"Algorithmic Rationality" and "Ultimate Concern": Reconfiguring Religious Rationality in the Age of Artificial Intelligence

Wan Xing [1,*]

[1] Donghua Chan Temple, Shaoguan 512600, China

* Correspondence:

Wan Xing

wanxing@cscholar.com

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—affectively, 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 reframings 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:

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.