Studies on Religion and Philosophy, 2025, 1(4), 29-40
https://doi.org/10.71204/axy5m078

Can Artificial Intelligence "Understand" Religion? A Philosophical Inquiry into Meaning, Symbol, and Transcendence

Othman Al-Farsi [1,*]

[1] 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.