

# Studies on Religion and Philosophy



# Studies on Religion and Philosophy

ISSN: 3067-0853

Volume 1, Issue 4, 2025

Quarterly (Issue No. 4)



**Editor-in-chief:** Wan Xing

**Associate Editor:** Dr. Lei Chuan

## Editorial Board Members:

Dr. Carimo Mohomed

Dr. Hong Chang

Dr. Li Jia

Dr. Xiqi Li

Dr. Caixia Liu

Dr. Hui Shen

Dr. Haiyang Zeng

Dr. Patricia A . Argüelles

Dr. Al-Tahir Muhammad Abd Al-Halim

Dr. Maziar Mozaffari Falarti

**Cover Design:** Li Wenjie

**Publishing Unit:** ConnectSix Scholar Publishing INC

**Publisher's website:** <http://www.cscholar.com/>

## Publisher's address:

6547 N Academy Blvd #2265

Colorado Springs CO 80918

US

**Website of the journal *Studies on Religion and Philosophy*:**

<https://srp.cscholar.com/>

(The content of this journal may not be reproduced without permission.)

## Table of contents

Religious Subjectivity in the Age of Artificial Intelligence: The Reshaping of Religious Experience under Technological Rationality from a Buddhist Perspective	Wan Xing Lei Chuan	1-15
Algorithmic Rationality” and “Ultimate Concern”: Reconfiguring Religious Rationality in the Age of Artificial Intelligence	Wan Xing	16-28
Can Artificial Intelligence “Understand” Religion? A Philosophical Inquiry into Meaning, Symbol, and Transcendence	Othman Al-Farsi	29-40
From Sacred Sanctuary to Intelligent System: A Religious-Studies Reflection on the Transformation of Temple Space under Artificial Intelligence	Ouyang Shaokang	41-52
Artificial Intelligence and the Religious Imagination of the Future: A Comparative Study of Eschatology, Soteriology, and Technological Utopianism	Zhanyi Li	53-62

# Religious Subjectivity in the Age of Artificial Intelligence: The Reshaping of Religious Experience under Technological Rationality from a Buddhist Perspective

Wan Xing <sup>1,\*</sup> Lei Chuan <sup>2</sup>

<sup>1</sup> Donghua Chan Temple, Shaoguan 512600, China

<sup>2</sup> Donghua Academy of China Studies, Guangzhou 510630, China;

**\* Correspondence:**

Wan Xing

wanxing@cscholar.com

*Received: 9 December 2025/Accepted: 26 December 2025/Published online: 31 December 2025*

## Abstract

Artificial intelligence (AI) has moved from being an external instrument to an environment that increasingly mediates attention, interpretation, and judgment. In religious life, this shift is not merely practical; it is structural. It reconfigures how religious experience is organized, how authority is distributed, and how religious subjectivity is formed. This article examines these transformations through a Buddhist philosophical lens, arguing that the rise of “algorithmic rationality” (a data-driven, predictive, optimizing mode of reason) reshapes the experiential conditions presupposed by Buddhist practice—especially the cultivation of mindfulness (*smṛti*), concentration (*samādhi*), and wisdom (*prajñā*). Drawing on core Buddhist doctrines—dependent origination (*pratītyasamutpāda*), non-self (*anātman*), the dynamics of craving and grasping (*trṣṇā/upādāna*), and the ethical primacy of intention (*cetanā*)—the paper analyzes how AI reorganizes the field of experience by: (1) externalizing and commodifying attention, (2) delegating hermeneutic labor to machine mediation, (3) accelerating temporality and weakening the pedagogical value of “slow practice,” (4) introducing a new form of authority that is opaque yet persuasive, and (5) intensifying subtle forms of attachment to metrics, personalization, and cognitive convenience. The article distinguishes the legitimate instrumental use of AI from its deeper tendency to colonize meaning-making and proposes a Buddhist-informed normative framework: “technological non-appropriation,” “digital restraint,” and “contemplative accountability.” It concludes that Buddhism can engage AI constructively only by reasserting the irreducibility of liberation to optimization and by safeguarding the soteriological structure of practice against the reification and outsourcing of subjectivity.

**Keywords:** Artificial Intelligence; Algorithmic Rationality; Buddhist Philosophy; Religious Subjectivity

## 1. The Problem of Religious Subjectivity Under Algorithmic Conditions

### 1.1. From “Tools” to “Conditions”: Why AI Changes the Question

Discussions of technology and religion often assume a stable structure of religious life into which new tools are inserted. A printing press distributes sutras; a microphone amplifies chanting; a website announces temple activities. In such cases, technology appears as an external means, and the theoretical question becomes largely ethical or institutional: Are the uses appropriate? Do they distort tradition? Who controls access?

Artificial intelligence differs in scale and structure. It is not simply another medium for distribution. AI increasingly functions as a cognitive environment: it filters what is seen, prioritizes what is attended to, summarizes what is read, recommends what is practiced, and predicts what is desired. In this sense, it affects the conditions under which meaning and agency arise. The relation between the religious subject and religious content is no longer direct, even when mediated by texts and teachers; it is increasingly triangulated by systems that optimize attention and interpretability (Tampubolon & Nadeak, 2024). This shift is especially significant for Buddhism because Buddhist practice is centrally concerned with the conditioning of experience itself—how attention is trained, how craving arises, how conceptualization reifies, and how insight dismantles mistaken views.

When AI becomes an interpretive and attentional infrastructure, it does not merely alter what Buddhists do; it may alter how Buddhist experience becomes possible at all. The theoretical issue therefore becomes a matter of religious subjectivity: what is the status of the practitioner’s agency, awareness, and responsibility under algorithmic mediation?

### 1.2. Religious subjectivity in Buddhism: a paradoxical focus

At first glance, “religious subjectivity” may seem an awkward category for Buddhism. Buddhism famously teaches non-self (*anātman*): the denial of a permanent, independent subject. Yet Buddhism is also intensely concerned with the transformation of experience—precisely the domain that modern thought associates with subjectivity. Buddhist practice presupposes that suffering (*duḥkha*) arises through ignorance and grasping, which are enacted through patterns of attention and identification. The path (*mārga*) is not a theoretical doctrine alone but a disciplined reconfiguration of lived experience through ethics (*sīla*), concentration (*saṃādhi*), and wisdom (*prajñā*). This implies a functional notion of subjectivity: not a metaphysical self-substance, but a dynamic process of awareness and response that can be cultivated toward liberation. Accordingly, the question is not whether Buddhism “has” a subject. Rather, the question is: what kind of experiential agency is presupposed by Buddhist soteriology, and how is that agency reorganized when algorithmic rationality becomes a pervasive mediator of cognition and desire?

### 1.3. The thesis and argumentative strategy

This paper advances three claims:

(1) AI embodies a distinct form of technological rationality—algorithmic rationality—that tends to reorganize attention, interpretation, temporality, and authority.

(2) These reorganizations structurally reshape Buddhist religious experience by encouraging externalized attention, outsourced hermeneutics, accelerated practice-temporality, and new attachment-patterns linked to personalization and metrics.

(3) Buddhist philosophy provides both diagnostic and normative resources to engage AI without surrendering the soteriological integrity of practice; however, this requires deliberate restraint and a reassertion of contemplative accountability.

Methodologically, the paper combines conceptual analysis of algorithmic rationality with Buddhist doctrinal and philosophical resources. It does not aim to offer sociological survey data or ethnographic descriptions. Instead, it provides a theoretical account of structural transformation, clarifying why the AI question for Buddhism is fundamentally a question about the formation of experience.

## 2. Algorithmic Rationality: From Instrumental Means to Cognitive Governance

### 2.1. Instrumental Rationality and Its Limits

Modern discussions often conceptualize technological rationality as instrumental: a neutral calculus of means to ends. On this view, a tool is morally and spiritually indifferent; what matters is the intention of its user. Much of Buddhist engagement with modern technology has implicitly followed this model, seeing technology as a set of supports (*upāya*, “skillful means”) that can be used for wholesome ends.

Yet instrumental rationality becomes inadequate when the tool begins to shape the user’s ends. AI systems are not passive instruments. They infer preferences, generate goals by optimizing engagement, and reorder choices through subtle incentives. The user’s intention remains relevant, but it is increasingly situated within an engineered environment that influences attention and desire.

### 2.2. Key Features of Algorithmic Rationality

Algorithmic rationality, understood as a mode of reasoning and governance, does not merely use numbers; it reorders reality so that what is real and actionable is what can be rendered into data. Quantification and legibility operate as a filtering principle: phenomena that resist measurement—ambivalence, silence, moral struggle, contemplative depth—are pushed to the margins, while what can be counted becomes privileged as “evidence” of truth and value. Prediction then extends this logic temporally by treating the future as a statistically manageable horizon, in which novelty, freedom, and transformation are conceptually reduced to variations of past behavioral regularities. Optimization turns prediction into governance: once the system can anticipate behavior, it can shape it, continuously adjusting informational environments to

maximize a metric that stands in for the “good” (attention, retention, affective satisfaction). Automation completes the circuit by delegating micro-decisions—what to see, what to read, what to practice, what to trust—to processes that bypass reflective deliberation. Finally, opacity and asymmetry introduce an epistemic imbalance: users are governed by outputs whose underlying reasoning is largely inaccessible, while those who design or control the system possess disproportionate power to define categories, defaults, and incentive structures. Taken together, these five features form a self-reinforcing architecture in which human perception and choice are increasingly pre-structured by computational systems, often without the subject’s explicit awareness of the conditioning at work.

These features matter religiously precisely because religion—especially Buddhism—cannot be reduced to information acquisition or behavioral compliance. Buddhist practice concerns the transformation of the conditions of experience: how craving arises, how attention is trained, how views are loosened, and how wisdom emerges through disciplined ethical and contemplative cultivation (Chattpadhyay, 2022). Algorithmic rationality, by contrast, tends to treat the inner life as a domain to be rendered legible, predicted, and optimized. In doing so, it converts qualitative transformation into measurable proxies (minutes meditated, streaks maintained, calmness scores, “personalized” progress), and it risks replacing the Buddhist telos of liberation from grasping with a subtler telos of managed well-being and efficient self-regulation. Moreover, Buddhism insists that insight (*prajñā*) is inseparable from intentionality (*cetanā*) and cannot be outsourced: the path requires the practitioner to encounter uncertainty, discomfort, and impermanence directly rather than to have these existential pressures engineered away through personalization and reassurance. The tension, then, is not that AI “contradicts” Buddhism at the level of doctrine, but that algorithmic rationality can reorganize the very experiential field in which Buddhist subjectivity is formed—externalizing attention, displacing hermeneutic responsibility, and fostering new attachments to prediction, optimization, and the illusion of control.

### 2.3. Algorithmic Rationality as a New Kind of Conditioning

Buddhism is, at its core, a theory of conditioning. Dependent origination describes how phenomena arise through conditions, including mental conditions: contact (*sparsā*), feeling (*vedanā*), craving (*trṣṇā*), clinging (*upādāna*), becoming (*bhava*) (Chowdhury, 2022). Buddhist practice is a disciplined intervention into conditioning: it aims to interrupt habitual chains and to cultivate wholesome conditions.

AI environments can be understood as novel conditioning systems. They systematically shape contact and attention (what is presented), modulate feeling (through affective design), stimulate craving (through novelty and personalization), and encourage clinging (through identity reinforcement, “for you” feeds, and quantified self-tracking). This does not mean AI is inherently “unwholesome,” but it means its default logics often align with craving and grasping rather than with restraint and insight.

## 2.4. Why AI Challenges Religious Subjectivity

Artificial intelligence challenges religious subjectivity not by denying belief or prohibiting practice, but by subtly reconfiguring the loci of responsibility that religious life presupposes (Umbrello, 2023; Ahmed et al., 2025). Religious subjectivity—particularly in Buddhism—depends on the practitioner’s cultivated capacity to own attention, to engage in interpretive labor, and to assume ethical responsibility through intentional action (*cetanā*). Algorithmic rationality, however, systematically displaces these capacities into external infrastructures: attention is steered by ranking and recommendation systems rather than by mindful resolve; interpretation is increasingly delegated to automated summaries and explanations that pre-structure understanding; and ethical deliberation is shaped through nudges, metrics, and social feedback mechanisms that operate prior to conscious reflection. The result is not merely a loss of focus or an increase in convenience, but a deeper reorganization of agency itself. In Buddhist terms, agency is neither an autonomous self nor a passive mechanism; it is a conditioned yet cultivable capacity to respond wisely to experience. When the conditions that shape intention and awareness are technologically engineered in advance—often invisibly—the practitioner’s field of moral and contemplative responsibility is narrowed and redirected. What is at stake, therefore, is the integrity of the experiential space in which Buddhist practice unfolds: a space that must allow for deliberate attention, interpretive struggle, and ethically grounded intention, rather than one pre-scripted by systems optimized for prediction, efficiency, and behavioral management.

## 3. Buddhist Religious Subjectivity and the Structure of Experience

### 3.1. Non-self as a Discipline, not a Senial of Agency

The doctrine of non-self denies a permanent, independent ego-substance, but it does not deny the efficacy of intentional action. Karma is unintelligible without intention; liberation is unintelligible without cultivation. Buddhism thus maintains a middle path: it refuses metaphysical selfhood while insisting on moral and contemplative responsibility.

Buddhist subjectivity can therefore be understood as a process of habituation and de-habituation. The practitioner is a nexus of conditions: mental factors, embodied routines, ethical commitments, and interpretive frames. Religious transformation occurs when these conditions are reoriented away from grasping and toward insight and compassion.

### 3.2. The Triad of Training and the Soteriological Structure

Buddhist practice is classically organized around the three trainings:

- (1) *Śīla* (ethical discipline): restraint, non-harm, and the formation of wholesome conduct.
- (2) *Samādhi* (concentration): stabilization and deepening of attention.
- (3) *Prajñā* (wisdom): insight into impermanence, non-self, and dependent origination.

This structure implies that religious experience is not primarily “content consumption.” It is transformative training. The subject is formed through repeated practice, not through episodic inspiration. AI systems, however, operate in an economy of consumption and engagement. Thus

the key question becomes: can AI-mediated environments sustain the disciplines that Buddhism regards as essential to liberation?

### 3.3. Attention as the Axis of Buddhist Subject Formation

Mindfulness (*smṛti*) is not merely present-moment awareness; it is a cultivated ability to remember the task, to return to the object, to discern wholesome from unwholesome, and to recognize arising phenomena without grasping. Attention is therefore both epistemic and ethical. It is the medium through which craving is known and loosened. In algorithmic environments, attention is commodified. It becomes a scarce resource competed for by platforms optimized for engagement. Even when Buddhist content is delivered through such platforms, the underlying architecture tends to fragment attention, prioritize novelty, and reward quick gratification. This creates a structural conflict with the Buddhist cultivation of sustained, non-grasping awareness.

### 3.4. Interpretation and “Right View”

Buddhism is not anti-conceptual, but it treats concepts as provisional. “Right view” is not merely correct doctrine; it is a guiding orientation that supports practice. Interpretation is therefore integral to subject formation: how one reads a teaching shapes how one practices.

Traditional Buddhist hermeneutics involves teacher-student transmission, communal practice, and repeated reflection. Interpretation is enacted through time and relationship. AI-generated explanations may offer accessibility, but they risk converting interpretation into a consumable product. When “understanding” is delivered instantly, the practitioner may bypass the slow, transformative labor of grappling—precisely the labor through which views are refined and attachment to conceptual certainty is weakened.

### 3.5. Temporality: the Pedagogy of Slowness

Buddhist practice is temporally thick. It presupposes repetition, patience, and gradual cultivation. Even sudden insight traditions presuppose long preparation and ethical grounding. The temporal structure matters because it reshapes desire: it trains the practitioner to tolerate discomfort, to remain with arising phenomena, and to see impermanence directly.

AI systems, by contrast, are temporally thin. They accelerate access, compress learning, and create expectations of immediate results. This acceleration can be useful (e.g., access to teachings), but it risks weakening the formation of patience and the capacity to endure the slow unfolding of insight.

## 4. How AI Reshapes Buddhist Religious Experience: Mechanisms of Transformation

This chapter analyzes concrete mechanisms by which algorithmic rationality restructures the field of Buddhist practice. The aim is not to demonize technology but to clarify how mediation works.

#### 4.1. Externalization of Attention: from Mindfulness to Managed Focus

In Buddhist practice, mindfulness (*smṛti*) is not a passive state but an actively cultivated capacity that involves remembering the object of practice, sustaining attention through effort, and repeatedly returning to awareness when distraction arises. This cultivation presupposes that attention is something for which the practitioner bears responsibility: one learns, through discipline, to notice impulses, to resist compulsive reactivity, and to remain present with phenomena as they unfold. Algorithmic systems, by contrast, are designed to manage attention from the outside. Through ranking mechanisms, push notifications, and personalized feeds, they pre-structure what appears, when it appears, and how urgently it demands response. When Buddhist teachings and practices are primarily encountered within such AI-mediated environments, attention subtly shifts from being intention-led to being stimulus-led. Engagement is no longer grounded primarily in the practitioner's deliberate resolve to practice, but in externally generated prompts that invite, interrupt, and redirect focus. Over time, this alters the phenomenology of practice itself: attention becomes reactive rather than reflective, responsive to cues rather than anchored in chosen commitment.

This externalization of attention also transforms the temporal and affective texture of Buddhist practice. Sustained cultivation—sitting with a single object, returning patiently to the breath, or dwelling with difficult mental states—gives way to fragmented micro-engagements dispersed across digital sessions. Practice risks being reorganized into brief, consumable units that fit the rhythms of the attention economy rather than the pedagogical rhythms of contemplative training (O'Donnell, 2015). Moreover, because algorithmic environments reward emotional salience and novelty, even religious or “mindfulness” content is incentivized to hook attention through reassurance, inspiration, or aesthetic appeal. From a Buddhist perspective, this fosters subtle forms of craving (*trṣṇā*): the mind becomes habituated to seeking the next gratifying stimulus, even under the guise of spiritual engagement. The danger, therefore, is not merely that practitioners become distracted, but that their capacity to remain with experience without grasping—precisely the capacity mindfulness is meant to cultivate—is structurally weakened by an environment that continually trains attention to move, select, and desire.

#### 4.2. Outsourcing Hermeneutics: the Machine as Interpreter

AI tools can translate sutras, summarize commentaries, answer doctrinal questions, and generate interpretations. This can lower barriers and broaden access—an undeniably positive possibility. Yet the philosophical risk is the outsourcing of hermeneutic responsibility.

Buddhist understanding is not merely cognitive; it is existential. Interpretive labor is part of transformation: struggling with a teaching reveals one's attachments, assumptions, and cravings for certainty. When AI delivers an interpretation that feels authoritative, the practitioner may mistake clarity for insight.

Moreover, AI interpretations are shaped by training data and optimization goals. They may produce plausible synthesis without lineage sensitivity, context awareness, or ethical accountability. In a tradition where “right view” supports liberation, an interpretive infrastructure

that is optimized for helpfulness rather than truth or soteriological integrity may subtly distort practice.

#### **4.3. Commodification Of Experience: Personalization, Metrics, And Spiritual Consumption**

Algorithmic rationality tends to personalize. In consumer contexts, personalization is marketed as empowerment: content “for you.” In Buddhist contexts, personalization can feel compassionate and skillful: tailored practices, customized teachings, mood-based meditations. However, personalization easily becomes commodification. The practitioner’s experience becomes a target for optimization: calmness, reduced stress, improved productivity. Buddhism does not deny conventional benefits, but it warns against confusing comfort with liberation. When practice is framed primarily as a lifestyle enhancement, it risks reinforcing the very selfing-process Buddhism seeks to expose. Metrics intensify this risk. If meditation apps quantify streaks, minutes, “focus scores,” or emotional states, practice becomes entangled with performance and identity. The practitioner may cling to metrics as proof of progress. In Buddhist terms, this is a subtle form of attachment (upādāna): clinging to spiritual identity and achievements.

#### **4.4. Acceleration of Temporality: the “Shortcut” Temptation**

AI promises efficiency: faster learning, smarter guidance, optimized routines. Applied to Buddhism, this can generate a “shortcut” mentality: if a system can deliver the best practice for one’s personality and schedule, why endure the uncertainty and difficulty of slow cultivation?

The Buddhist path, however, is not reducible to information selection. It is the transformation of craving and ignorance. The desire for shortcuts is itself a form of craving. An AI that caters to that craving may inadvertently strengthen it, turning the path into an instrument for self-improvement rather than liberation.

Acceleration also affects ritual and communal life. If digital systems streamline chanting, automate scheduling, or provide instant doctrinal answers, practitioners may lose the formative value of communal repetition and embodied participation. The ritual is not merely a means; it shapes attention and humility through its very slowness.

#### **4.5. Emergence of a New Authority: Opaque Persuasion and “Machine Charisma”**

Traditional Buddhist authority is complex: it involves lineage, ethical conduct, communal recognition, and the tested reliability of teachings. While institutions can fail, authority is, in principle, accountable: teachers can be questioned; communities can respond; texts can be traced.

Algorithmic authority is different. It is persuasive without being accountable. Its outputs feel neutral, objective, and efficient. Yet its reasoning is often opaque, and its incentives may be commercial. When an AI system answers a doctrinal question with confidence, it may acquire a form of “machine charisma.” Practitioners may defer to it because it seems omnipresent and instantly responsive.

This creates a structural risk: the displacement of the teacher-student relationship by a system that cannot embody ethics, cannot be held accountable in the same way, and cannot participate in

the lived context of practice. The danger is not that AI “lies,” but that it reorganizes epistemic trust in ways that weaken the relational and ethical dimensions of Buddhist learning.

#### **4.6. Reconditioning Craving: Novelty, Reassurance, and Identity Reinforcement**

Buddhist psychology emphasizes how craving arises through contact and feeling. Algorithmic environments intensify contact and manipulate feeling through novelty and reassurance. Even spiritual content can become a cycle of reassurance-seeking: the practitioner asks the machine for confirmation, comfort, and interpretation, repeatedly. This can produce a distinct modern pattern: the outsourcing of existential uncertainty. Buddhism, however, treats uncertainty and suffering as sites of insight. If AI continually soothes uncertainty with confident answers and personalized reassurance, it may reduce the practitioner’s encounter with *duḥkha* in its transformative dimension. Additionally, personalization often reinforces identity categories: “You are this kind of person; here is your practice.” Buddhism aims to loosen identity fixation. A system that constantly categorizes the user may subtly intensify selfing.

### **5. Philosophical Tensions and Risks for Buddhist Subjectivity**

#### **(1) Intention (cetanā) and responsibility under delegated judgment**

In Buddhist ethics, intention is central. Karma is not merely behavior; it is volition and orientation. When AI guides decisions—what to read, what to practice, how to interpret—responsibility becomes distributed. One might object: the user chooses to follow recommendations, so intention remains. Yet Buddhism also recognizes conditioning: choices arise within fields of influence (Wuthnow & Cadge, 2004). If an environment systematically nudges certain behaviors, it can weaken the practitioner’s reflective intentionality. This is not determinism; it is erosion of agency through habituation. A Buddhist critique would therefore ask: does AI-mediated practice strengthen or weaken the capacity for wholesome intention? Does it cultivate mindfulness, or does it cultivate reliance? Does it enable ethical reflection, or does it substitute automated judgment for deliberate discernment?

#### **(2) Reification and the subtle transformation of “emptiness” into data**

A central insight in many Buddhist traditions is emptiness (*śūnyatā*): the lack of inherent existence in phenomena. Emptiness is not nihilism; it is the insight that things exist dependently and conceptually. Reification is the opposite tendency: treating constructs as inherently real. Algorithmic systems inherently reify. They must classify, label, and predict. This is not a moral failure; it is a functional requirement. Yet when applied to meditation and spiritual development, reification becomes spiritually dangerous. It can transform states of mind into “objects” to be acquired, optimize compassion into a score, and treat insight as a measurable outcome. From a Buddhist perspective, this confuses conventional truth (useful categories) with ultimate insight (non-grasping understanding). It encourages attachment to representations. The practitioner may come to relate to practice as a managed project rather than a relinquishment of grasping.

### **(3) The displacement of “knowing” by “being informed”**

Buddhist wisdom is not mere information. It is a transformative knowing that changes one’s relation to experience. AI can deliver information at scale, but information is not insight. There is therefore a risk of confusing “being informed about Buddhism” with practicing Buddhism. This confusion is amplified when AI outputs are fluent and coherent. Fluency produces epistemic confidence. Yet Buddhist learning often requires encountering paradox, difficulty, and the limits of conceptualization. A system optimized for helpfulness may smooth out precisely the rough edges that provoke deep reflection.

### **(4) The temptation of spiritual convenience: comfort as the new telos**

In many modern contexts, Buddhism is reframed as a therapeutic technique for calmness and productivity. AI-driven wellness platforms can intensify this reframing. The implicit telos becomes comfort and performance. The dharma becomes a tool for better functioning within existing conditions. Buddhism does not deny the relief of suffering. But it diagnoses suffering at a deeper level: attachment to self, impermanence, and the delusion of control. If AI turns Buddhist practice into a tool for self-optimization, it may strengthen the very delusion of control Buddhism seeks to dismantle.

### **(5) Community, sangha, and the risk of disembodied religiosity**

Buddhism traditionally locates religious life within the framework of the Three Jewels—Buddha, Dharma, and Sangha—thereby affirming that community is not an optional supplement to individual practice but a constitutive condition for ethical formation, doctrinal clarification, and spiritual accountability. The sangha provides embodied contexts in which conduct is visible, correction is possible, and compassion is enacted through shared ritual, service, and mutual care. By contrast, AI-mediated religiosity tends toward disembodiment: practice becomes increasingly solitary, on-demand, and customized to individual preference. This shift has far-reaching implications. When practice is removed from communal settings, ethical accountability weakens, as one’s conduct is no longer shaped through relational exposure and correction but remains largely private and self-evaluated. Compassion, severed from concrete interpersonal obligation, risks becoming primarily affective or sentimental rather than practiced through sustained engagement with others’ needs. Rituals, which in traditional settings function as repetitive, embodied disciplines that cultivate humility, patience, and collective memory, are reduced to optional content streams that can be consumed or skipped without consequence. Likewise, the teacher-student relationship—central to Buddhist transmission and the gradual correction of misunderstanding—may be displaced by self-guided exploration mediated by algorithms, where guidance is plentiful but responsibility is diffuse and challenge is easily avoided. None of this implies that digital engagement is inherently incompatible with Buddhism; rather, it highlights that Buddhist subjectivity presupposes relational and embodied conditions that algorithmic environments, oriented toward personalization and convenience, do not support by default and may actively erode unless consciously counterbalanced.

## **(6) A new form of attachment: the “algorithmic self”**

Closely related to this disembodiment is the emergence of a new form of attachment that may be described as the “algorithmic self.” Even as Buddhism offers a sustained critique of substantial selfhood, modern technological environments continuously generate new identity formations. Under algorithmic rationality, the self is rendered as a profile: a statistically inferred pattern of preferences, behaviors, moods, and tendencies that is continually predicted, reinforced, and mirrored back to the user. This algorithmic self is not imposed coercively; it is produced through feedback loops that feel intuitive and affirming—“you like this,” “this suits you,” “people like you practice this way.” For Buddhist practice, this is especially perilous because it intensifies identification precisely at the point where practice seeks to loosen it. The practitioner may become attached to being “a mindful person,” “a committed Buddhist,” or “an advanced meditator,” not as provisional designations for functional roles, but as defended self-concepts sustained by metrics, personalization, and social visibility. Buddhist texts have long warned against attachment to views, spiritual achievement, and subtle forms of pride; AI environments risk producing a technologically amplified version of these tendencies, where identity is stabilized through quantified progress, curated self-presentation, and constant affirmation. Liberation thus becomes more difficult not because the self is crudely reassured, but because attachment becomes more refined, ambient, and difficult to detect, embedded within systems that continuously invite the practitioner to recognize themselves in the very patterns Buddhism urges them to see as empty and contingent.

## **6. A Buddhist Theoretical Response: Reclaiming Practice under Technological Rationality**

If artificial intelligence reshapes the conditions under which religious experience and subjectivity are formed, a Buddhist response cannot be limited to either technological enthusiasm or outright rejection. Buddhism has historically engaged new cultural and material conditions with remarkable adaptability, yet this adaptability has always been guided by a rigorous soteriological criterion: whether a given practice or tool reduces grasping (*upādāna*), clarifies intention (*cetanā*), and supports the cultivation of liberation-oriented wisdom (*prajñā*). The challenge posed by algorithmic rationality therefore requires a response that is neither reactive nor naïve, but critically discerning. This chapter articulates a set of interconnected Buddhist principles through which religious practice may be reclaimed under conditions increasingly shaped by technological rationality.

### **6.1. The Principle of Technological Non-Appropriation**

At the core of Buddhist practice lies the discipline of non-appropriation: the capacity to relate to phenomena without clinging, identification, or reification. Applied to the contemporary technological environment, this discipline may be reformulated as technological non-appropriation: the use of digital and algorithmic tools without allowing them to become objects of attachment, sources of identity, or substitutes for existential responsibility. From a Buddhist perspective, the ethical and spiritual danger of AI does not primarily reside in its presence, but in the subtle ways it invites appropriation—encouraging the practitioner to seek certainty,

reassurance, efficiency, or self-definition through technological mediation rather than through the transformative work of practice.

Technological non-appropriation therefore does not prohibit the use of AI, nor does it romanticize pre-digital forms of religiosity. Instead, it insists on ontological clarity: AI belongs to the domain of conventional supports (*samvṛti-satya*), not to the domain of liberation. When algorithmic systems are tacitly treated as authoritative interpreters of the Dharma, as reliable indicators of spiritual progress, or as personalized guides to “what one truly needs,” they begin to function as objects of grasping. The practitioner may then appropriate technology as a source of identity (“this is how I practice”), control (“the system knows what is best for me”), or reassurance (“my progress is verified”), thereby reinforcing precisely those patterns of attachment that Buddhist practice aims to loosen. Practically, the principle of technological non-appropriation thus implies disciplined limits: AI should not be treated as a definitive dharma authority, optimization metrics should not be mistaken for spiritual development, personalization should not be conflated with wisdom, and convenience should never be allowed to redefine the measure of the path. What is at stake is not technological usage as such, but the preservation of a non-grasping orientation toward all conditioned supports, including the most sophisticated ones.

## 6.2. Digital Restraint (*Samvara*) as Contemporary *Śīla*

Buddhist ethics has always emphasized restraint (*samvara*), particularly the guarding of the sense doors and the careful regulation of conditions that give rise to unwholesome mental states. In the context of an AI-driven attention economy, restraint acquires renewed centrality—not as ascetic moralism, but as a pragmatic protection of the conditions necessary for mindfulness and insight. Algorithmic systems are explicitly designed to capture, direct, and monetize attention; without intentional counter-practices, the practitioner’s attentional field becomes increasingly fragmented, reactive, and externally governed.

Digital restraint may therefore be understood as a contemporary extension of *śīla*, responsive to novel forms of conditioning. It involves limiting exposure to algorithmically optimized feeds that privilege novelty and emotional arousal over depth and continuity; establishing clear boundaries around notifications, compulsive checking, and passive scrolling; prioritizing spaces for study and practice that are not structured by engagement-maximization metrics; and cultivating intentional modes of entry into and exit from digital environments rather than remaining in a state of perpetual availability. Such restraint is not a rejection of technology, but a recognition that attention is the primary medium of Buddhist training. If attention is continuously captured, redirected, and exhausted by algorithmic demands, the cultivation of mindfulness becomes structurally undermined. Digital restraint thus functions as an ethical safeguard for contemplative life, preserving the interior space in which awareness can stabilize and insight can mature.

## 6.3. Contemplative Accountability: Preserving the Teacher, the Sangha, and Lived Correction

One of the most significant risks of algorithmic rationality lies in the displacement of accountability. Algorithmic authority is persuasive precisely because it appears neutral, efficient, and omnipresent, yet it remains fundamentally opaque and ethically unanswerable. Buddhist

practice, by contrast, has always relied on forms of accountability that are relational, embodied, and corrigible. For this reason, a Buddhist response to AI must insist on preserving dimensions of accountability that technological systems cannot replicate.

Teacher accountability remains central: authentic guidance is grounded not merely in informational competence, but in ethical conduct, lived experience, and contextual sensitivity to the practitioner's condition. Sangha accountability is equally indispensable: communal practice exposes blind spots, interrupts self-deception, and grounds compassion in concrete relationships rather than abstract ideals. Embodied accountability—through ritual participation, service, and ethical action—resists the reduction of Buddhism to a purely cognitive or psychological technique. While AI may assist learning by increasing access to texts and explanations, it cannot substitute for these relational structures through which Buddhist subjectivity is formed, challenged, and refined. Without such accountability, practice risks becoming privatized, self-validating, and insulated from the corrective friction that genuine transformation requires.

#### **6.4. Hermeneutic Humility: Resisting the Collapse of Insight into Explanation**

A further requirement for reclaiming Buddhist practice under technological rationality is the cultivation of hermeneutic humility. Buddhist traditions consistently distinguish between conceptual understanding and liberating insight, warning against the premature closure of inquiry through reified explanations. AI-generated interpretations and summaries can be valuable aids, but they must be held provisionally. The danger arises when explanation is mistaken for realization, and fluency for wisdom.

Hermeneutic humility involves returning repeatedly to primary texts and traditional commentarial contexts, treating AI outputs as preliminary orientation rather than final judgment, and allowing unresolved questions to remain open as sites of practice rather than problems demanding immediate resolution. It also involves recognizing that not-knowing, confusion, and interpretive struggle are not failures but integral dimensions of Buddhist pedagogy. Insight often emerges not through the accumulation of answers, but through sustained engagement with what resists easy comprehension. In this sense, humility before the limits of explanation is not anti-intellectual; it is a disciplined refusal to allow conceptual clarity to substitute for existential transformation.

#### **6.5. Reorienting “Skillful Means” (upāya) in the AI Era**

Buddhism has long demonstrated a capacity to adapt its forms through the principle of skillful means (upāya), adjusting pedagogical strategies to the needs and capacities of different contexts and audiences (Ma Rhea, 2018). However, upāya has never meant the uncritical adoption of whatever is effective or efficient. Skillful means are skillful only insofar as they are guided by wisdom and oriented toward liberation. In the age of artificial intelligence, this principle must be rearticulated with particular care.

AI tools may function as skillful means only if they genuinely support the three trainings—ethical discipline, concentration, and wisdom—rather than subtly undermining them. Accordingly, the use of AI in Buddhist contexts must be evaluated through normative questions: Does it strengthen ethical intention, or does it displace responsibility into automated systems? Does it

deepen attention, or does it fragment awareness through constant stimulation? Does it cultivate wisdom, or does it replace insight with readily consumable information? Does it strengthen sangha and compassion, or does it intensify isolation, self-curation, and identity fixation? Efficiency alone cannot serve as a criterion of skillfulness. A tool that accelerates access while weakening ethical intentionality, attentional stability, or communal accountability cannot be considered upāya in a genuinely Buddhist sense, regardless of its technological sophistication.

Taken together, these principles articulate a Buddhist theoretical response that neither rejects artificial intelligence nor allows it to redefine the structure of religious life. They affirm that liberation cannot be optimized, subjectivity cannot be outsourced, and wisdom cannot be automated. What Buddhism offers in the age of AI is not a competing technology, but a disciplined way of inhabiting conditioned systems without being dominated by them—a way of using tools without surrendering the work of awakening to the logic of calculation, prediction, and control.

## 7. Conclusion

The rise of artificial intelligence confronts Buddhism with a decisive question: can the tradition preserve the soteriological structure of practice in an environment increasingly optimized to capture attention, outsource interpretation, and accelerate desire?

This paper has argued that algorithmic rationality reshapes Buddhist religious experience structurally. It externalizes attention, outsources hermeneutic labor, accelerates temporality, introduces opaque authority, and fosters new forms of attachment tied to personalization and metrics. These shifts are not incidental. They reorganize the very conditions under which Buddhist subjectivity—understood as a cultivated process of ethical intention, mindful awareness, and liberating insight—can be formed.

Yet Buddhism also offers powerful resources for response. Dependent origination clarifies that AI environments are conditioning systems; non-self clarifies that liberation cannot be outsourced to a system; the primacy of intention clarifies that responsibility remains central even under mediation; and the three trainings clarify that wisdom is not reducible to informational optimization.

The most constructive engagement with AI is therefore neither rejection nor uncritical adoption. It is disciplined discernment: technological non-appropriation, digital restraint, contemplative accountability, and hermeneutic humility. On this basis, Buddhism can use AI as a conventional support while resisting its tendency to colonize meaning-making. The ultimate criterion remains the same as in classical Buddhism: whether a practice reduces grasping, deepens compassion, and opens the possibility of liberation.

### Author Contributions:

All authors have read and agreed to the published version of the manuscript.

## Funding:

This research received no external funding.

## Institutional Review Board Statement:

Not applicable.

## Informed Consent Statement:

Not applicable.

## Data Availability Statement:

Not applicable.

## Conflict of Interest:

The author declares no conflict of interest.

## References

- Ahmed, S., Sumi, A. A., & Aziz, N. A. (2025). Exploring Multi-Religious Perspective of Artificial Intelligence. *Theology and Science*, 23(1), 104-128.
- Chattopadhyay, M. (2022). Contemplation: its cultivation and culmination through the Buddhist glasses. *Frontiers in Psychology*, 12, 800281.
- Chowdhury, S. B. (2022). Unfolding Dependent origination: a Psychological analysis for Disclosing the root of the afflictive State of Mind. *Journal of International Buddhist Studies*, 13(1), 11-43.
- Ma Rhea, Z. (2018). Buddhist pedagogy in teacher education: cultivating wisdom by skillful means. *Asia-Pacific Journal of Teacher Education*, 46(2), 199-216.
- O'Donnell, A. (2015). Contemplative pedagogy and mindfulness: Developing creative attention in an age of distraction. *Journal of Philosophy of Education*, 49(2), 187-202.
- Tampubolon, M., & Nadeak, B. (2024). Artificial intelligence and understanding of religion: A moral perspective. *International Journal of Multicultural and Multireligious Understanding*, 11(8), 903-914.
- Umbrello, S. (2023). The intersection of Bernard Lonergan's critical realism, the common good, and artificial intelligence in modern religious practices. *Religions*, 14(12), 1536.
- Wuthnow, R., & Cadge, W. (2004). Buddhists and Buddhism in the United States: The scope of influence. *Journal for the scientific study of religion*, 43(3), 363-380.

## License:

Copyright (c) 2025 Author.

All articles published in this journal are licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited. Authors retain copyright of their work, and readers are free to copy, share, adapt, and build upon the material for any purpose, including commercial use, as long as appropriate attribution is given.

# Algorithmic Rationality” and “Ultimate Concern”: Reconfiguring Religious Rationality in the Age of Artificial Intelligence

Wan Xing <sup>1,\*</sup>

<sup>1</sup> Donghua Chan Temple, Shaoguan 512600, China

**\* Correspondence:**

Wan Xing

wanxing@cscholar.com

*Received: 10 December 2025/ Accepted: 26 December 2025/ Published online: 31 December 2025*

## Abstract

The rapid diffusion of artificial intelligence has not only transformed economic production, social communication, and political governance, but has also reconfigured the epistemic and existential conditions under which religious meaning is articulated and sustained. This article examines the tension between “algorithmic rationality”—a data-driven, predictive, and optimization-oriented mode of reasoning—and religious rationality, understood here through the category of “ultimate concern.” While technological rationality has long been analyzed as instrumental or functional, algorithmic rationality increasingly operates as a comprehensive cognitive environment that mediates attention, interpretation, and judgment. This paper argues that such rationality challenges religion not primarily by negating belief, but by subtly reshaping the conditions under which ultimate questions, transcendence, and existential orientation become intelligible. Drawing on philosophy of technology, contemporary theology, and religious studies, the article proposes a redefinition of religious rationality that neither retreats into anti-technological defensiveness nor capitulates to computational reductionism. Instead, it conceptualizes religious rationality as a non-optimizable, reflexive, and responsibility-bearing mode of reason that resists full translation into algorithmic terms. The paper concludes that the future viability of religion in highly computational societies depends on preserving the irreducibility of ultimate concern against the expanding horizon of algorithmic governance.

**Keywords:** Algorithmic Rationality; Ultimate Concern; Religious Rationality; Artificial Intelligence; Philosophy of Technology

## 1. Introduction

Artificial intelligence (AI) has increasingly moved beyond the status of a discrete technological tool to become an infrastructural condition shaping contemporary forms of cognition,

communication, and decision-making. Algorithmic systems filter information, rank relevance, predict behavior, and automate judgment across domains ranging from finance and medicine to education and governance. In this context, the question of religion and AI can no longer be adequately framed as a matter of ethical application or institutional adaptation. Instead, it raises a deeper theoretical problem: how does algorithmic rationality reshape the conditions under which religious meaning, belief, and commitment are rendered intelligible? Much of the existing discourse on religion and technology presupposes a relatively stable conception of religious rationality into which new technologies are introduced as external challenges or instruments. On this view, religion either resists technology, accommodates it pragmatically, or reinterprets it symbolically. Yet artificial intelligence complicates this framework. Unlike earlier technologies that primarily extended human capacities, AI increasingly participates in cognitive operations themselves—classification, inference, evaluation, and decision-making. It thereby intervenes not only in what religious actors do, but in how rationality itself is structured.

This paper argues that AI confronts religion with a distinctive form of rationality—algorithmic rationality—that differs qualitatively from classical instrumental reason. Algorithmic rationality does not merely calculate means toward given ends; it redefines relevance, value, and plausibility through datafication, prediction, and optimization. As such, it poses a structural challenge to religious rationality, particularly insofar as religion has traditionally articulated its claims through reference to what Paul Tillich famously called “ultimate concern”—that which demands unconditional seriousness and orients the totality of existence.

The central thesis of this article is that algorithmic rationality and religious rationality are not simply competing epistemic systems, but represent divergent modes of relating to meaning, contingency, and transcendence. While algorithmic rationality seeks to render the world predictable, optimizable, and governable, religious rationality is oriented toward what exceeds calculation: finitude, suffering, guilt, hope, and the question of ultimate meaning. The tension between these rationalities becomes especially acute in highly computational societies, where algorithmic mediation increasingly defines what counts as reasonable, relevant, and real.

To develop this argument, the paper proceeds in six stages. First, it clarifies the concept of algorithmic rationality and distinguishes it from earlier forms of technological and instrumental reason. Second, it reconstructs the notion of religious rationality through the category of ultimate concern. Third, it analyzes how algorithmic rationality reconfigures attention, temporality, and authority in ways that marginalize ultimate questions. Fourth, it examines the risk of reducing religion to functional or therapeutic roles within algorithmic systems. Fifth, it proposes a redefinition of religious rationality as reflexive, non-optimizable, and responsibility-centered. Finally, it reflects on the implications of this redefinition for the future of religion in AI-mediated societies.

## 2. Algorithmic Rationality: From Instrumental Reason to Cognitive Environment

Classical critiques of modern technology have long centered on the problem of instrumental rationality—the reduction of reason to a calculus of efficiency, predictability, and control. From

Max Weber's diagnosis of rationalization as the disenchantment of the world to the Frankfurt School's analysis of instrumental reason as a form of domination, technology has been interpreted as a mode of reasoning that subordinates values, meanings, and ends to technical calculability (Horkheimer & Adorno, 2002). In these critiques, technology appears primarily as an external means: a set of tools deployed by social systems to achieve predefined goals, often at the expense of ethical reflection and human autonomy. While such analyses remain indispensable for understanding modernity, they are insufficient for grasping the distinctive logic of algorithmic rationality that characterizes contemporary artificial intelligence. Algorithmic rationality represents not merely an intensification of instrumental reason, but a qualitative transformation in how rationality operates within social and cognitive life. Unlike classical technological rationality, which presupposed a relatively stable distinction between human subjects and technical instruments, algorithmic rationality increasingly collapses this distinction by embedding computational processes directly into perception, judgment, and decision-making. Algorithms do not simply execute human intentions; they participate in shaping what counts as relevant information, plausible interpretation, and reasonable action. As a result, rationality itself becomes partially automated, distributed across human-machine assemblages, and increasingly opaque to reflective scrutiny.

One decisive feature of algorithmic rationality is its data-centric ontology. Reality becomes intelligible primarily insofar as it can be translated into data points, variables, and statistical correlations. Experiences, behaviors, preferences, and even emotions are rendered as datasets to be collected, processed, and optimized. Phenomena that resist quantification—such as silence, ambiguity, transcendence, moral struggle, and existential anxiety—are not explicitly denied, but they are structurally marginalized. They appear as epistemically “thin,” methodologically inconvenient, or operationally irrelevant. In this sense, algorithmic rationality does not refute non-quantifiable dimensions of human life; rather, it deprioritizes them by redefining epistemic legitimacy itself. What cannot be measured is not false, but it is increasingly treated as negligible.

A second defining feature of algorithmic rationality is its predictive orientation. Whereas classical rationality aimed primarily at explanation and control of present phenomena, algorithmic systems are oriented toward forecasting future behavior. Through machine learning and pattern recognition, the future is no longer approached as an open horizon of possibility, but as a probabilistic extension of past data. Risk replaces contingency, and anticipation replaces deliberation. This predictive logic subtly reshapes the human experience of time. The future appears less as a domain of moral responsibility and existential openness, and more as a space to be preemptively managed through optimization. Such a temporal structure stands in tension with religious worldviews that understand the future in terms of promise, judgment, or eschatological hope—dimensions that cannot be reduced to statistical projection.

Third, algorithmic rationality is fundamentally optimizing. Decisions are evaluated according to performance indicators that function as proxies for value: engagement, efficiency, accuracy, satisfaction, or compliance. Normative questions—what is good, meaningful, or worthy of commitment—are translated into technical problems of optimization. This translation does not eliminate values; rather, it embeds them implicitly within metrics and design choices. Yet because

these values are encoded procedurally rather than articulated discursively, they become difficult to contest. The good is no longer debated; it is operationalized. In this way, algorithmic rationality narrows the space of ethical reflection by substituting calculable success for reflective judgment.

Crucially, algorithmic rationality operates not primarily as a tool but as a cognitive environment. Individuals do not simply use algorithms as external instruments; they increasingly inhabit algorithmically structured spaces in which perception, attention, and interpretation are continuously pre-shaped. Search engines, recommendation systems, predictive analytics, and generative models filter reality before it reaches conscious awareness. Visibility, relevance, and plausibility are no longer neutral givens, but outcomes of algorithmic mediation. This environmental character marks a significant departure from earlier technologies, which typically functioned as discrete devices external to the cognitive process. Algorithmic rationality, by contrast, operates at a pre-reflective level, shaping the conditions under which reflection itself becomes possible. This transformation has far-reaching implications for religion. Religious rationality has historically depended on practices and spaces that resist immediate utility and optimization: liturgy that repeats rather than innovates, ritual that values form over efficiency, silence that suspends productivity, confession that exposes vulnerability rather than performance, and contemplation that lingers rather than accelerates. These practices cultivate a mode of attention oriented toward depth, patience, and transcendence. When religious life is increasingly mediated by algorithmic environments optimized for engagement and convenience, the internal logic of these practices is subtly altered. Ritual risks becoming content, contemplation becomes a technique, and transcendence is reframed as psychological well-being.

Moreover, algorithmic rationality introduces a novel form of authority that differs fundamentally from both traditional religious authority and modern scientific expertise. Algorithmic authority is opaque yet persuasive. Its legitimacy does not rest on explicit truth claims, moral exemplarity, or institutional legitimacy, but on performance and usability. Outputs are trusted because they function smoothly, deliver results, and reduce friction. This “working” quality generates a form of epistemic confidence that does not require understanding. Users defer not because they are convinced, but because the system appears reliable and efficient. Authority thus shifts from justification to functionality. For religion, this shift poses a distinctive challenge. Religious authority has traditionally been grounded in narrative coherence, symbolic depth, ethical integrity, and communal recognition—forms of legitimacy that require interpretation, trust, and moral accountability. Algorithmic authority, by contrast, is largely unaccountable in existential or ethical terms. Its decisions are often inscrutable, and its value assumptions are embedded rather than argued. When such authority increasingly mediates access to religious texts, teachings, and communities, it risks reshaping religious rationality itself, encouraging deference to computational outputs rather than engagement with tradition, interpretation, and moral struggle.

In sum, algorithmic rationality cannot be adequately understood as a mere extension of instrumental reason. It represents a transformation in the structure of rationality itself—from a tool-oriented logic of means and ends to an environmental logic that governs attention, temporality, and authority. This transformation does not eliminate religion, but it alters the conditions under which religious rationality can be articulated and sustained. To grasp the

significance of artificial intelligence for religion, one must therefore move beyond questions of application and ethics to examine how algorithmic rationality reconfigures the cognitive and existential environment in which ultimate concern becomes meaningful at all.

### 3. Religious Rationality and the Concept of Ultimate Concern

To assess the impact of algorithmic rationality on religion, it is first necessary to clarify what is meant by religious rationality itself. In much modern discourse, religion is implicitly treated either as a pre-rational worldview superseded by scientific reason, or as an irrational residue that persists despite modernization. Both approaches misconstrue the nature of religion by reducing it to belief-content or cultural inheritance. Against this reduction, the present analysis understands religious rationality as a distinctive mode of reasoning oriented toward what Paul Tillich famously termed “ultimate concern” (Tillich, 1957). Ultimate concern refers not to a specific doctrinal object, but to that which claims unconditional seriousness and organizes the totality of a person’s existence. It functions as a horizon of meaning rather than a discrete proposition, shaping how all other concerns—practical, moral, or cognitive—are hierarchically ordered and interpreted. In this sense, ultimate concern is not merely something one happens to believe in; it is that in relation to which one lives. It structures priorities, defines what is worth sacrificing for, and determines how finitude, suffering, and mortality are understood. Religious rationality, therefore, is not exhausted by theological doctrines or institutional practices. It is a form of reason that articulates, sustains, and critically reflects upon the ultimate orientation of human existence. To speak of religious rationality is to acknowledge that religion involves a coherent, though non-instrumental, way of making sense of reality, one that addresses dimensions of meaning inaccessible to purely technical or scientific rationality.

Religious rationality, understood in this way, is not opposed to reason as such. Rather, it operates according to a different logic than technical, scientific, or algorithmic rationality. Whereas technical rationality seeks efficiency, prediction, and control, and scientific rationality seeks explanation and empirical adequacy, religious rationality addresses questions that cannot be resolved through optimization or empirical verification. Questions such as why there is something rather than nothing, what it means to live well in the face of finitude, how suffering and guilt should be interpreted, and whether hope can be justified beyond calculable outcomes, are not inefficient because they lack definitive solutions. They are constitutive of human self-understanding precisely because they resist closure. Religious rationality does not aim to eliminate uncertainty, but to inhabit it meaningfully. This orientation toward ultimate questions gives religious rationality a distinctive existential depth. It engages not only the intellect, but the whole person—affectionately, morally, and practically. Religious reasoning is inseparable from commitment, yet this commitment is not blind assent to propositions. Rather, it is an existential stance that integrates belief, practice, and self-understanding. Faith, in this context, is best understood not as the suspension of reason, but as a mode of trust that remains responsive to doubt, critique, and transformation. Religious rationality thus includes moments of questioning, crisis, and reinterpretation as integral to its own operation.

Historically, religious rationality has been mediated through narratives, rituals, ethical norms, and symbolic systems that articulate the relationship between the finite and the infinite, the conditioned and the transcendent. Myth and narrative do not function as primitive explanations of natural phenomena, but as symbolic frameworks that orient existence within a meaningful cosmos. Rituals enact patterns of time and embodiment that resist purely instrumental temporality, reaffirming communal memory and existential orientation. Ethical norms translate ultimate concern into concrete forms of responsibility, binding individuals to obligations that cannot be justified solely in terms of utility. Symbols, finally, function not as literal representations but as mediations that participate in what they signify, pointing beyond themselves while remaining historically situated. Such mediation presupposes a capacity to dwell with ambiguity, paradox, and tension. Religious rationality does not dissolve the finite–infinite distinction, nor does it collapse transcendence into immanence. Instead, it sustains a dynamic relation between them. This is why religious symbols must be continuously interpreted and reinterpreted; they are meaningful precisely because they are not exhaustively transparent. In Tillich’s terms, symbols both reveal and conceal the ultimate, preventing it from being reduced to an object among others. Religious rationality, therefore, is intrinsically hermeneutical. It involves ongoing interpretation in response to historical change, existential challenge, and internal critique. This hermeneutical dimension renders religious rationality fundamentally reflexive. It is capable of turning back upon its own forms, practices, and doctrines, subjecting them to critique in light of the very ultimate concern they are meant to express. Such reflexivity distinguishes religious rationality from dogmatism, which absolutizes particular formulations and forecloses reinterpretation. At the same time, it also distinguishes religious rationality from relativism, which dissolves ultimate concern into a plurality of incommensurable preferences. Religious rationality maintains that some concerns are ultimate and demand commitment, but it also recognizes that all human articulations of the ultimate are historically conditioned and therefore revisable.

This balance between commitment and critique is a defining feature of religious rationality. Ultimate concern demands existential seriousness; it cannot be treated as a matter of convenience or personal taste. Yet because the ultimate cannot be fully grasped or possessed, religious rationality remains open to self-correction. Doubt is not simply the negation of faith, but one of its internal moments. In this sense, religious rationality is neither absolutist nor arbitrary. It is structured by an orientation toward meaning that transcends immediate interests, while remaining attentive to the limitations of human understanding. It is precisely this fragile balance that becomes increasingly difficult to sustain in algorithmically mediated environments. Algorithmic rationality tends to equate reason with calculability, responsiveness, and optimization. Questions that cannot be translated into data-driven incentives or measurable outcomes are easily sidelined as inefficient or irrelevant. Commitment without immediate payoff appears irrational, and ambiguity becomes a problem to be resolved rather than a condition to be inhabited. In such a context, ultimate concern risks being reinterpreted as a subjective preference, a therapeutic resource, or a lifestyle option, rather than as a binding horizon of meaning.

Moreover, the reflexivity intrinsic to religious rationality presupposes spaces of distance from immediate functional demands—spaces in which questioning, repentance, lament, and

contemplation can occur. When rationality is increasingly governed by algorithmic environments that privilege speed, engagement, and continuous responsiveness, the temporal and cognitive conditions for such reflexivity are undermined. Religious rationality is not eliminated, but it is pressured to conform to modes of reasoning that are fundamentally alien to its logic. Understanding religious rationality as oriented toward ultimate concern thus provides a critical lens for evaluating the cultural impact of artificial intelligence. It reveals that the challenge posed by algorithmic rationality is not simply one of doctrinal conflict or moral risk, but one of existential reconfiguration. What is at stake is whether contemporary societies can still sustain forms of reasoning that acknowledge finitude, transcendence, and responsibility beyond calculation. Religious rationality, in this sense, does not compete with algorithmic rationality on its own terms. Instead, it exposes the limits of any rationality that seeks to exhaust meaning through prediction, optimization, and control.

#### **4. Algorithmic Mediation and the Marginalization of Ultimate Questions**

One of the most significant and far-reaching effects of algorithmic rationality is the reconfiguration of attention. In digital environments governed by recommendation systems, attention is no longer merely a psychological capacity or ethical discipline; it becomes a scarce economic resource to be captured, measured, and monetized. Algorithms rank content according to engagement metrics—clicks, viewing time, emotional response, and interaction frequency—and users are continuously nudged toward what systems predict will sustain their interest. While this logic is economically efficient and technologically sophisticated, it has profound and often unintended consequences for religious life, particularly for forms of questioning and reflection that do not conform to the dynamics of attention optimization.

Ultimate questions—questions of meaning, finitude, guilt, suffering, hope, and transcendence—are rarely “engaging” in the algorithmic sense. They demand patience rather than immediacy, repetition rather than novelty, and sustained reflection rather than emotional stimulation. They often involve discomfort, uncertainty, and the absence of clear resolution. Precisely for these reasons, they are structurally disadvantaged within attention economies designed to maximize retention and affective responsiveness. Algorithmic systems privilege content that can quickly capture interest and produce measurable engagement, whereas ultimate questions typically require silence, slowness, and a willingness to dwell with what resists closure. As a result, such questions are not simply ignored; they are rendered less visible, less salient, and less culturally plausible within algorithmically mediated environments. This marginalization does not necessarily take the form of overt exclusion. Rather, it operates through translation and reframing. Religious content that circulates successfully within algorithmic platforms often does so by adapting itself to the dominant logic of functionality. Ultimate concerns are reframed in therapeutic or motivational terms—stress reduction, emotional regulation, resilience, self-care, or personal growth. While these ref framings are not illegitimate in themselves, they subtly transform the meaning of religion by relocating transcendence within the horizon of individual well-being and performance. The question of ultimate meaning becomes a question of psychological utility; salvation is recoded as optimization; and faith is reinterpreted as a resource for coping. In this

way, algorithmic mediation does not eliminate religion, but it redefines its rationality by subordinating it to functional outcomes that are legible within data-driven systems. Attention, however, is not the only dimension reshaped by algorithmic mediation. Temporality is also profoundly transformed. Algorithmic rationality privileges immediacy, responsiveness, and real-time adjustment. Feedback loops operate continuously, and value is assigned to speed, adaptability, and constant availability. Time is experienced as a sequence of actionable moments rather than as a horizon of formation. Religious traditions, by contrast, are temporally thick. They rely on long-term cultivation, delayed fulfillment, and the gradual formation of character and understanding. Ritual repetition, liturgical cycles, and intergenerational transmission are not accidental features of religion; they are pedagogical structures that shape attention, memory, and ethical orientation over time.

When religious practice is reorganized around algorithmic temporalities, this pedagogical function is weakened. The slow cultivation of meaning is replaced by episodic consumption. Rituals risk becoming content items, detached from their formative rhythm and embedded instead within on-demand digital schedules. Reflection is fragmented into short intervals, and the experience of waiting—so central to many religious traditions—is displaced by the expectation of instant access. The future, rather than being anticipated as a space of promise or judgment, is increasingly managed through predictive analytics and optimization. This temporal compression undermines the capacity of religious rationality to sustain long-term existential orientation, replacing endurance with immediacy and formation with responsiveness.

Algorithmic mediation also reshapes the structure of authority in ways that further marginalize ultimate questions. Traditionally, religious authority has been embedded in institutions, sacred texts, and communal practices that allow for interpretation, contestation, and accountability. Even when such authority is hierarchical, it remains relational: teachers, clergy, and communities are identifiable agents who can be questioned, corrected, or held responsible. Algorithmic authority, by contrast, is diffuse, impersonal, and opaque. Its criteria of validity are rarely transparent, and its decisions are justified by performance rather than by explicit normative reasoning. When algorithmic systems increasingly shape what religious actors encounter—what texts are recommended, which interpretations are amplified, which communities become visible—they exert a subtle but powerful influence on theological imagination and ethical orientation. This influence is indirect, operating through selection and prioritization rather than through explicit instruction. Yet it is precisely this indirectness that makes algorithmic authority difficult to contest. Users are guided without being addressed, shaped without being persuaded. Responsibility for interpretive outcomes is dispersed across technical systems that cannot be held accountable in moral or existential terms. This transformation has significant implications for the status of ultimate questions. Religious authority has traditionally functioned not only to transmit answers, but to preserve spaces in which ultimate questions can be asked, sustained, and revisited. Algorithmic mediation, however, tends to favor answers over questions, solutions over inquiry, and clarity over ambiguity. Questions that cannot be resolved through actionable guidance or immediate reassurance are less likely to be surfaced or sustained. As a result, ultimate questions

risk being displaced by derivative questions—how to be happier, more productive, or more resilient—questions that are compatible with optimization and measurement.

Moreover, algorithmic mediation reshapes the epistemic environment in which plausibility is formed. What appears reasonable, relevant, or worthy of attention is increasingly determined by algorithmic filtering rather than by communal deliberation or tradition-based discernment. Ultimate questions, which often lack immediate social reinforcement and measurable outcomes, struggle to maintain plausibility within such environments. They come to appear abstract, impractical, or excessively demanding. Over time, this erosion of plausibility does not require explicit secularization; it operates through habituation. Individuals become accustomed to forms of reasoning that prioritize actionable information and immediate relevance, making it increasingly difficult to recognize the legitimacy of questions that exceed such criteria.

The marginalization of ultimate questions through algorithmic mediation thus represents a structural transformation rather than a cultural preference. It is not simply that contemporary individuals are less interested in meaning or transcendence; it is that the environments in which attention, time, and authority are organized systematically disadvantage the forms of questioning upon which religious rationality depends. This transformation is particularly significant because ultimate questions are not optional additions to religious life. They are constitutive of religion's capacity to orient existence, articulate responsibility, and sustain hope beyond calculable outcomes. In this sense, algorithmic mediation does not refute religious rationality, but it displaces it. It relocates meaning from the horizon of ultimacy to the domain of functionality, from existential orientation to behavioral management. Religion survives, but in a diminished form—no longer as a site where the deepest questions of existence are held open, but as a reservoir of techniques adaptable to the goals of optimization and self-regulation. Recognizing this displacement is essential for any serious engagement between religion and artificial intelligence. The challenge is not merely to preserve religious content within digital environments, but to safeguard the conditions under which ultimate questions can be asked at all. This requires more than ethical guidelines or institutional adaptation. It demands a critical reexamination of how algorithmic rationality structures attention, temporality, and authority, and how these structures can be resisted, reoriented, or complemented by forms of practice that sustain the irreducibility of ultimate concern.

## 5. Functionalization and the Risk of Religious Reduction

In highly computational societies, religion is increasingly evaluated according to the functions it appears to perform within broader social systems. These functions typically include the promotion of psychological well-being, the reinforcement of social cohesion, the regulation of ethical behavior, and the enhancement of individual resilience in the face of uncertainty. From the perspective of policy-making, organizational management, and technological governance, such evaluations are often framed as pragmatic and value-neutral. Religion is considered beneficial insofar as it contributes to measurable outcomes: reduced anxiety, increased compliance with social norms, improved cooperation, or enhanced adaptability to social and economic pressures.

While these functions are not trivial and have long been recognized within sociological analyses of religion, reducing religion to them entails a profound transformation in its meaning and rationality. The problem is not that religion has social or psychological functions. Historically, religious traditions have always shaped moral conduct, communal bonds, and individual coping strategies. The problem arises when these functions are treated not as secondary effects of religious life, but as its primary justification. In such cases, religion is no longer understood as a mode of orientation toward ultimate concern, but as an instrument whose value is assessed in terms of system performance. Algorithmic rationality powerfully reinforces this shift by privileging outcomes that can be quantified, predicted, and optimized. Within algorithmic frameworks, what matters is not whether a belief is true or meaningful in an ultimate sense, but whether it produces desirable behavioral or affective effects that can be monitored and scaled.

From this perspective, religion becomes valuable insofar as it contributes to system stability. It is appreciated for its capacity to reduce mental health burdens, foster prosocial behavior, and cultivate emotional resilience in populations exposed to rapid technological and economic change. Practices such as meditation, prayer, or communal worship are reframed as techniques for stress management and emotional regulation. Ethical teachings are interpreted as mechanisms for encouraging cooperation and reducing conflict. Even transcendence itself is subtly redefined as a form of personal fulfillment or subjective meaning-making that enhances individual well-being. Ultimate concern, in this context, no longer names a horizon that relativizes all instrumental goals; it becomes a resource mobilized in the service of those goals. This functionalization aligns religion closely with contemporary wellness culture and the ideology of self-optimization. Within such frameworks, the good life is increasingly defined in terms of balance, efficiency, and psychological fitness. Religion is integrated as one among many tools—alongside mindfulness apps, coaching programs, and therapeutic interventions—for managing the self. While this integration may appear inclusive and progressive, it carries a significant cost. Religion's critical distance from prevailing social rationalities is eroded. When religious practices are valued primarily for their capacity to enhance productivity or resilience, they are implicitly required to affirm the very conditions that generate stress, alienation, and inequality. The question shifts from whether existing systems are just or meaningful to how individuals can better adapt to them.

Algorithmic rationality intensifies this process by embedding religion within systems of governance that operate through data-driven assessment and behavioral modulation. In such systems, religion becomes part of a broader apparatus of population management. Religious engagement can be tracked, correlated with behavioral outcomes, and incorporated into predictive models of social risk and stability. From this standpoint, religion functions as a variable within governance strategies aimed at optimizing social order. Its success is measured not by its capacity to articulate truth or cultivate transcendence, but by its contribution to reducing volatility and enhancing compliance. The language of ultimate meaning is thus displaced by the language of risk management and optimization. This transformation carries profound implications for religious rationality. Historically, religious traditions have not only stabilized societies; they have also disrupted them. They have articulated visions of justice that exceed existing social arrangements and challenged dominant rationalities by appealing to norms and meanings that

cannot be derived from prevailing systems. Prophetic critique, eschatological hope, and ethical demands grounded in transcendence have enabled religion to question political power, economic exploitation, and cultural conformity. This critical function depends precisely on religion's refusal to be justified solely in functional terms. Its authority derives from a claim to ultimacy that relativizes all instrumental goals. When religion is fully absorbed into algorithmic governance as a tool for managing populations or enhancing productivity, this critical potential is neutralized. Religious symbols and narratives may continue to circulate, but their disruptive force is domesticated. Hope is reframed as optimism, justice as social harmony, and transcendence as inner peace. The capacity of religion to name suffering as unjust rather than merely unfortunate is weakened. Structural injustices are psychologized, and moral outrage is translated into coping strategies. In this way, functionalization does not eliminate religion; it transforms it into a supportive technology of adaptation.

The risk of religious reduction becomes especially acute when functionalization is coupled with personalization. Algorithmic systems excel at tailoring content to individual preferences and emotional states. Applied to religion, this personalization encourages an individualized, consumer-oriented spirituality detached from collective obligations and historical memory. Religious practices are selected based on personal benefit rather than shared commitment. Ultimate concern is fragmented into a series of individualized meaning projects, each optimized for subjective satisfaction. The communal and normative dimensions of religion—its capacity to bind individuals to obligations that transcend personal preference—are thereby weakened.

Moreover, functionalization obscures the normative question of whether the systems being stabilized are themselves worthy of preservation. Algorithmic rationality tends to treat existing social arrangements as given constraints within which optimization occurs. Religion, when reduced to function, is recruited to help individuals adapt to these constraints rather than to interrogate them. This stands in sharp contrast to religious rationality understood as orientation toward ultimate concern, which inherently raises questions about the legitimacy of any social order that absolutizes efficiency, growth, or control.

It is therefore crucial to distinguish between acknowledging the functional effects of religion and allowing those effects to define its essence. A religious rationality reduced to function ceases to be a source of ultimate critique. It becomes one subsystem among others, evaluated according to performance indicators rather than existential truth. Such a reduction may render religion more compatible with algorithmic governance, but it does so at the cost of its capacity to challenge the assumptions upon which that governance rests. In this sense, the functionalization of religion represents not merely a sociological trend, but a philosophical risk. It reflects a broader tendency of algorithmic rationality to absorb all forms of meaning into frameworks of optimization and control. Against this tendency, religious rationality insists that not all that matters can be measured, and not all that stabilizes is just. Its value lies precisely in its refusal to be fully instrumentalized. Preserving this refusal is essential if religion is to remain a meaningful interlocutor in highly computational societies. The challenge, then, is not to deny that religion has functions, but to resist the collapse of religion into function. Doing so requires a renewed articulation of religious rationality that affirms its orientation toward ultimate concern as

irreducible to system utility. Only on this basis can religion retain its critical, transcendent, and normatively disruptive role in a world increasingly governed by algorithmic reason.

## 6. Reconfiguring Religious Rationality in the Age of AI

The challenge posed by algorithmic rationality does not require religion to reject technology or retreat into anti-modern isolation. Instead, it calls for a rearticulation of religious rationality that clarifies its distinctive contribution within a computationally saturated world. This paper proposes three interrelated features of religious rationality in the age of AI. First, religious rationality must affirm its non-optimizable character. Ultimate concern cannot be reduced to metrics without ceasing to be ultimate. Religion must therefore resist pressures to justify itself solely in terms of efficiency or utility. Second, religious rationality must cultivate reflexivity. In an environment where algorithmic systems increasingly shape perception and judgment, religion can function as a site of critical reflection on the limits of calculation. By preserving practices of silence, confession, and lament, religion sustains spaces where the human encounter with contingency and transcendence remains possible. Third, religious rationality must emphasize responsibility. Algorithmic rationality tends to diffuse responsibility by embedding decisions within systems. Religion, by contrast, insists on personal and collective accountability before an ultimate horizon of meaning. This insistence is not incompatible with technology, but it challenges the tendency to outsource moral judgment to automated processes.

## 7. Conclusion

Artificial intelligence confronts religion not by disproving its doctrines, but by reshaping the rational conditions under which ultimate concern can be articulated. Algorithmic rationality, with its emphasis on datafication, prediction, and optimization, risks marginalizing questions that cannot be rendered computationally tractable. In response, religion must clarify its own rationality—not as an alternative technology, but as a mode of meaning that resists full incorporation into algorithmic governance. By reasserting the irreducibility of ultimate concern, religious rationality can remain a vital interlocutor in highly computational societies. Its task is not to compete with algorithms in efficiency, but to remind humanity of what cannot be optimized: responsibility, transcendence, and the fragile dignity of finite existence.

### Author Contributions:

Conceptualization, W. X; methodology, W. X; software, W. X; validation, W. X; formal analysis, W. X; investigation, W. X; resources, W. X; data curation, W. X; writing—original draft preparation, W. X; writing—review and editing, W. X; visualization, W. X; supervision, W. X; project administration, W. X; funding acquisition, W. X. All authors have read and agreed to the published version of the manuscript.

### **Funding:**

This research received no external funding.

### **Institutional Review Board Statement:**

Not applicable.

### **Informed Consent Statement:**

Not applicable.

### **Data Availability Statement:**

Not applicable.

### **Conflict of Interest:**

The author declares no conflict of interest.

### **References**

- Ahmed, S., Sumi, A. A., & Aziz, N. A. (2025). Exploring multi-religious perspectives on artificial intelligence. *Theology and Science*, 23(1), 104–128.
- Horkheimer, M., & Adorno, T. W. (2002). *Dialectic of enlightenment*. Stanford University Press.
- Tampubolon, M., & Nadeak, B. (2024). Artificial intelligence and the understanding of religion: A moral perspective. *International Journal of Multicultural and Multireligious Understanding*, 11(8), 903–914.
- Tillich, P. (1957). *Dynamics of faith*. Harper & Row.
- Umbrello, S. (2023). Artificial intelligence, responsibility, and the common good. *Religions*, 14(12), 1536.

### **License:** Copyright (c) 2025 Author.

All articles published in this journal are licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited. Authors retain copyright of their work, and readers are free to copy, share, adapt, and build upon the material for any purpose, including commercial use, as long as appropriate attribution is given.

# Can Artificial Intelligence “Understand” Religion? A Philosophical Inquiry into Meaning, Symbol, and Transcendence

Othman Al-Farsi<sup>1,\*</sup>

<sup>1</sup> Addis Ababa University, Algeria Av 1150, Ethiopia

**\* Correspondence:**

Othman Al-Farsi

othman.aau@gmail.com

*Received: 10 December 2025/ Accepted: 26 December 2025/ Published online: 31 December 2025*

## Abstract

Recent advances in artificial intelligence, particularly in large language models, have intensified debates about machine understanding, interpretation, and meaning. In religious contexts, these debates acquire particular philosophical urgency: can artificial intelligence meaningfully “understand” religious texts, rituals, and symbols, or does its competence remain confined to formal linguistic and statistical operations? Drawing on philosophical hermeneutics and the philosophy of religious language, this article argues that AI’s apparent interpretive capacities do not amount to genuine religious understanding. Religious meaning is not reducible to semantic coherence or predictive accuracy; it presupposes existential involvement, symbolic participation, and openness to transcendence. By examining religious texts, symbolic language, and ritual practices, the paper delineates the cognitive and ontological boundaries of AI with respect to religion. It concludes that while AI can function as a powerful auxiliary tool in religious studies, it cannot replace the interpretive subject nor access the dimension of meaning constitutive of religious understanding.

**Keywords:** Artificial Intelligence; Religious Understanding; Hermeneutics; Religious Language; Symbol; Transcendence

## 1. The Question of “Understanding”: Why Religion Poses a Special Problem for AI

The question of whether artificial intelligence can “understand” religion presupposes a prior philosophical clarification of what understanding itself entails. In contemporary discussions of artificial intelligence, especially in popular and applied contexts, understanding is frequently treated as a functional achievement. If a system can generate coherent responses, explain concepts in contextually appropriate ways, and adapt its outputs to new inputs, it is often said to “understand” the domain in question. This functionalist assumption is reinforced by the

impressive linguistic fluency of recent AI systems, which appear to demonstrate comprehension across a wide range of topics, including theology and religious studies. However, such usage conceals a deep conceptual ambiguity. It conflates performance with understanding, and in doing so obscures a distinction that becomes especially consequential when the object of understanding is religion. From a philosophical perspective, understanding has never been reducible to the ability to produce correct or plausible outputs. In the human sciences, understanding refers to a mode of grasping meaning that is situated, interpretive, and normatively laden (Dilthey, 1989; Gadamer, 2004). It involves not merely knowing that something is the case, but recognizing what it means within a broader horizon of significance. This horizon is shaped by history, language, embodiment, and forms of life. When contemporary AI discourse treats understanding as equivalent to functional competence, it implicitly adopts a thin, operational definition that abstracts understanding from these conditions. While such abstraction may be pragmatically useful in engineering contexts, it becomes philosophically inadequate when applied to religion.

Religion does not merely transmit information about the world. It articulates a way of inhabiting the world. Religious texts, rituals, and symbols are embedded in practices that orient individuals and communities toward fundamental questions of existence: finitude and mortality, guilt and responsibility, suffering and hope, meaning and transcendence. To understand religion, therefore, is not simply to grasp doctrinal propositions or historical facts. It is to apprehend how meaning is disclosed, lived, negotiated, and sometimes contested within a symbolic horizon that claims existential seriousness. Religious understanding involves recognizing how beliefs and practices shape a form of life, not merely how they can be described from an external standpoint. This immediately raises doubts about whether computational systems—however sophisticated—can access the kind of understanding religion presupposes. Artificial intelligence operates through formal operations on representations: it identifies patterns, correlates linguistic expressions, and generates outputs based on probabilistic models. These capacities enable AI to handle religious material at the level of syntax and semantic association. Yet religious understanding is not exhausted by such operations. It involves a relation to meaning that is inseparable from concern, commitment, and vulnerability. A system that does not care, cannot be addressed, and cannot be transformed by what it processes stands at a fundamental distance from this mode of understanding. The temptation to ascribe understanding to AI arises in part from the remarkable fluency of contemporary language models. These systems can translate sacred texts with high accuracy, summarize complex theological debates, generate sermons or prayers, and respond convincingly to questions about religious doctrines. From a surface perspective, this fluency appears indistinguishable from human understanding. However, as philosophers from Searle (1980) to Dreyfus (1992) have argued, linguistic competence alone does not suffice for understanding. Searle's Chinese Room argument famously illustrates that the manipulation of symbols according to formal rules does not entail comprehension of their meaning. Dreyfus, drawing on phenomenology, further emphasizes that human understanding is grounded in embodied, situated engagement with the world—an engagement that cannot be replicated by abstract symbol processing. Religion intensifies this problem because its meaning is not primarily propositional. While religious traditions certainly include doctrinal statements, their significance cannot be reduced to factual correctness. Religious language frequently operates symbolically,

metaphorically, and performatively. It aims not only to describe reality, but to disclose it in a particular light and to orient action and self-understanding accordingly. Understanding such language requires sensitivity to its normative and existential dimensions, not merely its semantic structure. AI systems, by contrast, treat religious language as data: strings of text to be processed, classified, and reproduced. The difference between these modes of engagement is not one of degree, but of kind.

Moreover, religious understanding is intrinsically normative. Religious claims do not merely assert that something is the case; they make demands. They call for trust, obedience, repentance, hope, or ethical transformation. To understand a religious claim is therefore to recognize its claim upon the self. As Tillich (1957) emphasizes, religion concerns “ultimate concern”—that which claims unconditional seriousness and orients the totality of existence. Ultimate concern is not a detachable belief-content; it is a lived orientation that structures priorities, commitments, and values. Understanding such concern requires acknowledging its normative force, not merely being able to restate it. This normative dimension further complicates the application of the concept of understanding to artificial systems. AI systems do not stand under obligations, experience guilt, or respond to calls for transformation. They can describe ethical demands, but they cannot be bound by them. They can generate language about hope or repentance, but they cannot hope or repent. This asymmetry matters because religious understanding is inseparable from the capacity to be claimed by meaning. A system that remains normatively indifferent to its outputs lacks a central condition of religious understanding.

From this perspective, the question of AI and religious understanding is not primarily empirical or technical. It is not a matter of whether future systems will be more powerful, more data-rich, or more context-sensitive. Rather, it is a conceptual and philosophical question concerning the conditions under which understanding is possible at all. If understanding presupposes being situated within a form of life, oriented by concerns, embedded in practices, and open to transformation, then the applicability of the concept to artificial systems must be critically examined. Hermeneutic philosophy has long insisted that understanding is inseparable from historicity, finitude, and self-involvement (Heidegger, 1962; Gadamer, 2004). These are not features that can be added to computational systems through technical refinement; they are constitutive of human existence. Religion thus functions as a limit case for artificial intelligence. It exposes the inadequacy of performance-based definitions of understanding by foregrounding dimensions of meaning that resist formalization. In religious contexts, understanding cannot be detached from embodiment, normativity, and transcendence without losing its essence. The apparent success of AI in handling religious language therefore risks generating a category mistake: mistaking simulation for participation, and fluency for understanding. This chapter has argued that religion poses a special problem for AI because it brings into sharp relief the difference between processing meaning and being oriented by it. By foregrounding ultimate concern, symbol, and existential commitment, religion forces a reconsideration of what it means to understand at all. The remainder of this paper develops the argument that genuine religious understanding cannot be reduced to computational operations, and that recognizing this limit is

essential not only for assessing AI, but for clarifying the nature of religion itself in a technological age.

## 2. Hermeneutic Understanding and the Irreducibility of Situated Meaning

Philosophical hermeneutics offers one of the most sustained and systematic critiques of reductive accounts of understanding. Against the assumption that understanding consists in the correct manipulation of representations or the successful execution of cognitive procedures, hermeneutics insists that understanding is a fundamentally situated phenomenon. From Dilthey through Heidegger to Gadamer and Ricoeur, the hermeneutic tradition emphasizes that understanding is historically embedded, linguistically mediated, and existentially implicated (Dilthey, 1989; Heidegger, 1962; Gadamer, 2004; Ricoeur, 1976). This tradition therefore provides a crucial conceptual framework for evaluating contemporary claims about artificial intelligence and understanding, particularly in religious contexts where meaning is inseparable from existential orientation.

### 2.1. Understanding as Lived and Historical: Dilthey and the Limits of Explanation

Wilhelm Dilthey's foundational contribution to hermeneutics lies in his distinction between explanation (*Erklären*) and understanding (*Verstehen*). While explanation seeks causal regularities and law-like relations characteristic of the natural sciences, understanding aims at grasping meaning as it is expressed in human life, history, and culture (Dilthey, 1989). Human expressions—texts, actions, institutions—cannot be adequately understood through causal explanation alone, because their significance is rooted in lived experience (*Erlebnis*). To understand a human expression is to situate it within the horizon of life from which it emerges. This distinction is highly relevant to contemporary debates about AI. Artificial intelligence excels at explanatory tasks in Dilthey's sense: it identifies patterns, correlations, and statistical regularities across vast datasets. When applied to religious texts, AI can explain linguistic usage, trace thematic developments, and classify doctrinal positions. Yet such explanatory power does not amount to understanding in the hermeneutic sense. Religious expressions are not merely data points; they are articulations of lived meaning within specific historical and cultural contexts. Understanding them requires sensitivity to how they respond to concrete existential conditions—suffering, injustice, mortality, hope—conditions that AI does not experience.

Dilthey emphasizes that understanding is inseparable from historicity. Human beings are historical beings who interpret the world from within traditions that shape perception and judgment. Religious meaning, in particular, is transmitted and transformed across generations through practices, narratives, and symbols. AI systems, by contrast, lack historical existence. Although they are trained on historical data, they do not belong to history. They do not inherit traditions, nor do they stand within a temporal continuity that gives meaning to interpretation. Their relation to history is external and archival, not lived and participatory. This limitation reveals a fundamental asymmetry between AI and human interpreters. While AI can model historical patterns, it does not experience history as a dimension of self-understanding. As a result, its “interpretations” of religious material remain detached from the lived horizons that confer

meaning upon that material. Dilthey's distinction thus already suggests that AI may explain religion, but it cannot understand it in the sense appropriate to the human sciences.

## **2.2. Understanding as Existential Projection: Heidegger's Ontological Turn**

Martin Heidegger radicalizes Dilthey's insight by rejecting the idea that understanding is a cognitive act performed by a subject upon an object. In *Being and Time*, Heidegger argues that understanding is an existential structure of *Dasein* itself (Heidegger, 1962). To be human is already to understand—to project oneself onto possibilities of being and to interpret the world in light of concerns, commitments, and finitude. Understanding is therefore not something we occasionally do; it is a way in which we exist. This ontological account has far-reaching implications for the question of AI understanding. Heidegger's notion of understanding is inseparable from being-in-the-world. To understand is to find oneself already involved in a meaningful world where things matter, where possibilities are disclosed, and where one's own being is at stake. This involvement is structured by care (Sorge), temporality, and mortality. Human understanding is always oriented toward the future as a horizon of possibilities shaped by finitude.

Artificial intelligence lacks this existential structure. AI systems do not project possibilities of being, nor do they interpret themselves in light of their own finitude or mortality. Their operations are not oriented by care or concern, but by optimization criteria externally imposed by designers and users. While AI can generate outputs that resemble interpretive judgments, these outputs are not grounded in existential self-relation. They are operations performed on representations, not disclosures of a meaningful world. Heidegger famously argues that only beings who have a world can understand. A "world" in this sense is not a collection of objects, but a meaningful totality in which entities show up as relevant, significant, or negligible. Tools, for example, are encountered not as neutral objects but as "ready-to-hand" within practical contexts. Religious understanding presupposes such worldhood in an intensified form, as religious meanings orient the whole of existence toward ultimate concerns.

AI systems, however, do not have a world in Heidegger's sense. They operate within environments, but these environments are not meaningful horizons disclosed through care and involvement. Consequently, AI lacks the ontological conditions for understanding as Heidegger conceives it. This is not a limitation that can be remedied through improved algorithms or larger datasets; it is a structural consequence of what AI is.

## **2.3. Understanding as Dialogue and Transformation: Gadamer, Ricoeur, and Religious Meaning**

Hans-Georg Gadamer further develops Heidegger's insights by emphasizing the dialogical and transformative character of understanding. In *Truth and Method*, Gadamer rejects the idea that understanding is a methodological procedure aimed at reconstructing an objective meaning. Instead, understanding is a dialogical event in which the interpreter and the text encounter one another within a shared linguistic horizon (Gadamer, 2004). This encounter involves what Gadamer calls a "fusion of horizons," in which the interpreter's preconceptions are challenged, revised, and sometimes overturned. A crucial implication of this view is that understanding

involves risk—the risk of being transformed by what one seeks to understand. Genuine understanding is not neutral or detached; it implicates the interpreter's own self-understanding. This dimension of vulnerability and openness is essential to hermeneutic understanding. One understands only insofar as one is willing to be addressed and affected. This insight is particularly salient in religious contexts. Sacred texts and traditions do not merely convey information; they address interpreters as morally and existentially responsible agents. They call into question existing assumptions and demand orientation toward what is taken to be ultimate. Understanding such texts is inseparable from the possibility of transformation—ethical, existential, or spiritual.

Paul Ricoeur deepens this analysis through his notion of the “surplus of meaning” (surplus de sens) (Ricoeur, 1976). Religious texts, Ricoeur argues, generate meanings that exceed the intentions of their authors and the capacities of any single interpretive framework. This surplus is not a defect to be eliminated, but a productive excess that sustains ongoing interpretation. Understanding, therefore, is never complete or final; it remains open-ended and contested. From this perspective, the limitations of AI become even clearer. AI systems do not engage in dialogue in the hermeneutic sense. They do not bring preconceptions that can be challenged, nor can they be transformed by the texts they process. Their outputs do not involve self-implication or risk. As a result, AI can at best simulate the forms of interpretation without participating in the event of understanding.

Religious understanding exemplifies this hermeneutic structure in an intensified form. It presupposes belonging to a tradition, participation in communal practices, and openness to transcendence. These conditions cannot be formalized or encoded without remainder. Consequently, from a hermeneutic perspective, AI can reproduce interpretive patterns, but it does not stand within a horizon of meaning. It lacks historicity, finitude, dialogical openness, and existential stake. These absences are not contingent technical limitations to be overcome by future innovation. They are constitutive of artificial intelligence as a non-existential system. Hermeneutics thus provides a powerful philosophical framework for articulating why AI cannot understand religion in the strong sense, and why claims to the contrary rest on a reduction of understanding to performance.

### **3. Religious Language and Symbol: Meaning Beyond Semantic Processing**

If hermeneutics reveals that understanding is irreducibly situated and existentially implicated, the philosophy of religious language further clarifies why religious meaning cannot be reduced to semantic processing. Religious language does not function primarily as descriptive discourse about the world; rather, it operates symbolically, metaphorically, and performatively, mediating a relation to transcendence that resists literalization and formalization. This chapter argues that religious language constitutes a decisive boundary for artificial intelligence, because its meaning cannot be exhausted by syntactic manipulation, semantic association, or probabilistic inference.

#### **3.1. Religious Language as Non-Descriptive and World-Disclosing**

A central insight of twentieth-century philosophy of religion is that religious language does not function in the same way as scientific or everyday descriptive language. Whereas descriptive

language aims to represent states of affairs that can be verified or falsified, religious language aims to disclose a world—a meaningful horizon within which existence is interpreted (Heidegger, 1962; Ricoeur, 1976). Statements such as “God is merciful,” “the Kingdom of God is near,” or “nirvana is beyond suffering” are not empirical reports. Their meaning cannot be assessed by correspondence to observable facts, but by their capacity to orient life, perception, and action. This insight undermines any attempt to equate religious understanding with semantic accuracy. Artificial intelligence excels at processing descriptive language because such language lends itself to formal representation and pattern recognition. However, when religious language is treated as a set of propositions to be parsed and recombined, its world-disclosing function is obscured. The meaning of religious language lies not only in what it says, but in how it reconfigures the interpreter’s relation to reality. Ricoeur emphasizes that religious discourse belongs to what he calls “limit-expressions”—forms of language that strain ordinary semantics in order to articulate experiences at the boundaries of meaning (Ricoeur, 1976). Such expressions cannot be paraphrased without loss, because their meaning emerges through tension, ambiguity, and interpretive openness. AI systems, by contrast, are optimized for semantic closure: they aim to resolve ambiguity into plausible continuations. In doing so, they risk neutralizing precisely the excess of meaning that constitutes religious language as religious.

### **3.2. Symbol and Participation: Tillich’s Ontology of Religious Meaning**

Paul Tillich’s theory of religious symbols provides one of the most influential accounts of why religious meaning cannot be reduced to representation. For Tillich, religious symbols do not merely stand for something else; they participate in the reality to which they point (Tillich, 1957). A genuine religious symbol opens a dimension of reality otherwise inaccessible and engages the whole person—intellectually, emotionally, and existentially. When a symbol ceases to participate in this way, it loses its power and becomes empty or idolatrous. This participatory ontology of symbols poses a decisive challenge for artificial intelligence. AI systems treat symbols as tokens within a representational system. They can identify symbolic associations, trace their historical usage, and generate coherent interpretations (Brey, 2020). Yet participation is not something that can be simulated through representation. To participate in a symbol is to be affected by it, to have one’s self-understanding reoriented through it. This presupposes vulnerability, concern, and existential exposure—conditions absent in artificial systems.

Moreover, Tillich insists that religious symbols emerge from and respond to ultimate concern. They are not arbitrarily constructed signs, but expressions of a community’s confrontation with finitude and meaning. AI, lacking ultimate concern, cannot stand in the existential relation that gives rise to symbolic meaning. It can reproduce symbolic language, but it cannot encounter symbols as symbols in Tillich’s sense. The distinction here is not epistemic but ontological: AI processes symbols without being claimed by what they disclose.

### **3.3. Metaphor, Paradox, and Negative Theology: Resistance to Formalization**

Religious language frequently employs metaphor, paradox, and negation to gesture toward transcendence. From biblical poetry to Buddhist koans and apophatic theology, religious traditions have insisted that ultimate reality cannot be captured by literal or univocal language.

Negative theology, in particular, asserts that what is ultimate can only be spoken of by way of negation—God is not this, not that (Pseudo-Dionysius; Marion, 1991). Such language resists semantic stabilization. Its function is not to convey information, but to interrupt ordinary modes of understanding and open space for contemplation. Jean-Luc Marion argues that saturated phenomena—such as revelation—exceed the capacity of conceptual grasp and resist reduction to objectifying knowledge (Marion, 1991). Understanding here involves receptivity rather than mastery.

AI systems, however, are oriented toward mastery in the form of completion, coherence, and plausibility. When confronted with paradoxical or apophatic language, they tend to normalize it—translating negation into affirmation, paradox into explanation. This tendency reflects a structural limitation: AI operates within a framework in which meaning must be rendered computationally tractable. Religious language, by contrast, often derives its meaning precisely from its resistance to such tractability.

This difference is especially evident in traditions that deliberately undermine conceptual understanding as a spiritual discipline. Zen koans, for example, are designed not to be solved, but to destabilize habitual patterns of thought. An AI may generate interpretations of koans, but such interpretations miss the point: the koan's meaning lies in its transformative effect, not in its semantic resolution. Here again, AI encounters a boundary that cannot be crossed through increased computational sophistication.

### 3.4. Language-Games, Forms of Life, and the Problem of Religious Use

Ludwig Wittgenstein's later philosophy offers a further lens through which to assess the limits of AI in relation to religious language. For Wittgenstein, meaning is not a matter of reference, but of use within a form of life (Wittgenstein, 1953). Religious language constitutes a distinctive language-game governed by practices, commitments, and forms of response that cannot be understood in isolation from the life in which they are embedded. To understand religious language, on this view, is not merely to know what words mean, but to know how they function within lived practices—prayer, worship, repentance, ethical decision-making. These practices are inseparable from embodiment, community, and normativity. AI systems, however, do not participate in forms of life. They do not pray, worship, or repent. They do not respond to religious language with obedience, trust, or resistance. As a result, they lack the practical criteria that constitute understanding within a language-game.

This insight clarifies why AI's success at generating religious language can be misleading. Fluency does not imply belonging. A system may produce grammatically correct and contextually appropriate religious utterances without understanding what it is to use such language meaningfully. The distinction is analogous to that between quoting a prayer and praying. Religious understanding involves the latter, not merely the former. Taken together, these considerations reinforce the central claim of this chapter: religious language and symbol operate in dimensions of meaning that exceed semantic processing. They presuppose participation, embodiment, normativity, and openness to transcendence. Artificial intelligence, constrained to representational manipulation, cannot access these dimensions. It can analyze, reproduce, and

simulate religious language, but it cannot inhabit the symbolic world in which that language is meaningful.

This conclusion does not diminish the value of AI for the study of religion. It does, however, set clear philosophical boundaries. Recognizing these boundaries is essential if we are to avoid conflating linguistic performance with understanding and if we are to preserve the integrity of religious meaning in an age of increasingly powerful computational systems.

#### **4. Ritual, Embodiment, and the Non-Propositional Dimension of Religious Meaning**

If the previous chapter demonstrated that religious language and symbolism exceed semantic processing, religious ritual reveals an even deeper limit to artificial intelligence: the irreducibly embodied, performative, and practical dimension of religious meaning. Rituals do not merely express beliefs; they enact and sustain forms of life. Understanding religion, therefore, requires more than linguistic or symbolic competence—it presupposes participation in embodied practices that shape perception, disposition, and communal identity. This chapter argues that religious ritual constitutes a decisive boundary for AI because ritual meaning cannot be accessed from a detached, representational standpoint.

##### **4.1. Ritual as Practice Rather Than Representation**

One of the most influential insights in contemporary ritual theory is that rituals are not primarily representational acts. They do not function mainly to communicate information or symbolize doctrinal content. Rather, rituals are practices that do something to participants: they discipline bodies, structure time, and cultivate dispositions (Bell, 1997). To understand a ritual is not simply to know what it signifies, but to grasp how it operates within a form of life. This insight challenges any model of understanding based solely on representation. Artificial intelligence approaches ritual as an object of analysis: sequences of actions, symbolic gestures, and textual accompaniments that can be categorized and described. From this external perspective, rituals appear as systems of signs whose meaning can be decoded. Yet such decoding misses the core of ritual meaning. Rituals are meaningful not because they represent something else, but because they form participants through repetition and embodied engagement.

Talal Asad (1993) emphasizes that religious practices cannot be abstracted from the disciplines and power relations that sustain them. Rituals shape sensibilities over time, producing forms of attention, obedience, and ethical responsiveness. Understanding ritual therefore presupposes being shaped by it. AI systems, however, cannot undergo formation. They can analyze ritual patterns, but they cannot be habituated, disciplined, or transformed through practice. Their relation to ritual remains observational rather than participatory.

##### **4.2. Embodiment and the Formation of Religious Meaning**

Religious rituals are inseparable from embodiment. Bodies kneel, fast, chant, bow, process, and remain silent. These bodily actions are not secondary expressions of belief; they are constitutive of religious meaning itself. Phenomenological approaches to religion emphasize that meaning is disclosed through bodily engagement with the world, not merely through conceptual reflection

(Merleau-Ponty, 1962). Embodiment matters because religious understanding often arises pre-reflectively. Through ritual participation, bodies learn rhythms of time, hierarchies of value, and modes of responsiveness that precede explicit belief. For example, the repetition of prayer structures attention and memory; fasting reorganizes desire; pilgrimage reorients spatial perception. These transformations cannot be reduced to propositional knowledge.

Artificial intelligence lacks bodies in this phenomenological sense. While AI systems may be embedded in physical devices, they do not inhabit bodies as sites of vulnerability, sensation, fatigue, or discipline. They do not experience hunger, pain, or exhaustion; nor do they acquire meaning through bodily repetition. As a result, AI cannot access the pre-reflective dimension of religious understanding that emerges through embodied practice. This absence is not merely practical but ontological. Embodied participation situates religious meaning within finitude and vulnerability—conditions essential to religious self-understanding. Rituals often enact precisely what it means to be finite: dependence, gratitude, mourning, repentance. AI systems, lacking finitude in this existential sense, cannot encounter ritual meaning as something that matters for their own being.

#### **4.3. Ritual Temporality, Repetition, and the Formation of Commitment**

Religious rituals are also temporally structured in ways that resist instrumental rationality. They rely on repetition, delay, and cyclical time rather than efficiency and immediacy. Liturgical calendars, fasting seasons, and rites of passage situate individuals within temporal rhythms that exceed individual preference and utility. Through repetition, rituals cultivate patience, endurance, and commitment. Understanding ritual therefore involves inhabiting a particular temporality. Meaning unfolds gradually through repetition, not instantaneously through comprehension. This stands in sharp contrast to algorithmic systems, which operate in real-time feedback loops optimized for speed and adaptability. AI systems process information episodically; they do not dwell in time or accumulate meaning through duration. Moreover, ritual repetition is not redundant. Each enactment reaffirms belonging and renews commitment. The meaning of ritual is inseparable from its recurrence. To understand a ritual is to recognize why it must be repeated even when its structure is already known. AI, oriented toward novelty and optimization, lacks a rationale for repetition without informational gain. This reveals a fundamental mismatch between ritual temporality and computational logic. Rituals also bind participants to communal histories and futures. They inscribe individuals into narratives that precede and outlast them. Understanding ritual thus involves recognizing oneself as part of a temporal continuum. AI systems, however, do not possess temporal self-continuity in this sense. Their “memory” is functional storage, not lived history.

#### **4.4. Community, Authority, and the Normative Force of Ritual Practice**

Religious ritual is inseparable from communal authority and normative accountability. Rituals are not private performances; they are socially regulated practices embedded in traditions. Participation entails submission to shared norms and recognition by others. Understanding ritual therefore includes knowing how one is answerable within a community. This normative dimension further distinguishes religious understanding from computational processing. AI

systems do not stand under communal authority. They are not accountable for correct or incorrect ritual performance in a moral or religious sense. They cannot fail ritually, repent, or be reconciled. Consequently, they cannot grasp the normative force that rituals exert over participants. Ritual authority also shapes interpretation. The meaning of a ritual is not decided solely by individual intention or analytical description, but by tradition-guided practice. Understanding ritual requires sensitivity to this authority structure. AI systems, however, operate without allegiance or submission. They can describe authority, but they cannot recognize it as binding. Taken together, these considerations show that religious ritual embodies a form of meaning inaccessible to artificial intelligence. Ritual meaning is practical rather than representational, embodied rather than abstract, temporal rather than episodic, and normative rather than neutral. AI can analyze rituals, simulate them, or assist in their documentation, but it cannot participate in the form of life through which ritual meaning is constituted. This does not imply that ritual meaning is mysterious or irrational. Rather, it underscores that understanding is inseparable from practice. Religious rituals teach not by explaining, but by forming. Artificial intelligence, lacking embodiment, finitude, and communal accountability, cannot undergo such formation. Ritual therefore marks a decisive boundary of AI's cognitive reach and reinforces the broader claim of this paper: that religious understanding cannot be reduced to computational processes.

## 5. The Limits of AI and the Meaning of Religion in a Technological Age

The preceding analysis leads to a clear conclusion: artificial intelligence cannot understand religion in the strong, hermeneutic sense. This does not diminish AI's value as a tool for religious studies, but it clarifies its limits. Treating AI as a religious interpreter risks reducing religion to information processing and obscuring its existential depth. More importantly, this conclusion sheds light on the nature of religion itself. Religion is not merely a system of beliefs or texts, but a mode of orientation toward meaning, symbol, and transcendence. Understanding religion involves participation, vulnerability, and openness to what exceeds calculation. In an age increasingly shaped by algorithmic rationality, recognizing this limit is ethically and philosophically significant. It reminds us that not all understanding can be automated, and not all meaning can be optimized. Religion, in this sense, functions as a critical counterpoint to technological reductionism, preserving dimensions of human existence that resist computation..

### Author Contributions:

All authors have read and agreed to the published version of the manuscript.

### Funding:

This research received no external funding.

### Institutional Review Board Statement:

Not applicable.

### **Informed Consent Statement:**

Not applicable.

### **Conflict of Interest:**

The author declares no conflict of interest.

### **References**

- Asad, T. (1993). *Genealogies of religion: Discipline and reasons of power in Christianity and Islam*. Johns Hopkins University Press.
- Bell, C. (1997). *Ritual: Perspectives and dimensions*. Oxford University Press.
- Brey, P. (2020). Artificial intelligence and human values. *Ethics and Information Technology*, 22(2), 75–84.
- Dilthey, W. (1989). *Introduction to the human sciences* (R. A. Makkreel & F. Rodi, Trans.). Princeton University Press. (Original work published 1883)
- Dreyfus, H. L. (1992). *What computers still can't do: A critique of artificial reason*. MIT Press.
- Gadamer, H.-G. (2004). *Truth and method* (2nd rev. ed., J. Weinsheimer & D. G. Marshall, Trans.). Continuum. (Original work published 1960)
- Heidegger, M. (1962). *Being and time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)
- Marion, J. L. (1991). *God without being* (T. A. Carlson, Trans.). University of Chicago Press.
- Merleau-Ponty, M. (1962). *Phenomenology of perception* (C. Smith, Trans.). Routledge & Kegan Paul. (Original work published 1945)
- Ricoeur, P. (1976). *Interpretation theory: Discourse and the surplus of meaning*. Texas Christian University Press.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–457.  
<https://doi.org/10.1017/S0140525X00005756>
- Tillich, P. (1957). *Dynamics of faith*. Harper & Row.
- Wittgenstein, L. (1953). *Philosophical investigations* (G. E. M. Anscombe, Trans.). Blackwell.

#### **License:** Copyright (c) 2025 Author.

All articles published in this journal are licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited. Authors retain copyright of their work, and readers are free to copy, share, adapt, and build upon the material for any purpose, including commercial use, as long as appropriate attribution is given.

# From Sacred Sanctuary to Intelligent System: A Religious-Studies Reflection on the Transformation of Temple Space under Artificial Intelligence

Ouyang Shaokang<sup>1,\*</sup>

<sup>1</sup> National University of Singapore, Singapore 119077, Singapore

**\* Correspondence:**

Ouyang Shaokang

ouyang\_sk1162@163.com

*Received: 10 December 2025/ Accepted: 26 December 2025/ Published online: 31 December 2025*

## Abstract

The increasing integration of artificial intelligence into religious institutions has generated extensive discussion concerning ethics, authority, and belief. However, comparatively little attention has been paid to the transformation of religious space itself. This article examines the intelligent transformation of temple space through the analytical lens of the sacred–profane distinction, arguing that the introduction of artificial intelligence represents not merely a technical modernization but a structural reconfiguration of religious spatiality. Drawing on classical theories of religion, particularly the works of Durkheim and Eliade, as well as contemporary discussions in the sociology of religion and philosophy of technology, the paper conceptualizes temples as symbolic, institutional, and experiential spaces constituted through boundaries, rituals, and normative orders. It then analyzes how algorithmic systems—such as surveillance technologies, data-driven management, and automated governance—reshape temple space by introducing system rationality oriented toward efficiency, visibility, and control. The study argues that this systematization tends to blur the distinction between sacred and profane, transform religious authority, and reconfigure the conditions of religious experience. Rather than rejecting technological mediation outright, the paper calls for a critical religious-studies perspective that recognizes both the adaptive possibilities and the risks of total system integration. It concludes that the future of temples as sacred spaces depends on their capacity to preserve symbolic density, ritual autonomy, and experiential transcendence within increasingly intelligent environments.

**Keywords:** Artificial Intelligence; Temple Space; Sacred–Profane Distinction; Religious Institutions; System Rationality; Religion and Technology

## 1. Introduction

In recent years, artificial intelligence and related digital technologies have increasingly entered religious spaces across different cultural contexts. In Buddhist temples, Christian churches, and other institutional religious sites, technologies such as facial-recognition security, automated donation systems, smart surveillance, digital ritual scheduling, algorithmic visitor management, and AI-based religious services have become progressively normalized. These developments are often justified in terms of efficiency, safety, transparency, and modernization. Yet from the perspective of religious studies, they raise a more fundamental question: what happens to religious space when it is reorganized according to system rationality?

Temples are not merely functional buildings. They are symbolic spaces, ritual centers, and institutional embodiments of the sacred. Classical theories of religion have long emphasized that religious space is structured through a distinction between the sacred and the profane (Durkheim, 1995; Eliade, 1957). The temple, as a sacred space, marks a qualitative rupture within ordinary social space. It is a place where different norms apply, different temporalities operate, and different forms of authority are recognized. Entering a temple traditionally entails a shift in posture, behavior, and meaning-orientation.

The integration of artificial intelligence into temple management challenges this structure in subtle but profound ways. Intelligent systems do not merely assist religious activities; they introduce a different mode of rationality—one oriented toward optimization, prediction, surveillance, and system integration. This raises the central question of this paper: does the intelligent transformation of temples represent a neutral modernization of religious space, or does it fundamentally reconfigure the sacred–profane distinction on which such spaces depend?

Existing scholarship on religion and technology has focused primarily on ethical debates (e.g., AI morality), theological questions (e.g., machine agency), or digital religion (e.g., online rituals). Comparatively little attention has been paid to religious space as space, especially under conditions of algorithmic governance. This paper seeks to fill that gap by offering a theoretical reflection on temple space as it moves “from sanctuary to system.” Methodologically, this study adopts a theoretical and interpretive approach grounded in classical and contemporary religious studies. It does not present ethnographic data but instead examines structural transformations through conceptual analysis. The argument proceeds through six chapters, tracing the transformation of temple space from sacred enclosure to system-managed environment and reflecting on its implications for religious meaning, authority, and experience.

## 2. Sacred and Profane: Classical Theories of Religious Space

### 2.1. Durkheim and the Social Construction of Sacred Space

Émile Durkheim’s foundational insight in *The Elementary Forms of Religious Life* is that religion organizes reality through a fundamental distinction between the sacred and the profane (Durkheim, 1995). Sacred things are those set apart, protected by prohibitions, and invested with collective significance, while profane things belong to the domain of ordinary, utilitarian, and

everyday life. Crucially, this distinction is not merely a matter of belief or doctrine; it is enacted through practices, norms, and spatial arrangements. Religion, for Durkheim, is above all a social fact, and sacred space is one of its most visible materializations. From this perspective, the temple functions as a socially constructed boundary that stabilizes the sacred–profane distinction in spatial form. The demarcation between inside and outside, center and periphery, pure and impure is not accidental but normatively regulated. Behavioral rules—such as requirements of silence, specific dress codes, bodily gestures, ritual purification, and regulated movement—are enforced through spatial design. These rules do not merely express reverence; they actively produce the experience of the sacred by suspending ordinary social habits. Entering a temple thus involves crossing into a space governed by a different normative order.

For Durkheim, the significance of sacred space lies in its social function. Sacred places embody collective representations and sustain social cohesion by materializing shared values. The temple is not sacred because of inherent metaphysical qualities, but because a community collectively recognizes and maintains it as such. Sacredness, therefore, depends on continuous social reinforcement through ritual practice and normative regulation. If these mechanisms weaken or are displaced, the sacred character of space is at risk.

The introduction of intelligent systems complicates this Durkheimian function in subtle but important ways. Surveillance cameras, biometric access controls, automated crowd management, and algorithmic monitoring introduce forms of regulation typically associated with profane institutions such as airports, corporate buildings, or shopping centers. These technologies operate according to principles of efficiency, risk prevention, and visibility rather than collective symbolic meaning. When such systems become normalized within temples, regulation no longer derives primarily from religious norms but from technical governance. This shift risks transforming sacred regulation into a subset of secular administration. The authority governing behavior increasingly appears impersonal and system-driven rather than communal and symbolic. As a result, the temple’s capacity to function as a distinct social space—one that suspends ordinary forms of control in favor of religious normativity—is weakened. From a Durkheimian perspective, this represents a potential erosion of the social mechanisms that sustain sacredness itself.

## 2.2. Eliade and the Phenomenology of the Sacred Place

While Durkheim emphasizes the social construction of sacred space, Mircea Eliade approaches religious space from a phenomenological perspective that foregrounds lived experience and symbolic orientation. For Eliade, sacred space emerges through hierophany—the manifestation of the sacred that interrupts the homogeneity of profane space (Eliade, 1957). In a world that is otherwise spatially continuous and functionally neutral, the appearance of the sacred introduces a qualitative rupture. A sacred place is not simply one location among others; it is a center that orients the world. Within this framework, the temple functions as a “center of the world” (*axis mundi*). It anchors cosmic order, providing orientation not only in physical space but also in existential meaning. Architectural layout, ritual movement, spatial hierarchies, and symbolic objects are not primarily functional; they are cosmological. Through repetitive ritual practices, sacred space continuously re-enacts the structure of the cosmos and situates human existence

within it. The meaning of the temple lies in its capacity to disclose a different order of reality—one that transcends ordinary utility and calculation.

From Eliade's perspective, sacred space resists instrumental rationality. Its value does not lie in efficiency, optimization, or management but in symbolic orientation and repetition. Sacred time and sacred space are cyclical, not linear; they prioritize renewal over progress and meaning over function. Any transformation of temple space must therefore be evaluated in terms of its impact on symbolic orientation rather than operational performance. Intelligent systems, however, reorganize space according to fundamentally different logics. Algorithmic technologies prioritize traffic flow optimization, risk management, data collection, and system integration. Space is treated as an operational resource to be monitored and controlled. Within such frameworks, there is no privileged center—only nodes within a network. Visibility replaces mystery, and predictability replaces symbolic depth. This introduces a structural tension between symbolic centrality and systemic integration. When temples are increasingly governed by intelligent systems, their spatial logic risks being redefined in non-symbolic terms. The temple becomes less a cosmic center and more a managed facility. From an Eliadean standpoint, this transformation threatens the phenomenological conditions under which sacred space can be experienced as sacred. The sacred does not disappear, but it becomes harder to encounter as a rupture within ordinary reality.

### 2.3. Sacred Space as Boundary, Threshold, and Liminal Zone

Despite their methodological differences, both Durkheim and Eliade converge on a crucial insight: sacred space functions as a boundary. It separates the sacred from the profane, but it also mediates between them. The temple is neither entirely detached from everyday life nor fully absorbed into it. It operates as a threshold—a liminal zone in which ordinary categories are suspended and reconfigured. Victor Turner's theory of liminality further develops this insight by emphasizing the experiential dimension of thresholds (Turner, 1969; Asad, 1993). Liminal spaces are transitional zones in which individuals temporarily step outside ordinary social structures and undergo symbolic transformation. In religious contexts, crossing into sacred space entails a change in status, perception, and comportment. The threshold itself—gates, doorways, courtyards, ritual preparations—is often heavily ritualized to mark this transition. These spatial transitions are essential for ritual efficacy. They prepare participants to encounter the sacred by gradually distancing them from profane concerns. The threshold is not merely a point of entry; it is a process of reorientation. Religious experience depends on such staged transitions, which cultivate attentiveness, humility, and receptivity.

When temples are redesigned as intelligent systems, thresholds are increasingly replaced by automated checkpoints. Entry may be regulated by sensors, QR codes, facial recognition, or ticketing systems rather than ritual gestures or symbolic preparation. While such mechanisms improve efficiency and security, they fundamentally alter the meaning of transition. The act of entering the temple becomes administratively processed rather than ritually enacted. This transformation has significant implications for religious experience. Automated access treats entry as a logistical problem rather than a symbolic passage. The liminal quality of the threshold is reduced, and the experiential shift from profane to sacred is weakened. From a theoretical

perspective, this suggests that intelligent systems do not merely coexist with sacred space but actively reshape the conditions under which sacredness can be experienced.

Taken together, Durkheim's social theory, Eliade's phenomenology, and Turner's concept of liminality provide a powerful framework for understanding why the intelligent transformation of temple space is not a neutral process. Sacred space depends on social recognition, symbolic orientation, and ritualized thresholds. When these elements are reconfigured by system rationality, the very structure of the sacred–profane distinction is placed under pressure. This theoretical insight sets the stage for subsequent chapters, which examine how institutional logic, algorithmic governance, and lived religious experience are transformed as temples move from sanctuary to system.

### 3. The Temple as Institution and Symbolic Space

While established the temple as a sacred space constituted through social boundaries, symbolic orientation, and ritual thresholds, a further analytical step is necessary. Temples are not only spaces of religious experience; they are also institutional actors embedded in broader social, economic, and political contexts. At the same time, their institutional functions are inseparable from dense symbolic structures that exceed purely administrative logic. This chapter argues that the temple must be understood as a hybrid formation: simultaneously an institution that organizes resources and authority, and a symbolic space that produces and sustains religious meaning. This dual character is crucial for understanding why the introduction of artificial intelligence generates deep structural tensions rather than merely technical adjustments.

#### 3.1. The Temple as a Religious Institution: Organization, Authority, and Continuity

From a sociological perspective, temples function as religious institutions that ensure continuity over time. They manage personnel, property, finances, ritual schedules, education, and relations with the surrounding society. Max Weber's analysis of religious institutions emphasizes that routinization and organization are unavoidable if religious movements are to endure beyond charismatic origins (Weber, 1978). Temples embody this process of institutional stabilization. Institutionalization, however, does not imply secularization. In religious contexts, organizational structures are legitimized not merely through efficiency but through reference to tradition, sacred authority, and ritual competence. Clerical hierarchies, monastic rules, and ritual offices derive their legitimacy from their role in preserving correct practice and transmission. Authority is thus both organizational and symbolic. This dual legitimacy distinguishes religious institutions from purely bureaucratic organizations. Decisions within temples are not evaluated solely in terms of outcomes but in terms of their conformity to religious norms. Institutional rationality is therefore internally constrained by symbolic obligations. The temple as an institution operates within a normative horizon defined by the sacred. The introduction of AI-driven management systems challenges this balance. Algorithmic tools evaluate success through metrics—attendance numbers, donation flows, risk reduction, efficiency gains. These criteria introduce an external standard of legitimacy that may not align with religious norms. Over time, institutional decision-making risks being subtly reoriented away from symbolic fidelity toward system performance.

### 3.2. Symbolic Space and the Production of Religious Meaning

Beyond institutional organization, temples function as symbolic spaces in which meaning is produced and reproduced. Religious symbols are not merely decorative; they structure perception, memory, and action. Architecture, spatial hierarchy, iconography, soundscapes, and bodily movement all contribute to a symbolic environment that shapes religious subjectivity. Clifford Geertz famously described religion as a system of symbols that establishes powerful, pervasive, and long-lasting moods and motivations by formulating conceptions of a general order of existence (Geertz, 1973). The temple materializes such a symbolic system spatially. Its layout guides movement, directs attention, and communicates values without explicit instruction. Symbolic space operates through condensation and overdetermination. A single altar, statue, or ritual path may carry multiple layers of meaning accumulated over centuries. These meanings are not always consciously articulated; they are learned through repeated exposure and participation. The temple thus functions as a pedagogy of meaning that precedes reflection. AI-mediated systems, however, tend to translate symbolic environments into functional categories. Spaces become “zones,” rituals become “events,” and symbolic density is rendered as data points. While such translation may facilitate management, it risks flattening symbolic meaning by subordinating it to operational logic. The symbolic surplus that sustains religious depth is not easily compatible with system optimization.

### 3.3. Authority, Legitimacy, and the Spatialization of Power

Authority within temples is spatially enacted. Access to certain areas may be restricted; ritual roles may be differentiated spatially; movement may be hierarchically ordered. These spatial arrangements do not merely reflect power; they constitute it. Michel Foucault’s insights into the spatialization of power help illuminate how authority operates through architecture and regulation (Foucault, 1977). In religious contexts, however, spatial authority is not primarily disciplinary but symbolic. Restrictions signify holiness rather than surveillance. Control is justified by reference to purity, reverence, or ritual order. Authority is thus experienced as meaningful rather than coercive. The introduction of intelligent systems alters this experiential quality. Surveillance cameras, access controls, and automated monitoring introduce a mode of power that is continuous, invisible, and data-driven. Even when implemented for benign reasons, such systems shift the felt source of authority from religious office to technical infrastructure. This transformation has significant implications. Authority becomes less personal and less contestable. Decisions appear as outputs of systems rather than judgments grounded in tradition. From the perspective of religious studies, this represents a reconfiguration of how legitimacy is spatially produced and recognized within temple life.

### 3.4. Institutional Rationality and Symbolic Fragility in Intelligent Environments

The coexistence of institutional rationality and symbolic meaning within temples has always been fragile. Modernization has repeatedly pressured religious institutions to justify themselves in secular terms. Artificial intelligence intensifies this pressure by offering powerful tools for governance that operate according to system rationality. Niklas Luhmann’s theory of social systems highlights that modern institutions tend to operate according to functionally differentiated

logics (Luhmann, 2012). When religious institutions adopt system technologies, they risk internalizing logics foreign to their symbolic mission. What begins as support infrastructure can gradually redefine institutional priorities. This does not imply that temples should reject all technology. Rather, it underscores the need for reflexivity. Temples are institutions precisely because they preserve symbolic meaning against purely instrumental rationality. If intelligent systems are allowed to reorganize temple space without symbolic constraints, the institutional core of religion may be hollowed out from within.

This chapter has argued that the temple's dual character as institution and symbolic space makes it particularly sensitive to technological intervention. Artificial intelligence does not merely affect how temples are managed; it reshapes how authority is legitimized, how meaning is produced, and how space functions symbolically. Recognizing this dual structure is essential for understanding why the intelligent transformation of temples constitutes a profound religious issue rather than a neutral administrative reform.

#### 4. Artificial Intelligence and the Systematization of Temple Space

If Chapters 2 and 3 established the temple as a sacred boundary and a symbolic institution, the present chapter addresses how artificial intelligence introduces a qualitatively different mode of spatial organization. AI does not merely add new tools to temple administration; it brings with it a form of rationality that tends to reorganize space as a system. This process may be described as the systematization of temple space: the gradual transformation of a sanctuary oriented toward symbolic meaning into an environment governed by operational logic, data flows, and algorithmic control.

##### 4.1. System Rationality and Algorithmic Governance

Artificial intelligence operates through what systems theory describes as system rationality: the coordination of heterogeneous functions toward goals such as optimization, predictability, and stability (Luhmann, 2012). Rather than responding to meaning or normativity, systems respond to inputs and outputs, deviations and corrections. When AI is introduced into religious contexts, this rationality does not remain neutral; it actively reframes how space is perceived and governed. Applied to temples, system rationality tends to reconceptualize religious space as a set of interrelated processes to be managed. Visitor movement becomes a logistical problem, ritual scheduling becomes a coordination task, and religious gatherings are evaluated in terms of capacity, efficiency, and risk. From the standpoint of algorithmic governance, the temple is no longer primarily a sanctuary set apart from ordinary rationalities, but a complex environment requiring continuous regulation. This shift has important theoretical implications. Sacred space, as described in classical religious studies, is defined by qualitative difference rather than functional integration. It resists total transparency and instrumental control. System rationality, by contrast, aims precisely at eliminating opacity and unpredictability. As AI systems increasingly mediate decision-making within temples—determining access, allocating resources, and regulating behavior—the logic of sanctuary is gradually replaced by the logic of management.

Moreover, algorithmic governance tends to be self-reinforcing. Once introduced, systems generate data that justify further system expansion. For example, data collected for crowd management may later be used for behavioral analysis or financial optimization. This dynamic risks transforming temples into nodes within broader socio-technical systems, subject to the same rationalities as commercial or governmental spaces. The boundary between sacred and secular governance thus becomes increasingly porous.

#### **4.2. Surveillance, Datafication, and the Sacred**

One of the most visible manifestations of system rationality in temple space is the expansion of surveillance technologies. Cameras, sensors, facial recognition systems, and behavioral tracking tools are often justified on grounds of safety, theft prevention, or crowd control. From a technical standpoint, such measures appear reasonable. From a religious-studies perspective, however, they introduce a profound transformation in how sacred space is experienced. Surveillance renders behavior continuously visible and measurable. In doing so, it introduces a logic of suspicion into spaces traditionally governed by trust, reverence, and moral self-regulation. Religious spaces have historically relied on internalized norms rather than external monitoring. Devotion, repentance, and ethical reflection presuppose a degree of interiority—an awareness of being accountable before the sacred rather than before an observing system. Datafication further intensifies this shift. When movement, attendance, and ritual participation are translated into data, religious practice becomes legible to algorithmic evaluation. Patterns are identified, deviations flagged, and behaviors optimized. While such processes may improve organizational efficiency, they also redefine what counts as relevant religious activity. That which cannot be measured risks becoming invisible or marginalized.

This transformation affects the symbolic structure of sacred space. In a surveilled environment, the sense of withdrawal from ordinary social scrutiny is weakened. The temple increasingly resembles other monitored public spaces, undermining its experiential distinctiveness. From a Durkheimian perspective, this weakens the social mechanisms that sustain sacredness; from a phenomenological perspective, it alters the lived experience of transcendence. Importantly, surveillance does not simply add an external layer of control; it reshapes subjectivity. Visitors may become more self-conscious, less inclined toward spontaneous or vulnerable religious expression. The sacred space becomes a site of performance under observation rather than a refuge from visibility. This shift illustrates how system rationality penetrates not only spatial organization but also religious experience itself.

#### **4.3. Automation and the Reconfiguration of Ritual Space**

Beyond governance and surveillance, artificial intelligence increasingly automates aspects of ritual life. Scheduling systems determine ritual timing, automated kiosks facilitate offerings, digital interfaces guide devotional practices, and AI-generated content assists in religious instruction. These innovations are often framed as supportive tools that enhance accessibility and participation. Yet automation introduces a subtle but significant reconfiguration of ritual space. Rituals, as religious studies have emphasized, are not merely sequences of actions; they are formative practices that rely on repetition, bodily engagement, and symbolic excess. Their

efficacy depends not on efficiency but on the cultivation of attentiveness, patience, and transformation over time.

Automated systems, however, operate according to procedural logic. They emphasize standardization, speed, and predictability. When rituals are increasingly mediated by automated systems, there is a risk that they become procedural tasks to be completed rather than transformative encounters. The logic of “correct execution” begins to replace the logic of meaningful participation. This shift is particularly evident when access to ritual space is regulated by systems rather than by symbolic preparation. Automated entry systems, timed slots, and digital instructions structure participation in ways that resemble service delivery rather than sacred encounter. While such measures may increase throughput, they alter the experiential texture of ritual engagement.

Moreover, automation tends to externalize ritual authority. Decisions about timing, sequence, and access are increasingly embedded in systems rather than embodied in religious specialists. Authority becomes technical rather than interpretive. From the perspective of religious studies, this represents a displacement of ritual knowledge from lived tradition to algorithmic procedure. Taken together, these developments suggest that artificial intelligence does not merely coexist with ritual space; it actively reshapes it. Ritual becomes more manageable but potentially less transformative. The temple, in turn, risks becoming a site of religious service provision rather than a space of symbolic and existential reorientation.

## 5. Sacred Experience under Intelligent Conditions

If the previous chapter analyzed how artificial intelligence systematizes temple space at the level of governance, surveillance, and ritual organization, the present chapter turns to the experiential consequences of this transformation. Religious studies has long emphasized that religion cannot be reduced to institutions or doctrines alone; it is fundamentally concerned with lived experience—how individuals encounter the sacred, interpret transcendence, and orient their lives accordingly. This chapter argues that intelligent systems reshape not only the external structure of temple space but also the internal conditions under which sacred experience becomes possible.

### 5.1. The Transformation of Religious Experience: From Withdrawal to Continuous Connectivity

Classical theories of religious experience emphasize withdrawal from ordinary modes of rationality and perception. Sacred space traditionally functions as a site of interruption—an environment in which everyday concerns, instrumental reasoning, and social roles are temporarily suspended. Entering a temple often involves slowing down, observing silence, and adopting postures of receptivity. Such practices cultivate an experiential distance from the profane world, making room for transcendence.

Intelligent systems fundamentally alter this condition by extending regimes of connectivity, responsiveness, and monitoring into sacred space. Digital interfaces, real-time displays,

automated announcements, and continuous data flows reduce experiential discontinuity. The temple becomes less a space of withdrawal and more an extension of the connected world. Even when technologies are unobtrusive, their presence contributes to an ambient awareness of systems operating in the background (Zuboff, 2019).

This shift has significant experiential implications. Religious experience becomes embedded in a context of continuous mediation rather than rupture. The sacred is encountered not through interruption but through managed access. While this may increase accessibility, it risks undermining the experiential depth traditionally associated with sacred space. Transcendence, instead of appearing as a qualitative break, is reinterpreted as an optimized experience within a system. From a phenomenological perspective, this transformation weakens the contrast necessary for religious experience to emerge as distinctive. If the temple no longer feels qualitatively different from other system-managed environments, the experiential markers of sacredness—stillness, awe, reverence—become harder to sustain.

## **5.2. Presence, Mediation, and the Reconfiguration of the Sacred Encounter**

Religious experience is often described in terms of presence: the sense of being addressed by, or standing before, a transcendent reality. This presence is not reducible to belief; it is an experiential orientation cultivated through ritual, symbol, and space. Importantly, religious presence has traditionally been mediated through material forms—architecture, icons, sounds—but these mediations function symbolically rather than instrumentally. Artificial intelligence introduces a new kind of mediation: algorithmic mediation. Unlike traditional religious media, which point beyond themselves, algorithmic systems tend to foreground their own operational logic. Screens, sensors, and automated processes draw attention to management rather than mystery. The mediation of the sacred thus risks being displaced by the mediation of systems.

This displacement has subtle effects on how presence is experienced. Instead of feeling addressed by the sacred, participants may become more aware of being processed by systems. The experiential axis shifts from vertical transcendence to horizontal regulation. Even when religious symbols remain unchanged, the background infrastructure alters how they are encountered.

Moreover, algorithmic mediation introduces predictability into experiences that have traditionally valued openness and indeterminacy. Religious encounters often involve uncertainty, ambiguity, and waiting. Intelligent systems, by contrast, aim to minimize unpredictability. This tension raises the question of whether sacred presence can still be experienced as gift rather than outcome when mediated by systems optimized for control.

## **5.3. Agency, Responsibility, and the Reshaping of Religious Subjectivity**

A further dimension of religious experience concerns agency and responsibility. Religious traditions typically emphasize personal engagement—acts of devotion, repentance, ethical decision-making—that presuppose active participation. Sacred space supports this agency by framing individuals as responsible subjects before the sacred. Intelligent systems subtly reshape this subjectivity. Automated guidance, procedural instructions, and system-regulated participation reduce the need for interpretive judgment. Religious actors increasingly follow prompts rather

than discern meaning. While such systems may lower barriers to participation, they also risk diminishing the sense of personal responsibility central to many religious traditions. This transformation affects how individuals relate to religious authority. Authority becomes embedded in systems rather than embodied in interpreters or traditions. The locus of decision-making shifts from reflective agency to procedural compliance. Over time, religious subjectivity may be reconfigured from a responsive, discerning self to a system-compliant user.

From the perspective of religious studies, this shift raises normative concerns. Religious experience is not merely about comfort or accessibility; it involves ethical and existential accountability. If intelligent systems absorb too much agency, the formative dimension of religious experience—its capacity to shape moral character and self-understanding—may be weakened. Taken together, these experiential transformations suggest that intelligent systems do more than modernize religious practice. They reconfigure the phenomenological conditions under which sacred experience occurs. The temple remains a religious site, but the mode of experiencing the sacred shifts—from withdrawal to connectivity, from symbolic mediation to system mediation, and from active agency to procedural participation.

## 6. Conclusion

This study has argued that the increasing integration of artificial intelligence into temple spaces should not be understood merely as a process of technological modernization, but as a structural transformation that reconfigures the very conditions under which religious space, authority, and experience are constituted. By tracing the classical sacred–profane distinction and examining the temple as both institution and symbolic space, the paper has shown that artificial intelligence introduces a form of system rationality oriented toward optimization, visibility, and control that stands in tension with the symbolic, ritual, and experiential logics of sacred space. Intelligent systems do not simply assist religious practice; they reshape spatial boundaries, redefine authority, and alter the phenomenological texture of religious experience by extending regimes of connectivity and procedural governance into environments traditionally marked by withdrawal, liminality, and transcendence. While such transformations do not eliminate the sacred, they risk rendering it increasingly compatible with secular systems of management, thereby weakening the experiential and symbolic rupture upon which religious meaning depends. The future of temples in an age of intelligent systems, therefore, does not hinge on the acceptance or rejection of technology as such, but on the capacity of religious communities to critically negotiate the limits of system integration, preserving spaces of symbolic density, ritual autonomy, and existential openness within increasingly algorithmic environments.

### Author Contributions:

All authors have read and agreed to the published version of the manuscript.

### Funding:

This research received no external funding.

### **Institutional Review Board Statement:**

Not applicable.

### **Informed Consent Statement:**

Not applicable.

### **Conflict of Interest:**

The author declares no conflict of interest.

### **References**

- Asad, T. (1993). *Genealogies of religion: Discipline and reasons of power in Christianity and Islam*. Johns Hopkins University Press.
- Durkheim, É. (1995). *The elementary forms of religious life* (K. E. Fields, Trans.). Free Press. (Original work published 1912)
- Eliade, M. (1957). *The sacred and the profane: The nature of religion* (W. R. Trask, Trans.). Harcourt, Brace & World.
- Foucault, M. (1977). *Discipline and punish: The birth of the prison* (A. Sheridan, Trans.). Pantheon Books.
- Geertz, C. (1973). *The interpretation of cultures*. Basic Books.
- Luhmann, N. (2012). *Theory of society, Volume 1* (R. Barrett, Trans.). Stanford University Press. (Original work published 1997)
- Turner, V. (1969). *The ritual process: Structure and anti-structure*. Aldine.
- Weber, M. (1978). *Economy and society* (G. Roth & C. Wittich, Eds.). University of California Press.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.

#### **License:** Copyright (c) 2025 Author.

All articles published in this journal are licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited. Authors retain copyright of their work, and readers are free to copy, share, adapt, and build upon the material for any purpose, including commercial use, as long as appropriate attribution is given.

# Artificial Intelligence and the Religious Imagination of the Future: A Comparative Study of Eschatology, Soteriology, and Technological Utopianism

Zhanyi Li <sup>1,\*</sup>

<sup>1</sup> Lanzhou University, Lanzhou 730000, China

**\* Correspondence:**

Zhanyi Li

[lizhanyi028@163.com](mailto:lizhanyi028@163.com)

*Received: 10 December 2025/ Accepted: 26 December 2025/ Published online: 31 December 2025*

## Abstract

Contemporary discourses surrounding artificial intelligence are increasingly framed in terms of radical futures: superintelligence, existential risk, human transcendence, and civilizational rupture. These narratives often present themselves as secular and scientific, yet they display striking structural similarities to classical religious doctrines of eschatology and soteriology. This article offers a philosophical and comparative analysis of artificial intelligence futures through the lens of religious studies, arguing that contemporary AI imaginaries reproduce key elements of religious end-time and salvation narratives in a secularized form. By comparing AI discourse with traditions of eschatology, soteriology, and apocalyptic expectation, the study reveals the implicit religious dimensions embedded within technological utopianism and catastrophe narratives. The paper further argues that these AI-driven future visions reshape contemporary understandings of hope, redemption, and responsibility, often displacing ethical agency from communal and moral practices to technical systems. Rather than dismissing AI futures as mere ideology, the article contends that recognizing their quasi-religious structure is essential for critically assessing their cultural power and normative implications.

**Keywords:** Artificial Intelligence; Eschatology; Soteriology; Technological Utopianism; Religion and Technology; Futures Imaginaries

## 1. Introduction

Few contemporary technologies have generated as many speculative future narratives as artificial intelligence. Public and academic discussions of AI frequently invoke scenarios of civilizational transformation: the emergence of superintelligence, the obsolescence of humanity, the end of work, or the dawn of a post-biological era. These narratives are often framed as rational

projections grounded in scientific extrapolation. Yet their rhetorical structure, moral stakes, and temporal orientation suggest that something more than technical forecasting is at work.

This paper begins from the observation that contemporary AI discourse is saturated with ultimate futures. AI is not merely presented as a tool that will improve specific domains; it is portrayed as a force that will decide the fate of humanity as such. This framing bears a striking resemblance to religious traditions in which history is oriented toward an ultimate horizon—whether judgment, redemption, or cosmic renewal. The central question of this study is therefore not whether AI is a religion, but whether AI futures discourse functions as a secularized form of religious imagination.

Recent scholarship in religion and technology has increasingly noted these parallels (Geraci, 2018; Tirosh-Samuelson, 2021; Noble, 2020). However, much of this work remains descriptive or metaphorical. What is still lacking is a systematic philosophical comparison between AI futures and classical religious doctrines of eschatology and soteriology. This paper seeks to provide such a comparison by focusing on three interrelated dimensions: (1) temporal structure, (2) redemption narratives, and (3) moral agency.

## 2. Eschatology and AI Futures: The End as Singularity or Catastrophe

### 2.1. Classical Eschatology: Structure and Meaning

In religious traditions, eschatology refers to doctrines concerning the “last things”: the end of history, final judgment, resurrection, or cosmic transformation. Importantly, eschatology is not merely about chronology or speculative futurity; it is fundamentally about meaning. As Karl Löwith (1949) famously argued, eschatological thinking provides history with direction, coherence, and intelligibility by orienting it toward an ultimate horizon. Similarly, Jürgen Moltmann (1993) emphasizes that eschatology is not an appendix to theology but its organizing center, shaping how the present is understood and lived.

Eschatology structures time teleologically rather than causally. Historical events do not merely follow one another; they are interpreted in light of an anticipated end. This orientation transforms the present into a space of decision, responsibility, and expectation. Human action acquires significance not because it accumulates progress, but because it is lived toward an ultimate fulfillment or judgment. In this sense, eschatology is deeply practical: it motivates ethical commitment, endurance, and hope.

A defining feature of eschatological time is its discontinuity. The end does not emerge smoothly from historical development but interrupts it. Apocalyptic traditions in Judaism and Christianity, for example, portray the end as a radical break in which existing structures of power, injustice, and suffering are overturned. This rupture carries moral weight. Judgment separates good from evil, the faithful from the unfaithful, and truth from illusion. History is thus lived under the sign of the end—not as passive waiting, but as charged anticipation.

Crucially, classical eschatology also introduces limits to human control. The end is not something humanity engineers; it arrives as an event that exceeds human mastery. Even in

traditions that emphasize human cooperation with divine purposes, eschatological fulfillment ultimately transcends technical calculation. This feature distinguishes eschatology from modern notions of progress. Whereas progress assumes continuous improvement through human effort, eschatology insists on a decisive moment that redefines the meaning of all preceding history. From a philosophical perspective, eschatology thus performs a double function: it critiques present arrangements by measuring them against an ultimate horizon, and it resists the reduction of the future to technical planning. These functions will prove crucial when comparing religious eschatology with contemporary AI futures.

## 2.2. AI Eschatology: Singularity, Extinction, and Rupture

Contemporary discourse surrounding artificial intelligence reproduces the basic structure of eschatological thinking with striking consistency, even as it presents itself in secular and scientific terms. Central to this discourse is the notion of radical rupture—the idea that AI development will lead to a point at which existing historical categories no longer apply. The concept of the technological singularity, popularized by Ray Kurzweil (2005) and widely debated in subsequent literature, exemplifies this logic. The singularity marks a threshold beyond which human intelligence, agency, and social organization are fundamentally transformed or rendered obsolete. Like religious eschatology, AI singularity discourse is not merely predictive but meaning-laden. It reorients the present by framing current technological development as movement toward an ultimate horizon. Ordinary political, ethical, and social concerns are reinterpreted as preparatory stages for an impending transformation. The future thus becomes the decisive site of meaning, and the present is evaluated in terms of readiness for that future.

In parallel, existential risk narratives frame AI as a potential agent of catastrophic rupture. Scholars such as Nick Bostrom (2014) and Stuart Russell (2019) argue that advanced AI poses an unprecedented threat to human survival. Here again, the future is imagined in binary terms: salvation or annihilation, flourishing or extinction. This framing closely mirrors apocalyptic dualism, in which history culminates in final division. There is little room for gradual adaptation or plural futures; the end is total. Importantly, these narratives introduce moral urgency. Like religious apocalypticism, AI eschatology mobilizes fear and hope to motivate present action. Calls for alignment research, governance frameworks, and global coordination are often justified by reference to the stakes of the end. The future catastrophe or salvation functions as a regulative horizon that disciplines present behavior.

Recent critical scholarship has emphasized that such narratives are not neutral forecasts but moral imaginaries shaped by cultural assumptions and normative commitments (Danaher, 2020; Cave & Dihal, 2020). They tell stories about what ultimately matters, who bears responsibility, and what kind of future is worth striving for. From a religious-studies perspective, this reveals AI eschatology as a secularized form of apocalyptic expectation. The divine is replaced by technology, but the structure of ultimate concern remains.

## 2.3. Continuities and Transformations: Secularized Apocalypse

The comparison between classical eschatology and AI futures reveals both continuity and transformation. On the one hand, AI discourse inherits key eschatological features: an ultimate

horizon, a decisive rupture, moral urgency, and a reorientation of the present. On the other hand, it transforms these features by relocating agency and meaning within technical systems. One crucial difference lies in the source of rupture. In religious eschatology, the end arrives from beyond history—it interrupts human control. In AI eschatology, by contrast, rupture emerges from within human technological activity. Humanity becomes both the agent and the object of eschatological transformation. This internalization of the end marks a significant shift in the structure of responsibility.

Furthermore, AI eschatology often collapses transcendence into immanence. There is no beyond-history judgment; there is only system optimization or failure. The moral horizon narrows accordingly. Ethical questions become problems of risk management, alignment, and control rather than questions of justice, repentance, or reconciliation. As a result, the end loses its capacity to radically critique the present; instead, it reinforces a technocratic orientation toward governance and prediction. At the same time, AI eschatology preserves the affective power of apocalyptic thinking. It generates awe, anxiety, and hope, often amplifying them through media and institutional discourse. This affective continuity explains why AI futures resonate so strongly in contemporary culture. They respond to enduring human concerns about finitude, meaning, and destiny, even as they translate those concerns into secular language. From a philosophical standpoint, recognizing the eschatological structure of AI futures allows for a more critical engagement with their normative implications. It becomes possible to ask not only whether AI predictions are accurate, but what kind of future they imagine as ultimate, and what they ask humanity to sacrifice in the present.

#### **2.4. Eschatology without Transcendence: A Critical Assessment**

The secularization of eschatology in AI discourse raises a final critical question: what is lost when eschatological structure is detached from transcendence? Classical eschatology, for all its diversity, typically situates the end beyond human calculation and control. This distance creates space for humility, ethical responsibility, and critique of present power structures. AI eschatology, by contrast, risks absolutizing technical rationality. The end becomes a problem to be solved rather than a mystery that judges the present. This shift may undermine the critical function of eschatology by aligning ultimate meaning with system success. Salvation becomes survival, and judgment becomes optimization failure. From the perspective of religious studies, this transformation warrants careful scrutiny. AI futures do not simply replace religious eschatology; they inherit its power while reshaping its moral logic. Understanding this continuity and rupture is essential for assessing how contemporary societies imagine their ultimate futures—and how those imaginaries shape present ethical and political choices.

### **3. Soteriology and Technological Redemption: AI as a Secular Promise of Salvation**

#### **3.1. Classical Soteriology: Salvation, Transformation, and the Limits of Human Agency**

In religious traditions, soteriology concerns the question of salvation: how human beings are delivered from conditions perceived as fundamentally negative or threatening, such as suffering, guilt, finitude, and death. While doctrines of salvation vary widely across traditions, they share

several structural features. First, salvation is not merely improvement but transformation. It entails a qualitative change in the human condition—new life, liberation, reconciliation, or enlightenment. Second, salvation is not fully reducible to human effort. Even where ethical practice is central, salvation typically involves grace, transcendence, or participation in a reality that exceeds instrumental control.

In Christian theology, salvation addresses the problem of sin and alienation, culminating in reconciliation with God and the renewal of creation. In Buddhist traditions, liberation (*nirvāṇa*) concerns release from suffering and ignorance through transformative insight rather than technical mastery. In both cases, salvation presupposes an acknowledgment of human limitation. The human condition is marked by vulnerability and dependency; salvation cannot be engineered through technique alone.

Philosophically, this structure imposes limits on human agency. Salvation requires transformation of desire, perception, and orientation, not merely the optimization of existing capacities. It also introduces a moral dimension: salvation is inseparable from responsibility, ethical commitment, and self-reflection. Importantly, soteriology does not promise the elimination of all suffering through control; rather, it offers a way of meaningfully confronting suffering.

This framework stands in sharp contrast to modern narratives of progress, which tend to interpret human problems as solvable through technical intervention. The tension between soteriology and progress becomes especially pronounced in contemporary AI discourse.

### 3.2. AI and the Reimagining of Salvation as Optimization

Contemporary AI discourse increasingly presents artificial intelligence as a solution to humanity's deepest problems. AI is expected to cure disease, eliminate poverty, optimize governance, and even overcome death through life extension, cognitive enhancement, or digital immortality. In these narratives, salvation is reimagined as the successful optimization of biological, cognitive, and social systems. This technological vision of salvation differs fundamentally from religious soteriology. Rather than transformation through ethical or spiritual reorientation, AI promises redemption through enhancement. Human limitations are reframed as technical constraints, and suffering is treated as a system failure. Mortality becomes a bug to be fixed rather than a condition to be interpreted. Scholars have described this orientation as a form of "technological messianism" or "secular soteriology" (Geraci, 2018; Tirosh-Samuelson, 2021). AI occupies the role once assigned to divine or transcendent agents. It is imagined as the force that will redeem humanity—not by calling it to moral transformation, but by redesigning its conditions of existence. Crucially, this form of salvation is immanent. There is no appeal to transcendence beyond history; redemption unfolds entirely within technological development. This immanence reshapes the meaning of hope. Hope becomes confidence in innovation rather than trust in transformation. The future is no longer awaited; it is engineered. Recent journal literature has noted that such narratives often obscure their normative assumptions. By framing salvation as optimization, they privilege efficiency, control, and scalability as ultimate values (Danaher, 2020; Floridi, 2022). Questions of justice, meaning, and vulnerability are subordinated

to system performance. From a religious-studies perspective, this represents not merely secularization but a redefinition of what counts as salvation.

### **3.3. Transhumanism, Immortality, and the Denial of Finitude**

The soteriological dimension of AI discourse becomes especially visible in transhumanist visions of radical enhancement and immortality. Advocates of mind-uploading, digital consciousness, and post-biological existence often frame these projects as liberation from the constraints of the human body. Death, in this context, is no longer an existential horizon but a technical obstacle. Such visions echo religious longings for eternal life, yet they diverge in crucial ways. Religious traditions typically interpret immortality symbolically or relationally—life beyond death is meaningful because it is reconciled, judged, or transformed. In transhumanist discourse, immortality is primarily quantitative: more time, more capacity, more control.

Philosophical critiques have argued that this denial of finitude undermines the very conditions that make meaning possible (Hauskeller, 2016). Finitude structures responsibility, urgency, and care. By attempting to eliminate finitude, technological salvation risks emptying human existence of its moral texture. Moreover, the promise of technological immortality raises questions of exclusion and inequality. Salvation through AI is contingent on access to technology. Unlike religious salvation, which at least normatively claims universality, technological redemption is stratified. This introduces a quasi-soteriological hierarchy in which some are saved by enhancement while others remain vulnerable. From a religious perspective, this development resembles what theological traditions would describe as false salvation: a promise that addresses symptoms while ignoring deeper conditions of alienation and injustice. The focus on survival and enhancement displaces questions of reconciliation, meaning, and shared responsibility.

### **3.4. Salvation without Transformation? A Critical Comparison**

Comparing religious soteriology with AI-driven redemption narratives reveals a fundamental divergence in how salvation is conceived. Religious soteriology emphasizes transformation of the self and the community, often through practices that cultivate humility, compassion, and ethical responsibility. AI soteriology, by contrast, emphasizes transformation of systems rather than persons. The locus of change shifts from moral agency to technical infrastructure. This shift has significant ethical implications. When salvation is outsourced to systems, human responsibility is attenuated. Moral failure becomes system failure, and ethical deliberation is replaced by technical correction. As several recent studies have argued, this displacement risks depoliticizing and demoralizing social problems (Benjamin, 2019; Crawford, 2021).

At the same time, the persistence of soteriological structure in AI discourse indicates that secular modernity has not abandoned questions of salvation. Rather, it has relocated them. The longing for deliverance from suffering, finitude, and uncertainty remains, but it is expressed through technological imagination. Recognizing this continuity allows for a more nuanced critique. The problem is not that AI aspires to alleviate suffering, but that it often does so by narrowing the meaning of salvation to what can be optimized. Religious traditions remind us that salvation, if it is to be meaningful, must address not only conditions but orientation—how humans relate to one another, to time, and to their own limits.

## 4. Technological Utopianism and the Moralization of the Future

If the previous chapters have shown how contemporary AI discourse reproduces eschatological rupture and soteriological promise, this chapter examines how these elements converge into broader forms of technological utopianism. Utopian narratives do not merely describe desirable futures; they prescribe values, legitimize power, and orient collective action. From a religious-studies perspective, technological utopianism functions as a secularized moral horizon that reshapes how societies understand progress, responsibility, and the good life.

### 4.1. Utopia as Secularized Hope: From Religious Futures to Technological Horizons

Utopian thinking has long been intertwined with religious eschatology. Ernst Bloch famously argued that modern utopian imagination inherits the structure of religious hope even when stripped of explicit theological content (Bloch, 1986). The anticipation of a radically transformed future—free from suffering, conflict, and scarcity—originates in religious visions of redemption and the coming kingdom. Modernity did not eliminate these hopes; it translated them into secular idioms. Technological utopianism represents one of the most influential contemporary translations of this religious inheritance. In AI discourse, utopia is no longer grounded in divine intervention or moral renewal but in technological acceleration. Intelligent systems are expected to resolve social conflict, optimize governance, and eliminate inefficiency. The future is imagined as a state of systemic harmony achieved through computation rather than conversion.

This transformation reshapes the meaning of hope. In religious traditions, hope is often tied to patience, endurance, and ethical struggle in the present. It acknowledges uncertainty and the limits of human control. Technological hope, by contrast, is oriented toward prediction and planning. The future becomes a design problem rather than a promise that judges the present. This shift has profound normative implications: hope is no longer a virtue cultivated through practice, but a confidence placed in systems. Recent scholarship has noted that such utopian narratives tend to naturalize particular visions of social order by presenting them as technologically inevitable (Danaher, 2020; Cave & Dihal, 2020). From a comparative perspective, this mirrors religious millenarian movements that treated the coming kingdom as imminent and unavoidable. In both cases, the future acquires moral authority over the present.

### 4.2. AI Utopianism and the Reconfiguration of the Good Society

Technological utopianism is not merely about individual salvation; it articulates a vision of the good society. AI is frequently presented as the solution to structural social problems: biased governance, inefficient markets, human error, and moral inconsistency. Algorithmic systems promise fairness through data, neutrality through automation, and harmony through optimization. From a philosophical standpoint, this vision entails a significant redefinition of justice. Justice is increasingly framed as correct system behavior rather than as a contested moral ideal. Political disagreement appears as noise to be minimized, and ethical pluralism as inefficiency. The utopian horizon is one of frictionless coordination rather than moral deliberation. Critical scholars have argued that this orientation risks depoliticizing social life (Benjamin, 2019; Crawford, 2021). When social problems are framed as technical challenges, questions of power, inequality, and historical responsibility are obscured. The future appears morally resolved in advance by

intelligent systems, leaving little space for dissent or transformation. From the perspective of religious studies, this resembles what might be called immanentized eschatology: the belief that ultimate harmony can be achieved within history through the correct application of technique. Classical religious traditions have often warned against such visions, precisely because they conflate salvation with control. By promising a perfect society through optimization, AI utopianism risks reproducing a form of secular absolutism.

#### **4.3. Moral Displacement and the Externalization of Responsibility**

One of the most significant consequences of technological utopianism is the displacement of moral responsibility. In religious eschatology, the future judges the present, but it does so in a way that intensifies ethical responsibility. Human beings are called to repentance, justice, and care precisely because the end matters. In AI-driven futures, responsibility is increasingly externalized to systems, designers, and abstract governance frameworks. Moral failure becomes system failure. Instead of asking how humans ought to live together, discourse focuses on how systems ought to be aligned. Ethical agency shifts from communal practice to technical design.

Recent philosophical literature has emphasized this shift as a central danger of AI ethics (Floridi et al., 2018; Danaher, 2020). When morality is framed as an engineering problem, ethical life risks being reduced to compliance with system outputs. This reduction undermines the role of moral judgment, narrative understanding, and contextual responsibility. From a religious perspective, this displacement is particularly significant. Religious traditions typically insist that salvation and justice cannot be outsourced. Even when divine grace is central, human responsibility remains irreducible. By contrast, technological utopianism encourages a passive posture toward the future: if systems are built correctly, moral problems will resolve themselves. This posture mirrors what theologians have criticized as “cheap salvation”—a promise of redemption without transformation. The future is imagined as morally perfected, but the present is relieved of ethical struggle. Such narratives risk eroding the moral seriousness that eschatological thinking originally cultivated.

#### **4.4. Utopia without Transcendence: Limits of Technological Hope**

A final tension in technological utopianism concerns the absence of transcendence. Classical religious utopias, even when worldly, retain a reference to something beyond human calculation—divine justice, ultimate truth, or cosmic order. This transcendence functions as a critical limit on power. It prevents any human system from claiming final authority. AI utopianism lacks this limiting horizon. Because the future is produced entirely within technological systems, there is no external standpoint from which to judge their legitimacy. Optimization becomes its own justification. As a result, technological utopia risks collapsing into technocracy: rule by systems that cannot be meaningfully questioned from within their own logic. Religious studies offers a critical resource here. By revealing the implicit religious structure of technological utopianism, it becomes possible to question its claims to inevitability and moral sufficiency. The issue is not whether AI can contribute to human flourishing, but whether a future imagined solely in terms of system performance can sustain meaning, responsibility, and justice.

This chapter has argued that technological utopianism functions as a secularized moral horizon that inherits religious hope while transforming its ethical grammar. It promises harmony without transcendence, salvation without transformation, and justice without judgment. Recognizing these limits prepares the ground for the concluding chapter, which reflects on how religious traditions might contribute to a more responsible imagination of technological futures.

## 5. Conclusion

This study has argued that contemporary discourses on artificial intelligence are best understood not only as technical or scientific projections, but as future-oriented imaginaries that inherit and transform fundamental structures of religious thought. By comparing AI narratives with classical traditions of eschatology, soteriology, and utopian hope, the paper has shown that contemporary technological imagination reproduces the logic of ultimate rupture, redemption, and moral horizon while relocating these themes within immanent, system-driven frameworks. Artificial intelligence becomes a secular bearer of ultimate concern: it promises salvation through optimization, frames the future as a decisive threshold, and redefines moral responsibility as a problem of design and governance. Yet in doing so, it also displaces key elements central to religious futures—transcendence, ethical transformation, and communal accountability—risking a form of technological absolutism in which hope is reduced to prediction and justice to system performance. Recognizing the implicit religious dimensions of AI futures does not entail rejecting technological development; rather, it invites a more reflective engagement with the narratives that guide it. Religious traditions, with their long engagement with finitude, hope, and moral responsibility, offer critical resources for resisting the totalization of technological utopianism and for reimagining futures in which intelligence, whether human or artificial, remains accountable to meanings that cannot be fully optimized or controlled.

### Author Contributions:

Conceptualization, Z. L; methodology, Z. L; software, Z. L; validation, Z. L; formal analysis, Z. L; investigation, Z. L; resources, Z. L; data curation, Z. L; writing—original draft preparation, Z. L; writing—review and editing, Z. L; visualization, Z. L; supervision, Z. L; project administration, Z. L; funding acquisition, Z. L. All authors have read and agreed to the published version of the manuscript.

### Funding:

This research received no external funding.

### Institutional Review Board Statement:

Not applicable.

### Informed Consent Statement:

Not applicable.

## Conflict of Interest:

The author declares no conflict of interest.

## References

- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity Press.
- Bloch, E. (1986). *The principle of hope* (N. Plaice, S. Plaice, & P. Knight, Trans.). MIT Press.
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Cave, S., & Dihal, K. (2020). Hopes and fears for intelligent machines in fiction and reality. *Nature Machine Intelligence*, 2(2), 74–77.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Danaher, J. (2020). Automation and utopia: Human flourishing in a world without work. *Philosophy & Technology*, 33(3), 389–409. <https://doi.org/10.1007/s13347-019-00386-2>
- Floridi, L. (2022). AI ethics: A rethinking of its key concepts and challenges. *Minds and Machines*, 32(1), 1–15.
- Floridi, L., Cowls, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
- Geraci, R. M. (2018). *Temples of the future: Religion and the culture of artificial intelligence*. Oxford University Press.
- Hauskeller, M. (2016). *Palgrave handbook of posthumanism in film and television*. Palgrave Macmillan.
- Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. Viking.
- Löwith, K. (1949). *Meaning in history: The theological implications of the philosophy of history*. University of Chicago Press.
- Moltmann, J. (1993). *Theology of hope: On the ground and the implications of a Christian eschatology* (J. W. Leitch, Trans.). Fortress Press. (Original work published 1964)
- Noble, D. F. (2020). *The religion of technology: The divinity of man and the spirit of invention*. Penguin Books. (Original work published 1997)
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
- Tirosh-Samuelson, H. (2021). Transhumanism as a secularist faith. *Religions*, 12(4), 241.

**License:** Copyright (c) 2025 Author.

All articles published in this journal are licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited. Authors retain copyright of their work, and readers are free to copy, share, adapt, and build upon the material for any purpose, including commercial use, as long as appropriate attribution is given.