like Benford's law of leading digits (Mir, 2011). The strength of these macroscopic approaches is their ability to reveal universal patterns in human belief systems using large-scale empirical data. Nevertheless, a major weakness is that they abstract away the actual ethical content of the religions, treating human agents as identical statistical particles. Our work bridges this gap by asserting that the underlying ethical teachings of these religions actively shape the behavioral dynamics and socio-technical interactions of their adherents.

Method/Approach

Philosophical Mapping and Taxonomy

The first component of our structured framework involves a rigorous philosophical mapping of the ethical tenets inherent in Hinduism, Jainism, Buddhism, Christianity, and Islam. This qualitative module extracts the paramount moral imperatives from each tradition: the strict adherence to Ahimsa (non-injury) in Jainism, the pursuit of Dharma (duty/righteousness) in Hinduism, the Eightfold Path of Buddhism, the ethics of Agape (selfless love) in Christianity, and the pursuit of Adl (divine justice) in Islam. By categorizing these tenets into a shared philosophical taxonomy, we establish a standardized matrix of moral values. This matrix allows for cross-religious comparison, identifying both shared embedding spaces (such as a universal prohibition against unprovoked harm) and distinct doctrinal regions.

Structured Computational Framework

To operationalize this philosophical taxonomy, we propose a numbered pipeline that simulates how a diverse array of religious ethics can be aggregated and evaluated.

1. **Theological Extraction:** Parse authoritative religious texts and scholarly commentaries to define weighted ethical constraints (e.g., placing absolute maximal weight on non-violence for the Jain persona).
2. **Moral Uncertainty Aggregation:** Apply non-fanatical aggregation algorithms—such as the Highest k-trimmed Mean—to combine these different religious perspectives into a consensus recommendation for a given moral dilemma (Szabo et al., 2023).
3. **Persona Encoding:** Map these aggregated constraints into the representational space of a hypothetical neural framework, explicitly ensuring that distinct religious ideologies are situated in distinct dimensional regions (Cintas et al., 2025).

Key Design Choices and Rationale

The core rationale behind utilizing non-fanatical aggregation methods in step two is specifically to prevent ethical fanaticism, which can occur when navigating highly divergent religious doctrines (Szabo et al., 2023). If standard expected choiceworthiness were used, the absolute strictness of Jain Ahimsa regarding microscopic life might computationally overpower all other practical ethical considerations in a generalized model. By applying trimmed means or median-based social welfare functionals, the framework respects the strictness of individual religious laws while ensuring a balanced, democratic aggregation of global moral perspectives. Furthermore, the decision to separate the theological extraction from the computational encoding ensures that philosophical integrity is maintained before any mathematical reduction occurs.

Evaluation Plan

To evaluate the efficacy of this framework, we propose a hypothetical benchmark dataset consisting of 1,000 complex moral dilemmas, ranging from environmental sustainability challenges (Holzapfel et al., 2022) to biomedical privacy conflicts (Azencott, 2018). In this hypothetical setup, simulated agents imbued with the five distinct religious personas will evaluate each dilemma. The evaluation metrics will measure the degree of alignment between the agents' outputs and the established orthodox theological rulings for each respective religion. Additionally, we will analyze the internal representation layers of the simulated system to verify that the distinct religious ethical perspectives are properly isolated, avoiding the polysemy and overlapping activations that could lead to theological confusion (Cintas et al., 2025).

Discussion

Practical Implications and Deployment

Deploying a robust comparative framework for religious ethics has significant practical implications for the design of culturally aware socio-technical systems. For instance, developers can utilize this taxonomy to build custom AI tools that assist engineers in incorporating diverse moral positions during the software design phase, moving beyond purely business-driven motives (Olson, 2024). Additionally, understanding these religious ethical lenses is critical for managing Crowdsourced Context Systems (CCS) on social media platforms (Lloyd et al., 2025). When digital communities collaboratively curate context around culturally sensitive topics, recognizing the underlying

religious moral frameworks can help platform designers structure presentation and platform treatment in ways that minimize sectarian conflict.

Limitations and Failure Modes

Despite its philosophical and computational rigor, this approach faces several inherent limitations and failure modes.

4. **Theological Reductionism:** Attempting to distill vast, multi-millennial religious traditions into a discrete set of parameterized ethical constraints risks severely oversimplifying complex and often internally debated doctrines.
5. **Representational Polysemy:** When encoding these complex philosophical traits into vector spaces, overlapping activations may occur between fundamentally different religious concepts that use similar linguistic expressions, leading to inaccurate moral representations (Cintas et al., 2025).
6. **Demographic Data Inaccuracies:** Determining the exact weighting of a religious perspective relies on historical adherence distributions, but tracking the precise evolution and original time of apparition of religious groups is notoriously difficult and often hampered by unreliable demographic data (Ausloos & Petroni, 2008).

Ethical Considerations and Risks

Integrating religious ethics into computational models introduces profound ethical risks that must be carefully managed.

7. **Deception and Exploitative Framing:** Anthropomorphizing an AI system to express specific religious ethical views can create a false sense of spiritual authority, potentially deceiving users or fostering an exploitative relationship framing (Ferrario et al., 2026).
8. **Illusion of Ethical AI:** Utilizing empirical interpretability tools to explain AI decisions based on religious personas might paradoxically obstruct genuine ethical discourse, providing a superficial illusion of ethical compliance while actually demonstrating a low denunciatory power against unethical outputs (John-Mathews, 2021).

Future Work

Future research must expand the boundaries of this comparative framework to encompass a wider array of global belief systems.

9. **Inclusion of Indigenous Philosophies:** Subsequent studies should formalize the ethical perspectives of indigenous and localized belief systems, moving beyond the major world religions to capture highly distinct ecological and communal moral codes.
10. **Modeling Religious Syncretism:** Future technical work must develop dynamic agent-based models capable of simulating religious syncretism, analyzing how individual human agents reconcile competing religious ethics when they are exposed to multiple overlapping theological domains over their lifetimes.

Conclusion

The intersection of comparative religious ethics and formal computational modeling represents a vital frontier for understanding global moral pluralism. By systematically examining the ethical tenets of Hinduism, Jainism, Buddhism, Christianity, and Islam, we have established a foundational taxonomy that honors the profound theological distinctions between these ancient traditions. We have proposed a hypothetical framework that translates these qualitative moral doctrines into structured modules, effectively mitigating the risks of moral fanaticism while accommodating complex moral uncertainty.

As technology increasingly mediates human interactions and social organizations, acknowledging the underlying religious ideologies that govern human behavior is no longer optional. Moving forward, the integration of comparative religious philosophy into AI development and statistical sociophysics will be essential. This interdisciplinary approach not only preserves the rich heritage of global moral reasoning but also ensures that the socio-technical systems of the future remain ethically robust, deeply interpretable, and authentically aligned with the diverse populations they serve.

References

1. Ausloos, M., & Petroni, F. (2008). On World Religion Adherence Distribution Evolution. pp. 72-78
2. Szabo, Jazon, Such, Jose, Criado, Natalia, & Modgil, Sanjay (2023). Moral Uncertainty and the Problem of Fanaticism.

3. Mir, Tariq Ahmad (2011). Law of the leading digits and the ideological struggle for numbers. *Physica A: Statistical Mechanics and its Applications* 391 (2012) pp. 792-798.
4. Cintas, Celia, Rateike, Miriam, Miehl, Erik, Daly, Elizabeth, & Speakman, Skyler (2025). Localizing Persona Representations in LLMs.
5. Holzapfel, André, Jäskeläinen, Petra, & Kaila, Anna-Kaisa (2022). Environmental and Social Sustainability of Creative-Ai.
6. Azencott, Chloé-Agathe (2018). Machine learning and genomics: precision medicine vs. patient privacy.
7. Lloyd, Travis, Nguyen, Tung, Levy, Karen, & Naaman, Mor (2025). Beyond Community Notes: A Framework for Understanding and Building Crowdsourced Context Systems for Social Media.
8. Ferrario, Andrea, Vinay, Rasita, Casserini, Matteo, & Facchini, Alessandro (2026). A Scoping Review of the Ethical Perspectives on Anthropomorphising Large Language Model-Based Conversational Agents.
9. John-Mathews, Jean-Marie (2021). Some Critical and Ethical Perspectives on the Empirical Turn of AI Interpretability.

Contact Author Email: duttapankojkumar@gmail.com

DOI: <https://doi.org/10.54121/2021111526>