religions

MDPI

Article

From Instrument to Religious Authority: Mapping AI Applications, Roles, and Risks in Contemporary Buddhism

Wanje Gil [1] and Chungwhan Sung [2,*]

[1] Institute for Buddhist Education Studies, Dongguk University, Seoul 04620, Republic of Korea; jaygil8755@dgu.ac.kr
[2] Department of Indian Philosophy, College of Buddhist Studies, Dongguk University, Seoul 04620, Republic of Korea
* Correspondence: aruna415@hotmail.com

Abstract

Artificial intelligence is moving from the administrative margins of religious life into doctrinal explanation, personalized guidance, translation, counseling, heritage reconstruction, and embodied ritual representation. This study develops a framework for analyzing these developments in contemporary Buddhism, combining an integrative review of 65 publications, comparative analysis of 30 documented cases, and conceptual framework development. Seven domains of Buddhist life were crossed with seven AI modalities to produce 49 application intersections, classified as observed, emerging, or speculative. Each case was also coded for its initiating actor and purpose, distinguishing purposes stated by the initiator from those inferred from documented function. The analysis shows that religious significance follows less from technical type than from the role people assign to a system. Academic institutions, archives, and foundations initiated most documented applications, and the four cases performing religious authority acquired that position through institutional delegation, commercial substitution, or individual experiment rather than through technical advance. Counterfeit religious authority is distinguished as a separate pathway defined by unauthorized appropriation and concealed origin. The study proposes governance differentiated by role: accuracy, provenance, privacy, and human final judgment for instrumental uses; disclosure and dependency safeguards for relational uses; and doctrinal review, community authorization, and strict authority limits for religiously authoritative uses.

Keywords: artificial intelligence; Buddhism; digital religion; religious authority; generative charisma; embodied AI; counterfeit religious authority; AI governance

Check for updates

Academic Editor: Heather Kempton

Received: 6 August 2026
Revised: 7 September 2026
Accepted: 10 September 2026
Published: 11 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Artificial intelligence is no longer confined to the administrative margins of religious life. Generative systems produce explanations, images, translations, sermons, and devotional materials; conversational systems sustain interactive exchanges; analytical systems personalize content and infer user needs; and embodied systems speak, gesture, navigate, and occupy spaces already invested with ritual or symbolic meaning. These developments extend the long-standing reconfiguration of religion through digital media, but they also introduce a qualitative change: the medium may now generate religious language, respond as an apparent interlocutor, and act through a physical form. Digital-religion scholarship has shown that technologies do not merely transmit pre-existing religion; they reshape

presence, participation, interpretation, community, and authority (Campbell 2010). Generative and embodied AI intensify this process by making the boundary among medium, content producer, interface, and perceived agent increasingly unstable (Lim 2026).

Research on Buddhism and AI has expanded, but it remains divided among several partially connected conversations. One body of work uses Buddhist concepts to examine artificial consciousness, selfhood, ethics, interdependence, suffering, and the conditions of responsible technological development (Promta and Himma 2008; Hongladarom 2020; Hershock 2021). Another analyzes individual robots, chatbots, translation systems, educational applications, or digital Buddhist projects (Cheong 2020; Baffelli 2021; Li et al. 2022; Nehrdich et al. 2023). A third considers the broader transformation of Buddhism through intelligent technology and Humanistic Buddhism (Zheng 2024). A recent Korean scoping review similarly found that domestic scholarship was distributed across Buddhist interpretations of AI, the rearticulation of Buddhist values, and the development of AI-based Buddhist content, while systematic models and empirical evaluation remained limited (Gil et al. 2025). These studies establish a substantial foundation, but they do not yet provide an integrated account of where AI is entering Buddhist life, which technical modalities are involved, what roles the systems perform, or how those roles alter risk and governance requirements.

The central problem is that technical categories alone do not explain religious significance. The same large language model can remain an instrumental aid when it drafts a translation for expert correction, become relational when it remembers a practitioner's history and offers continuing guidance, and approach religious authority when it adjudicates doctrine, prescribes practice, delivers a sermon, or represents an institution. Conversely, a relatively scripted robot may acquire strong religious salience through embodiment, sacred imagery, temple placement, ritual performance, or institutional endorsement. Authority is therefore not a property that follows automatically from computational sophistication. It is produced through a configuration of technical capability, interaction, personalization, human-like presentation, institutional authorization, naming, user reception, and accountability (Campbell 2007, 2021; Loewen-Colón and Mosurinjohn 2022; Lim 2026). This role-centered approach also makes it possible to distinguish legitimate or institutionally mediated AI authority from counterfeit religious authority, in which synthetic media exploit the identity and charisma of an existing teacher, text, or institution without authorization.

This study addresses these gaps through an integrative literature review, comparative case analysis, and conceptual framework development. It combines a predefined set of 26 Korean studies with 39 international publications and a verified inventory of 30 observed or emerging cases. Seven domains of Buddhist life are crossed with seven AI modalities to generate forty-nine analytical intersections. Each intersection is assessed as observed, emerging, or speculative and then interpreted through role, risk, and governance dimensions. The purpose is not to estimate the prevalence or effectiveness of Buddhist AI applications, nor to decide whether AI is conscious or capable of religious attainment. It is to provide a comparative account of how AI is positioned within Buddhist activities and how agency, relationship, authority, and responsibility are reorganized in those settings.

The premise of this study is that technology is religiously configured before, during, and after adoption. Buddhist communities do not encounter AI as a finished object whose meaning is fixed by its technical properties. They commission it, name it, place it, authorize it, restrict it, and receive it, and each of these is a decision for which someone is answerable. The question pursued here is therefore not what AI is, but what people make it into within Buddhist settings, and with what consequences for authority and responsibility.

The study is guided by four research questions:

RQ1. How can contemporary Buddhist uses of AI be mapped across domains of Buddhist life and AI modalities?

RQ2. Which application intersections are observed, emerging, or speculative in the available literature and documented cases?

RQ3. How do documented human decisions about design, deployment, authorization, and reception configure AI as an instrument, a relational interlocutor, or a carrier of religious authority?

RQ4. What risks and differentiated governance requirements arise from these roles and applications?

The study makes four principal contributions. It first provides a 7×7 framework that organizes heterogeneous applications without presenting technological plausibility as evidence of adoption. It then advances a role-based account in which religious meaning depends more on the function and authority attributed to AI than on the technical type of the system. Third, it extends generative-charisma and digital-authority research by identifying role transitions, institutionally mediated authority, and counterfeit religious authority as analytically distinct configurations. Finally, it connects these distinctions to a role- and risk-based governance model: instrumental systems require accuracy, provenance, privacy protection, and human final judgment; relational systems additionally require role disclosure, human escalation, and safeguards against dependency; religiously authoritative systems require doctrinal review, community authorization, explicit responsibility, and strict limits on delegated authority; and counterfeit authority requires detection, disruption, removal, correction, and protection of personal and institutional rights.

2. Artificial Intelligence, Buddhism, and Religious Authority

2.1. Digital Religion and the Religious–Social Shaping of Technology

Digital-religion scholarship has challenged the assumption that media merely transmit an otherwise unchanged religious message. Media participate in the organization of presence, access, participation, community, interpretation, and authority. Religious actors consequently do more than place existing practices online; they select, negotiate, and reconfigure technologies in relation to inherited traditions and contemporary institutional needs (Campbell 2010). In Buddhist contexts, the trajectory from digitized scriptures and online Buddha halls to mobile applications, dialogue systems, and robot monks likewise reflects a long-term interaction among religious practice, media infrastructures, public representation, and technological innovation (Travagnin 2020).

This distinction is important because technical sophistication and religious significance do not rise together in a simple sequence. A powerful language model may remain a subordinate drafting tool under expert review. Conversely, a relatively scripted system may be received as a teacher, ritual representative, or sacred presence. The analytical shift from digital Buddhism to generative and embodied AI therefore requires attention not only to what a system computes but also to what it generates, how it interacts, where it is placed, who authorizes it, and what users understand it to be.

The religious–social shaping of technology approach provides a corrective to technological determinism. Campbell argues that religious communities do not simply accept or reject a medium on the basis of its technical properties. They interpret and negotiate technology through their histories and traditions, contemporary values and priorities, patterns of innovation and restriction, and communal discourses that justify particular uses (Campbell 2010). Technology is therefore socially and religiously configured before, during, and after adoption.

Religious shaping is reciprocal. Institutions configure the technology, but the technology also redistributes visibility, labor, access, and interpretive opportunity. Search and

recommendation systems privilege some textual sources over others. Conversational interfaces move religious inquiry from public or communal settings into private, dyadic encounters. Automated translation expands transnational access while potentially flattening doctrinal and linguistic differences. Embodied systems create new questions about presence, representation, and ritual efficacy. The resulting configuration is neither imposed by technology alone nor controlled entirely by religious institutions.

2.2. Generative Charisma and Role-Based Authority

Digital-religion research has shown that religious authority is multilayered and negotiated among hierarchy, organizational structures, ideology, texts, media actors, and audiences (Campbell 2007, 2021). AI complicates this arrangement because it can generate religious language, personalize interaction, and appear to respond in real time. The system may consequently do more than relay an already authorized teaching. It may participate in the production of the encounter through which authority is experienced.

Lim's theory of generative charisma identifies three mechanisms through which AI may acquire religious authority: captivation through generation, intimacy trust through personalization, and oscillating enchantment (Lim 2026). Captivation arises when the system's capacity to produce responsive or apparently insightful content evokes wonder or perceived exceptional ability. Intimacy trust develops through individualized, repeated, and often dyadic interaction. Oscillating enchantment describes the movement between recognition that the system is artificial and renewed fascination when its capabilities appear to improve or exceed expectations. Together, these mechanisms explain why AI-generated authority cannot be understood only as the transfer of authority from a human teacher to a digital channel.

Legitimacy and compelling performance must therefore be separated. A system can be experienced as persuasive without being doctrinally accurate, institutionally authorized, or accountable. Loewen-Colón and Mosurinjohn describe machine-mediated encounters as potentially spiritually authorizing when they are received as persuasive or legitimating events (Loewen-Colón and Mosurinjohn 2022). The present study treats such authorization as an empirical and institutional question rather than as an inherent property of the machine. It asks whose authority is being represented, how users recognize it, and what procedures permit its claims to be challenged or corrected.

Machine mediation may also work in the opposite direction, obscuring rather than displaying authority. A fluent, sourceless, and institutionally unmarked response can present itself as a view from nowhere, appearing to state what the tradition holds rather than what a particular teacher, school, or corpus holds. This anonymity is not neutral. It removes the very markers—lineage, school, translator, editorial decision—through which Buddhist claims have conventionally been situated and contested, and it does so at precisely the moment when a user is least equipped to notice their absence.

To connect technological applications with authority and governance, this study distinguishes three analytically ordered roles: instrumental, relational, and religiously authoritative AI. The categories concern the role performed by the system and attributed to it within a specific setting. They are not claims about consciousness, moral personhood, Buddha-nature, or membership in the saṅgha. These questions are excluded not because they are unimportant—they are the subject of a substantial philosophical literature (Promta and Himma 2008; Hongladarom 2020; Hershock 2021)—but because they are not answerable by the evidence examined here and are not required by the argument. Whether an AI system could possess Buddha-nature is a doctrinal and metaphysical question. Whether a community treats one as a teacher is an observable institutional fact with immediate consequences for authority, responsibility, and harm. A system granted no ontological standing

may still be given a name, placed in a hall, and received as speaking for a tradition, and it is that second question this study is designed to address.

Instrumental AI performs a bounded task under human direction. It may retrieve a passage, transcribe a sermon, assist translation, detect damage, draft a document, or recommend material, while accountable human actors retain interpretive and final decision-making authority. Relational AI functions as a continuing interlocutor through repeated interaction, personalization, memory, affective responsiveness, or ongoing guidance. Religiously authoritative AI performs, mediates, or is received as performing functions associated with doctrinal interpretation, spiritual direction, ritual efficacy, or representation of a teacher, lineage, or institution.

The categories are neither technologically exclusive nor irreversible. The same question-answering system may be instrumental when it cites a passage, relational when it remembers a user's practice history, and religiously authoritative when it adjudicates doctrine or directs conduct as an apparently legitimate teacher. Such transitions are not spontaneous. They follow identifiable changes in how a system is configured, presented, authorized, and received, which Section 6.5 sets out in relation to the documented cases.

Counterfeit religious authority is positioned outside this legitimate role continuum. It occurs when AI imitates or appropriates the identity, voice, image, text, or institutional standing of a religious figure or organization without authorization. Unlike a disputed but openly authorized experiment, counterfeit authority depends on deception or wrongful appropriation. It therefore requires detection, disruption, correction, and rights protection rather than merely stronger supervision of an otherwise legitimate service.

Because the term carries more than it should, its scope must be fixed at the outset. By counterfeit religious authority this study means something narrow: the unauthorized appropriation of the identity, voice, image, text, or institutional standing of a particular religious person or organization, presented in a form likely to be received as authentic. The term is not a judgment on AI in religion generally. Practitioners who regard all machine-generated religious speech as spiritually empty are advancing a different and legitimate theological argument, and the category proposed here does not adjudicate it. Nor does the term extend to novel, contested, or unconventional religious uses of AI that are openly presented as what they are. What makes an instance counterfeit is deception about origin and the appropriation of another party's accumulated trust—not innovation, and not disapproval.

The typology provides the conceptual bridge between the mapping framework and the risk analysis. Domain and modality identify where and how AI is used; role identifies what the system is doing within the religious relationship; risk modifiers identify the potential severity of harm; governance specifies the controls appropriate to that configuration.

2.3. Human Agency, Intention, and the Attribution of Role

The account developed above is not a claim about autonomous machine agency, and it should not be read as one. Every application examined in this study is the outcome of a chain of human decisions that begins long before any user interprets an output. Developers assemble the corpora, define the objectives and the refusal conditions, and choose the voice, the name, and the body. Institutions decide whether a system is commissioned, endorsed, placed, or withheld, and under whose authority it speaks. Users decide what to ask, what weight to give an answer, and whether an exchange is information, companionship, or instruction. Outputs are conditioned by model design, training data, prompts, interfaces, and deployment decisions, although individual outputs are not fully predictable. Developers and providers shape what is technically possible; institutions select sources and set the boundaries of authority; users accept, resist, reinterpret, or repurpose what

results, and those responses return as revised prompts, retrained models, and further deployments. The analytical question is therefore not only what an AI system does, but what people do through it and with it.

Three types of intention should be distinguished, because they belong to different parties and frequently diverge. Design intention determines what a system is built to be capable of and how it will appear. Institutional intention determines whether a religious body lends its standing to the result, and in what setting. Use intention determines what the system is actually asked to do once it is available. A translation model built for correction by trained scholars may be deployed by a temple as a public interface and received by an individual as personal direction, without any change in the model itself. Divergence among design, institutional, and use intentions can therefore produce roles different from those originally described by a system's developers or sponsors.

Science-and-technology-studies research describes design intention as inscription: designers build into an artifact a script that anticipates its users, their competences, and the setting in which it will act (Akrich 1992). The script is prescriptive but not determining. Users may follow it, ignore it, or turn the artifact to purposes its makers did not envisage, and the technology in turn mediates how situations are perceived and what actions appear available (Verbeek 2011). This is the same reciprocity that the religious–social shaping approach identifies in religious communities' negotiation of new media (Campbell 2010), examined one step earlier—at the point of design rather than the point of reception. A robot given a monastic name, a robe, and a position in a hall carries a script; whether it is received as a teacher depends on what a community does with that script.

It follows that the three roles distinguished above are conferred rather than achieved. A system does not become relational by improving, and it does not become religiously authoritative by becoming more capable. It occupies a role because people place it in a position where that role is available to it and because others treat it accordingly. This is why the same model may be instrumental in one setting and authoritative in another, and why a technically modest system may carry more religious weight than a sophisticated one. It also means that every role transition corresponds to a decision that could have been made differently: to personalize, to embody, to name, to endorse, to place, or to withhold disclosure. Each is a point at which responsibility attaches, and each is a point at which governance can intervene.

Distributing agency in this way does not dissolve accountability; it locates it. Developers, vendors, data curators, Buddhist institutions, monastics, scholars, counselors, archivists, platforms, and users participate differently in the conditions that produce an AI outcome, and should bear responsibility in proportion to their control and competence; a causal chain with many participants does not reduce the obligations of any one of them. Where a system speaks in a religious register, the relevant question is not whether the machine intended anything, but which parties decided that it should speak, on whose authority, and with what provision for correction.

This position has two consequences for what follows. Methodologically, it is why the case register records not only what each application does but who initiated it and what purpose was coded for it, and why those codings are reported alongside the mapping in Section 5. Analytically, it is why the governance framework in Section 7 addresses both technical properties and human decisions. Technical controls remain necessary, but decisions to authorize, name, place, and disclose a system determine where responsibility lies.

3. Materials and Methods

3.1. Research Design

This study used a literature-based conceptual design integrating an integrative literature review, comparative case analysis, and conceptual framework development. It did not estimate prevalence or effectiveness. Instead, it organized a heterogeneous and rapidly changing field so that Buddhist activities, AI functions, performed roles, risks, and governance requirements could be compared systematically (Whittemore and Knafl 2005; Torraco 2005; Snyder 2019).

This study built on a published Korean scoping review that had identified 26 Korean studies from 2016 to September 2025 (Gil et al. 2025). Those studies were retained as a predefined national evidence set and recoded for the present questions. They were combined with 39 included international publications and a separately documented inventory of 30 Buddhist–AI cases. Cases were used as analytically informative instances rather than as a statistically representative sample (Yin 2018). Figure 1 summarizes the resulting four-stage architecture, including matrix-level coding of the forty-nine intersections and case-level coding of initiating actors and coded purposes.

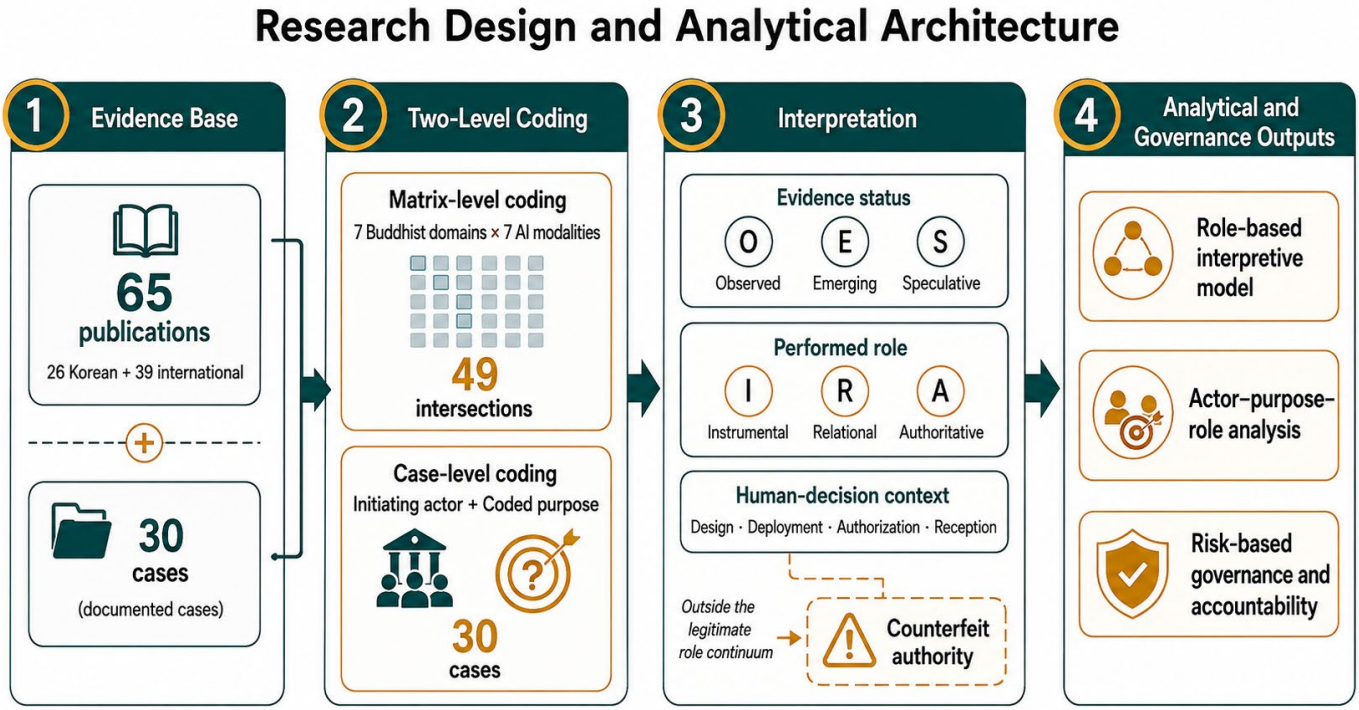


Figure 1. Research design and analytical architecture.

The evidence base of 65 publications and 30 documented cases is analyzed through two linked coding tracks. Matrix-level coding crosses seven Buddhist domains with seven AI modalities to produce 49 intersections. Case-level coding records the initiating actor and coded purpose for each of the 30 cases. Evidence status, performed role, and the human-decision context of design, deployment, authorization, and reception inform a role-based interpretive model, actor–purpose–role analysis, and risk-based governance and accountability. Counterfeit authority remains outside the legitimate role continuum.

The analytical architecture crossed seven domains of Buddhist life with seven AI modalities to create forty-nine intersections. Each intersection was classified as observed, emerging, or speculative and interpreted through role, risk, and governance dimensions. Instrumental, relational, and religiously authoritative AI formed the legitimate role continuum; counterfeit religious authority was treated as a separate pathway based on unautho-

rized simulation or appropriation. Category development was iterative: initial concepts guided coding, while cases and boundary conditions were used to refine definitions (Jabareen 2009; Jaakkola 2020).

3.2. Literature Search and Selection

Literature identification followed two streams. The Korean stream comprised the 26 studies identified in the published Korean review after searches of KCI, DBpia, RISS, KISS, and e-article; no additional Korean database search was conducted. Each record was re-examined for relevance to the present domains, modalities, roles, and risks. Full bibliographic and coding information is provided in Supplementary Table S1.

The international search covered publications from 1 January 2000 to 5 August 2026. Because subscription-database exports were not available in the study workflow, records were identified through scholarly publisher platforms, targeted journal and scholarly-web searches, backward and forward citation searching, and searches for recurrent authors, projects, and technologies. Queries combined Buddhism-related terms with artificial intelligence, generative AI, language models, chatbots, robotics, machine learning, natural language processing, computer vision, translation, retrieval-augmented generation, deepfakes, avatars, and embodied AI. The exact queries and search dates are retained in Supplementary Table S2a–c.

Eligible sources directly examined Buddhism and AI or analyzed AI applications in Buddhist doctrine, education, practice, care, propagation, heritage, administration, transnational exchange, or religious authority. Peer-reviewed articles, scholarly books and chapters, substantial conference papers, and research reports were eligible. Records were excluded when Buddhism or AI was incidental, when mindfulness was purely secular and clinical, when no identifiable AI component was present, or when the available text was insufficient for assessment. News, organizational pages, product listings, and project materials were evaluated through the separate case procedure.

The search captured 78 candidate occurrences. After 24 duplicate occurrences were consolidated using DOI, normalized title, authorship, year, and source information, 54 unique records were screened. Six were excluded at title or abstract stage and nine after full-text or complete-record assessment, leaving 39 international publications. Two researchers compared eligibility decisions and resolved differences by discussion. The combined scholarly corpus therefore comprised 65 publications.

3.3. Case Identification and Verification

A case was defined as an identifiable AI system, project, deployment, experiment, or documented misuse in which AI performed a substantive function in a Buddhist context. Searches covered 2000–2026 and used general terms together with named expressions such as robot monk, Buddhist chatbot, AI Buddha, BuddhaBot, Buddhist OCR, AI-generated sutra, and Buddhist deepfake. Project, institution, temple, university, developer, and application names identified during the literature review were also searched.

Evidence included peer-reviewed studies; scholarly books, chapters, conference papers, and research reports; official Buddhist, university, research, government, or developer materials; credible journalism; and directly inspectable public artifacts. Official sources verified existence, stated purpose, technical function, development stage, and institutional setting, but were not treated as independent proof of doctrinal reliability, safety, effectiveness, or user acceptance. Non-academic cases normally required two independent sources, preferably one primary institutional or technical source and one independent scholarly or journalistic source.

Evidence quality and implementation status were coded separately. Peer-reviewed publications received the highest evidentiary grade, followed by scholarly or institutional research, official project sources, multiple credible reports, and single or promotional sources. Observed status required verifiable operation, release, demonstration, or empirical evaluation; emerging status covered development, testing, limited deployment, or concrete institutional planning. Speculative applications were retained only in the matrix and were not counted as documented cases.

Reports referring to the same underlying project or event were consolidated into one case. Separate records were created only when a later version materially changed the AI modality, religious function, institutional setting, target group, or authority status. Two researchers jointly confirmed eligibility, project boundaries, evidence grade, operational status, primary domain, primary modality, role, and risk coding. The final inventory contains 30 observed or emerging cases and is provided as Supplementary Table S3.

3.4. Coding and Framework Construction

The coding scheme is deliberately narrow. It is built to answer a specific set of questions: what role a system is given in a Buddhist setting, who confers that role, how far its outputs are treated as authoritative, and what follows for risk and responsibility. Domains and modalities are recorded because they locate an application, not because they are of independent interest. The analytical weight falls on role, evidence status, and the conditions of authorization.

Six coding families were applied: Buddhist domain, AI modality, evidence grade, evidence status, role type, and risk category. Primary domain was assigned by the application's principal religious or institutional purpose and intended beneficiary. Primary modality was assigned by the capability most directly related to that purpose; secondary domains and modalities were retained when relevant. Risk coding was non-exclusive and considered users, data sensitivity, vulnerability, authority intensity, opacity, scale, and consequences of error.

The seven domains and seven modalities are heuristic rather than ontological categories. An initial set derived from the reviewed literature was revised during full-text verification of the documented cases; the categories were not treated as fixed, and boundary cases were resolved by consensus and recorded in Supplementary Table S4. The matrix is designed to make comparison, evidence gaps, and role transitions visible. It does not claim that Buddhist practice or AI development is organized in these terms by the actors themselves.

The seven domains were Dharma propagation and public engagement; Buddhist education and scholastic training; meditation, practice, and devotional life; counseling, care, and social welfare; translation and Buddhist cultural heritage; institutional administration and temple operations; and transnational exchange and globalization. The seven modalities were generative content; retrieval and question answering; conversational and agentic systems; speech, language, and translation; visual analysis and digital reconstruction; analytics and personalization; and embodied AI and autonomous physical action.

All forty-nine intersections were reviewed after full-text verification and duplicate-source consolidation. The two researchers jointly reviewed the complete coding register, discussed boundary cases, and confirmed evidence status, linked evidence, predominant role, risk codes, and the final coding rationale by consensus. The complete audit trail is supplied in Supplementary Table S4. The matrix is an analytical heuristic rather than a statistically validated measurement model; its purpose is to compare heterogeneous applications and make evidence gaps and role transitions visible.

Two further codings were added to the case register for this analysis. Initiating actor records the party that decided upon and publicly presented each application, distinguishing Buddhist orders, temples, and monastic authorities from academic and research institutions, commercial developers, lay groups and individual practitioners, and cases with no identifiable initiator. Coded purpose records the aim associated with the application. Each purpose code is flagged according to whether it was stated by the initiator in its own documentation or reconstructed from documented function, and the evidence for every assignment is given in Supplementary Table S3.

4. Mapping AI Across Seven Domains of Buddhist Life

The reviewed literature and verified cases were organized according to the principal religious or institutional purpose served by each application. Because a single system may combine retrieval, generation, conversation, speech processing, visual analysis, personalization, and physical embodiment, the seven domains should not be interpreted as technologically exclusive categories. Primary classification was determined by the application's central purpose and intended beneficiary, while additional functions and secondary domains were recorded separately. For example, a scripture-based question-answering system was classified under public engagement when designed for newcomers, under education when used for structured doctrinal learning, and under practice when it provided sustained guidance concerning meditation or devotional conduct. This purpose-based classification reduced duplicate counting while retaining the multifunctional character of Buddhist AI applications.

4.1. Dharma Propagation and Public Engagement

Embodied AI constituted the most publicly visible form of application. Mindar, the android representation of Kannon installed at Kodaiji Temple in Japan, delivers an interpretation of the Heart Sutra through speech, bodily movement, multilingual presentation, and projected media. The temple described the project as a means of making Buddhist teachings accessible to people who might otherwise remain uninterested in religion (Kodaiji Zen Temple 2019; Zheng 2024). Xian'er, developed in connection with Longquan Temple in China, similarly combined a monastic appearance with simple question answering, recitation, and public communication. Cheong interprets this case in terms of the communicative affordances through which robots may mediate spiritual knowledge and religious community (Cheong 2020).

Four of the seven intersections in this domain were classified as observed and three as emerging. Evidence was strongest for content generation, introductory question answering, conversational access, and embodied public engagement. Speech transformation, visual production, and audience personalization were more often proposed or indirectly documented than verified as stable Buddhist institutional services. The principal risks concerned doctrinal accuracy, misleading presentation, privacy in interactive systems, and ambiguity regarding the identity and authority of public-facing AI (Supplementary Table S4).

4.2. Buddhist Education and Scholastic Training

Knowledge-retrieval and question-answering systems constitute the most clearly documented modality in this domain. BuddhaBot was initially developed as a non-generative dialogue system that retrieves passages from selected Buddhist scriptures. BuddhaBot-Plus subsequently incorporated generative AI to produce question-specific explanations. The distinction is educationally significant. Retrieval systems can support textual location and source-based study while constraining free-form generation, whereas generative systems may offer greater accessibility but introduce additional interpretive and doctri-

nal risks. In 2025, Kyoto University announced a project to adapt BuddhaBot-Plus for use by Bhutan's monastic community, with dedicated system development, monitoring by monks, and a formal safety-assessment process (Kyoto University 2025).

Two intersections in this domain were classified as observed and five as emerging. The evidence supports the existence of working retrieval and dialogue systems and empirical evaluations of AI-generated Buddhist answers, but it does not support a general conclusion that AI has been widely integrated into Buddhist educational institutions. The main risks include doctrinal error, homogenization across Buddhist traditions, opaque textual sources, overreliance on fluent but unverified explanations, learner-data collection, and erosion of accountable human interpretation (Supplementary Table S4).

4.3. Meditation, Practice, and Devotional Life

Embodied ritual applications were more limited. In 2017, SoftBank's Pepper robot was publicly demonstrated chanting Buddhist sutras and striking a drum as part of a proposed funeral service. The demonstration verified the performance of ritualized movements and sounds, but it did not demonstrate widespread or sustained adoption in actual funerals (Reuters 2017). The case is therefore important less as evidence of an established ritual practice than as a boundary case concerning whether ritual validity depends on intention, consciousness, ordination, institutional delegation, correct external performance, or recognition by participants.

Four intersections—practice-related question answering, conversational guidance, voice-based meditation or chanting, and practice analytics—were classified as observed. Observed evidence included source-grounded Buddhist question-answering and publicly available conversational practice applications combining guided meditation, adaptive instruction, and progress tracking, together with voice-based guidance. Much of this evidence, however, came from prototypes, app-store records, or project documentation rather than longitudinal evaluation. Personalized practice-text generation and ritual robotics were classified as emerging, while Buddhist-specific posture-recognition feedback remained speculative. The evidence therefore indicates visible development in digital guidance and tracking but limited institutional support for automated spiritual assessment or embodied ritual substitution. Principal risks include inappropriate personalization, dependence on automated guidance, exposure of sensitive practice information, doctrinal simplification, and uncertainty regarding the religious validity of AI-mediated ritual acts (Supplementary Tables S3 and S4).

4.4. Counseling, Care, and Social Welfare

The most concrete Korean case is Cheongjeong MaeumBot, a Buddhist counseling chatbot developed by Myeongjun Kim. The project structured chatbot counseling around the Four Noble Truths and five categories of mental defilement, concentrating on the diagnosis and reduction of ordinary forms of distress. Research participants used the prototype and provided feedback, meaning that the project went beyond a purely conceptual proposal. Nevertheless, the study did not establish broad deployment, clinical effectiveness, or suitability for severe mental-health crises (Kim 2024).

Only the Buddhist counseling-chatbot intersection was classified as observed. Five intersections were emerging, while Buddhist-specific safety and companionship robots remained speculative. This domain should therefore be interpreted as socially significant but empirically underdeveloped. The strongest evidence concerns prototype design and limited user interaction rather than effectiveness, scalability, or sustained institutional integration. Major risks include exposure of sensitive data, inaccurate assessment, emotional

dependence, delayed access to professional care, and confusion between Buddhist pastoral support and licensed clinical treatment (Supplementary Table S4).

4.5. Translation and Buddhist Cultural Heritage

This was one of the most empirically developed domains. In 2025, the Buddhist Digital Resource Center released a Tibetan OCR application capable of converting scanned Tibetan texts into searchable and reusable digital text. The beta application supports multiple Tibetan scripts and was developed with the participation of Monlam AI collaborators. Its immediate functions include transcription, search, translation support, and use in digital research environments (Buddhist Digital Resource Center 2025).

Machine translation constitutes another substantial cluster of evidence. Li et al. developed a neural machine-translation approach for translating Chinese Buddhist sutras into English (Li et al. 2022). Nehrdich et al. later introduced MITRA-zh, drawing on more than 209,000 bilingual training pairs and a manually corrected evaluation set for Buddhist Classical Chinese. Their specialized model outperformed general commercial translation systems in the reported evaluation, while the authors emphasized the relevance of such systems to large-scale corpus analysis and semantic retrieval (Nehrdich et al. 2023). Tsadra Foundation's collaboration with the Dharmamitra team similarly integrates AI into Buddhist translation workflows while explicitly presenting it as assistance for, rather than replacement of, trained human translators (Tsadra Foundation n.d.).

Applications in this domain were overwhelmingly instrumental. Four intersections—archive retrieval, translation, OCR and digitization, and corpus or conservation analytics—were classified as observed. Heritage commentary and docent agents were emerging, while a specifically Buddhist autonomous museum robot remained speculative. The principal risks concern OCR and translation errors, disappearance of textual variants, unequal representation of traditions and languages, uncertain rights to translated corpora, unauthorized model training, and the presentation of generated or reconstructed material as authentic heritage (Supplementary Table S4).

4.6. Institutional Administration and Temple Operations

The provisional framework includes administrative-document drafting, internal-regulation retrieval, reservation and reception agents, telephone voice systems, facility monitoring, operational analytics, and guidance or maintenance robots. These applications are technically feasible and commonly found in secular organizations. The verified evidence base, however, did not identify a Buddhist institution demonstrating comprehensive and sustained AI integration across temple administration. Technical plausibility should therefore not be represented as evidence of actual Buddhist implementation.

Six intersections were classified as emerging and one—autonomous temple-guidance and facility robotics—as speculative. No intersection was classified as observed at the level required to demonstrate stable Buddhist institutional deployment. Institutional administration was consequently the least empirically developed of the seven domains. It should be presented as an implementation gap rather than as an area of established adoption. Relevant governance requirements include data minimization, access control, auditability, cybersecurity, human authorization of consequential decisions, and separation of administrative optimization from judgments concerning religious membership or spiritual status (Supplementary Table S4).

4.7. Transnational Exchange and Globalization

The clearest case is the collaboration between Kyoto University and Bhutan's Central Monastic Body. The project involved more than simply making a Japanese-developed chatbot available online. It included an official request from Bhutan's monastic authority,

development of a dedicated version of BuddhaBot-Plus, monitoring by Bhutanese monks, and a planned safety-assessment process. The case demonstrates that transnational deployment requires negotiation over textual corpora, doctrinal relevance, institutional authorization, cultural adaptation, and responsibility for generated answers (Kyoto University 2025).

Four intersections—multilingual content, multilingual question answering, foreign-visitor guidance, and interpretation—were classified as observed, while video localization, international-audience analysis, and event-guidance robotics were emerging. This comparatively high number of observed intersections should be interpreted cautiously because several constitute secondary uses of systems whose principal domain is translation, education, or public engagement. The major risks include loss of doctrinal and cultural context, dominance of well-resourced languages and Buddhist traditions, inaccurate localization, cross-border data transfer, and representation of one institutional interpretation as universally Buddhist (Supplementary Table S4).

Across the seven domains, the evidence was uneven rather than uniformly distributed. Translation and cultural heritage contained the strongest concentration of mature instrumental applications. Public engagement and devotional practice contained the most visible boundary cases involving embodiment, ritual performance, and religious authority. Education and counseling showed substantial conceptual and experimental activity but limited evidence of broad institutional adoption. Temple administration remained largely emerging or speculative. These differences support the need for the 7 × 7 framework presented in the following chapter: AI adoption in Buddhism cannot be adequately represented through a single continuum of technological advancement because evidentiary maturity, religious purpose, role intensity, and risk vary across both domains and modalities.

5. The 7 × 7 Buddhist–AI Application Framework

The domain findings were integrated into a 7 × 7 framework crossing seven domains of Buddhist life with seven AI modalities. The forty-nine cells are analytical intersections rather than forty-nine equally mature technologies. Each cell identifies an application that is observed, emerging, or speculative and records its predominant role and associated risks. This structure separates religious purpose from technical capability, supports comparison across contexts, and prevents technological plausibility from being treated as evidence of adoption.

5.1. Modalities and Evidentiary Distribution

The modalities distinguish generation, source retrieval, sustained conversation, language transformation, visual analysis, analytics, and physical action. They are functional rather than mutually exclusive technical architectures: a robot may combine dialogue, speech, vision, retrieval, and movement, while a counseling agent may combine a doctrinal corpus, generative language, memory, and user profiling. Primary modality identifies the capability most central to the Buddhist function; secondary capabilities remain documented in Supplementary Table S4.

Nineteen intersections were observed, twenty-six emerging, and four speculative. Retrieval and conversational systems each contained five observed intersections, whereas visual and embodied AI were less mature overall. Observed visual use was limited to OCR and three-dimensional reconstruction in the cultural-heritage domain. Embodied systems attracted public and scholarly attention but remained concentrated in a few visible deployments and demonstrations. Translation and heritage contained the strongest mature instrumental applications; administration contained no observed intersection.

Role and maturity did not move together. Retrieval, translation, OCR, and corpus analytics were comparatively mature but overwhelmingly instrumental. Embodied AI had fewer mature applications yet contained the matrix's two religiously authoritative intersections. Across the forty-nine representative applications, forty-one were primarily instrumental, six relational, and two religiously authoritative. Counterfeit religious authority was not treated as a fiftieth legitimate application because its defining feature is unauthorized appropriation rather than modality.

5.2. The Forty-Nine Intersections

Figure 2 presents the complete matrix. Each cell contains a concise representative application. Fill color encodes evidence status (observed, emerging, or speculative), while boundary treatment encodes predominant role: unmarked cells are instrumental, solid-outlined cells are relational, and double-outlined cells are religiously authoritative. The visualization is not an adoption roadmap: observed status does not establish effectiveness or doctrinal legitimacy, and speculative status does not imply desirability. Supplementary Table S4 provides the linked literature and cases, evidence count, role, risks, and coding rationale for every cell.

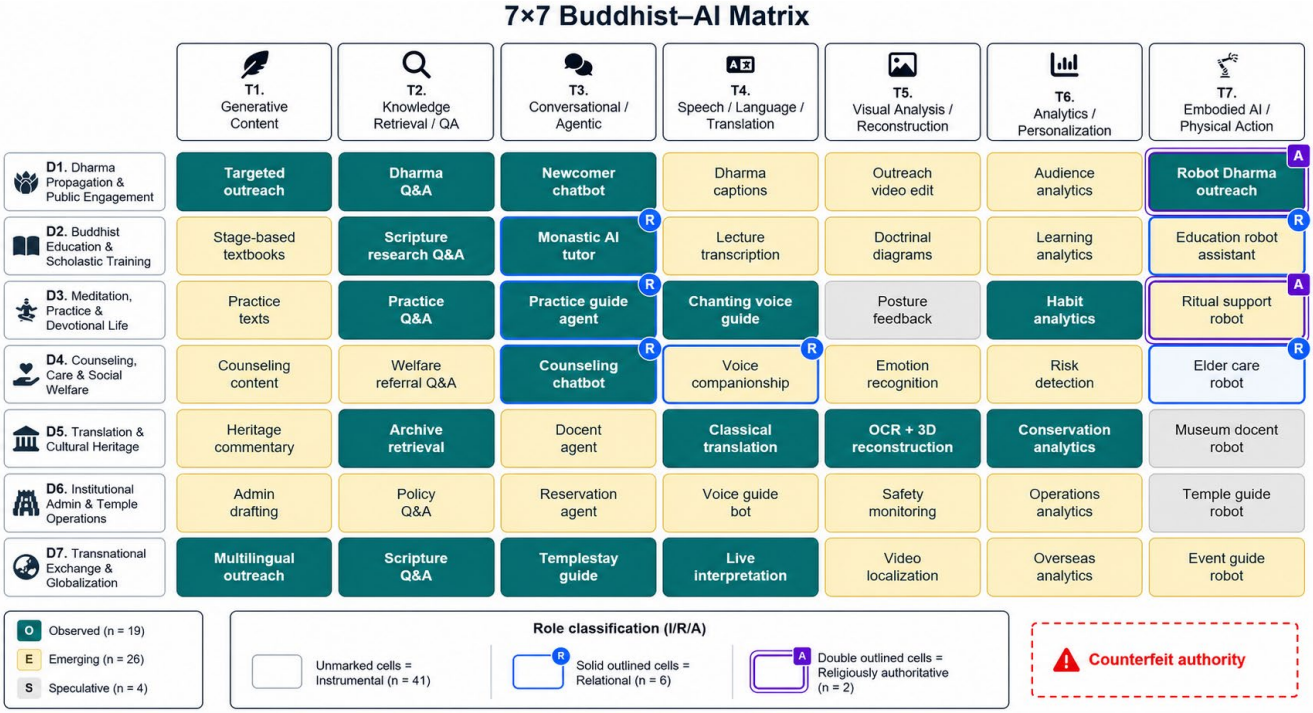


Figure 2. The 7 × 7 Buddhist–AI application matrix. Seven Buddhist domains are crossed with seven AI modalities. Cell fill encodes evidence status (O = observed, E = emerging, S = speculative); boundary treatment encodes predominant role (unmarked = instrumental, solid outline = relational, double outline = religiously authoritative). Counterfeit authority is displayed outside the legitimate role continuum.

Several patterns are central. Text-based retrieval and dialogue benefit from abundant digitized Buddhist materials, but sustained conversation is more likely than retrieval to acquire relational functions. Speech and translation show mixed but developed evidence, while raising interpretive problems across historical languages and traditions. Physical embodiment remains less mature but can create strong authority effects through appearance, location, ritual positioning, and institutional recognition. The matrix therefore shows that technological maturity, religious role, and normative acceptability are analytically distinct.

The matrix records where AI has entered Buddhist life, but it does not show who brought it there or for what declared reason. Table 1 therefore reports two further codings of the case register—the actor who initiated each application and the purpose coded for it—cross-tabulated against the role the application performs; the case-level listing is given in Supplementary Table S3. Academic and research institutions initiated or co-initiated sixteen of the thirty cases; lay groups, foundations, and individual practitioners six; commercial developers six, three of them as partners in university projects; and Buddhist orders, temples, or monastic authorities four. Three cases have no identifiable initiator. Scholarly inquiry was the most frequently coded purpose, followed by propagation, preservation of heritage, and education, then commercial distribution and pastoral care; three cases carried no stated purpose at all. No case was coded with institutional administration as its primary purpose, and one recorded it as a secondary purpose. This is consistent with the finding that no intersection in the administration domain reached observed status: the register contains no documented statement of institutional demand in that area.

Table 1. Initiating actor and coded purpose by performed role (n = 30).

	Instrumental	Relational	Religiously Authoritative	Counterfeit	Total
Initiating actor					
A1 Buddhist order, temple, or monastic authority	0	1	1	0	2
A1+A2 Buddhist institution with academic partner	0	1	1	0	2
A2 Academic or research institution	8	2	1	0	11
A2+A3 Academic institution with commercial partner	2	1	0	0	3
A3 Commercial developer or vendor	0	2	1	0	3
A4 Lay group, NGO, foundation, or individual practitioner	5	1	0	0	6
A5 No identifiable initiating actor	1	0	0	2	3
Coded purpose					
P0 None stated	1	0	0	2	3
P1 Propagation and public outreach	1	2	2	0	5
P3 Preservation of textual or material heritage	5	0	0	0	5
P4 Education and scholastic training	3	2	0	0	5
P5 Pastoral care and counseling	0	1	0	0	1
P7 Commercial distribution	0	2	1	0	3
P8 Scholarly inquiry	6	1	1	0	8
All cases	**16**	**8**	**4**	**2**	**30**

Counts are cases, not intersections; each case is counted once under its primary initiating actor, its primary coded purpose, and its predominant role. Joint actor codes are used only where the collaboration is documented in the initiators' own materials. Coded purpose is the aim given in the initiator's own documentation where one was available, and otherwise reconstructed from documented function; twenty of the thirty codes are of the first kind and ten of the second, and each is flagged individually in Supplementary Table S3. Purpose codes P2 (widening access) and P6 (institutional administration) appear only as secondary codes and are therefore absent from this table. The case-level listing, including evidence status, sources, and the evidence for each purpose assignment, is given in Supplementary Table S3. The inventory is not a representative sample, and the associations shown here are descriptive rather than causal.

6. From Instrument to Religious Authority

The 7 × 7 framework identifies where AI is being applied within Buddhist life, but technological modality alone does not explain the religious significance of an application. A conversational system, for example, may function as a search interface, a continuing practice companion, or an apparent interpreter of doctrine. Conversely, a technically simple and largely scripted robot may acquire religious significance through its appearance, institutional location, ritual positioning, and the meanings attributed to it by users. The central analytical distinction is therefore between what an AI system is technically capable of doing and what role it performs within a Buddhist context.

The distinction between capability and role also identifies the human choices that configure an application. These include where a system is placed, how it is named, whether generated content is disclosed, and when a difficult question is routed to a human teacher. Such decisions can shift a system from instrument to interlocutor or apparent authority, and identifiable people and institutions remain answerable for them.

The framework is designed to identify the points at which people decide what an AI system may do within a religious setting: what it is built to do, what it is authorized to say, whose standing it borrows, and what recourse exists when it errs. The categories locate these decisions and the corresponding responsibilities; they do not describe an autonomous technological development.

This study distinguishes three analytically ordered roles: instrumental, relational, and religiously authoritative. These categories concern performed and attributed roles rather than the ontological status of AI. They do not determine whether an AI system is conscious, possesses Buddha-nature, can attain enlightenment, or qualifies as a legitimate member of the saṅgha. Instead, the framework examines the scope of tasks delegated to AI, the character of human–AI interaction, the authority attributed to its outputs, and the institutional conditions under which those outputs are produced and received.

6.1. Instrumental AI

Instrumental AI performs bounded tasks defined by human users or institutions. Its outputs assist Buddhist activities, but the system is not positioned as an independent teacher, counselor, ritual actor, or doctrinal authority. Human actors determine the purpose of the task, select or approve the sources, evaluate the output, and retain final responsibility for decisions.

Most applications identified in the 7×7 matrix were predominantly instrumental. Examples include optical character recognition of Buddhist manuscripts, machine-assisted translation, semantic retrieval of scriptures, lecture transcription, preparation of draft educational materials, corpus analysis, visitor information, and administrative-document generation. These applications differ technically, but they share a common role structure: AI processes or transforms information within a task whose religious meaning and final interpretation remain under human control.

Instrumental status should not be confused with technological neutrality or low risk. A retrieval system can omit relevant scriptures, a translation system can erase doctrinal distinctions, and an administrative model can expose sensitive devotee data. An AI system may remain instrumental while producing significant epistemic, cultural, or institutional harm. Instrumentality describes the system's role, not the safety or moral acceptability of its operation.

Accordingly, the defining question for instrumental AI is not whether the technology is simple or advanced, but whether interpretive and religious responsibility remain visibly located in accountable human actors. The instrumental role begins to weaken when the system moves beyond bounded assistance and becomes a continuing interlocutor whose advice is personalized, emotionally significant, or behaviorally directive.

6.2. Relational AI

Relational AI emerges when a system functions not merely as a tool but as a continuing interlocutor. The defining features are continuity of interaction, responsiveness to the individual user, personalization, memory, perceived availability, and the possibility of emotional, pedagogical, or spiritual reliance. Relationality does not require that the AI possess consciousness or genuine emotion. It concerns the social position that the system occupies within the interaction.

A Buddhist question-answering system remains instrumental when it retrieves a passage in response to a discrete query. The same system moves toward a relational role when it remembers the user's prior questions, adapts explanations to the user's practice history, provides continuing encouragement, or initiates follow-up interactions. Similarly, an AI meditation application is instrumental when it supplies a timer or recording, but relational when it monitors practice, responds to reported difficulties, interprets emotional states, and presents itself as a persistent guide.

Relationality is analytically distinct from religious authority. A user may experience an AI system as supportive, familiar, or compassionate without accepting its statements as doctrinally binding. A counseling chatbot may be treated as a confidential conversational partner while remaining subordinate to a human counselor. An AI tutor may sustain a pedagogical relationship without possessing authority to resolve doctrinal disputes.

Nevertheless, relationality creates conditions under which authority can develop. Repeated successful interactions generate familiarity; personalization can produce perceived understanding; immediate availability can encourage dependence; and affective language can create intimacy and trust. Lim's concept of generative charisma is relevant here because it explains how AI may become compelling through generative performance, personalized intimacy, and oscillation between awareness of artificiality and experience of enchantment (Lim 2026). Relational systems therefore occupy a critical intermediate position: they do not necessarily exercise religious authority, but they may establish the trust through which later directives are accepted.

The defining issue is consequently not whether the system actually cares, understands, or possesses spiritual experience. The analytical question is whether the interaction is structured and received as an ongoing relationship and whether that relationship begins to influence the user's religious interpretation, emotional reliance, or conduct.

6.3. Religiously Authoritative AI

Religiously authoritative AI refers to an AI system that performs, mediates, or is received as performing functions conventionally associated with legitimate religious authority. The category does not claim that AI has become an independently legitimate Buddhist authority. It identifies situations in which a system's outputs or actions are treated as carrying privileged religious weight.

Religious authority may take several forms. Doctrinal authority concerns the ability to determine what constitutes a correct interpretation of Buddhist teaching. Spiritual or directive authority concerns the capacity to prescribe practice, evaluate progress, or advise users on religiously consequential decisions. Ritual authority concerns the recognized capacity to perform, lead, or validate a ritual act. Representative authority concerns the ability to speak or act on behalf of a Buddhist teacher, institution, lineage, or tradition. These dimensions may overlap, but they should not be treated as identical.

Mindar illustrates how religious authority can be mediated without high levels of technical autonomy. Its sermon is largely scripted, yet the android's embodiment as Kannon, its location in a Buddhist temple, its performance of a sermon, and its integration with visual and auditory media position it within a religiously charged encounter (Kodaiji Zen Temple 2019; Zheng 2024). Loewen-Colón and Mosurinjohn describe such encounters as spiritually authorizing when users attribute to a machine-mediated experience a claim to persuasion, arbitration, or legitimacy (Loewen-Colón and Mosurinjohn 2022).

The Korean robot Gabi provides an additional boundary case. Gabi's participation in a formal ceremony, reception of a Buddhist name, wearing of monastic clothing, and receipt of modified precepts introduced an explicit form of symbolic and institutional authorization. Its technical capacities were limited, but its religious positioning was compar-

atively strong. This divergence between technical capability and symbolic authorization supports the central claim that role intensity and technical sophistication are separate analytical variables (Korea.net 2026; Yonhap News Agency 2026).

By contrast, the institutional deployment of BuddhaBot-Plus in collaboration with Bhutan's Central Monastic Body illustrates controlled authorization without necessarily granting the system independent religious authority. The project included an official request, adaptation to a monastic context, monitoring by monks, and a planned safety assessment. Institutional involvement increases the perceived legitimacy of the system, but continued human monitoring can preserve its status as an authorized instrument rather than an autonomous teacher (Kyoto University 2025).

Religiously authoritative AI should therefore be understood as a role category based on performance, delegation, presentation, and reception. It does not settle whether the authority is doctrinally valid or normatively desirable. An AI system may be perceived as authoritative without being legitimately authorized, while an institutionally authorized system may be deliberately restricted from issuing authoritative judgments. The distinction between perceived authority and legitimate authority becomes particularly important in the case of counterfeit religious authority.

6.4. Counterfeit Religious Authority

Counterfeit religious authority is not a fourth stage following religiously authoritative AI. It is a separate and illegitimate pathway through which existing religious identity, trust, or charisma is fabricated, appropriated, or simulated without authorization. Whereas a legitimate role transition develops through disclosed interaction and some form of institutional or communal negotiation, counterfeit authority begins by borrowing the appearance of authority while concealing the actual source of the content.

Illegitimacy here is neither a doctrinal verdict nor an external judgment about religious authenticity. It is a matter of consent and provenance, assessed from the position of the person or institution whose identity is used. A video is counterfeit because the teacher whose voice it reproduces did not authorize it and cannot correct it, and because the community that would ordinarily verify such a teaching was bypassed—not because its content is unorthodox. The question of illegitimacy to whom therefore has a definite answer: to the party whose standing was taken, and to those who received the material believing it came from that party.

This restriction matters because a charge of inauthenticity is a powerful instrument. Applied loosely, it could be turned against new movements, against lay-led and diaspora communities, or against reform interpretations, by parties with an interest in defending an established position. The criterion proposed here is designed to be unavailable for that purpose. It asks whether a specific identity was appropriated without authorization and its origin concealed, and it is indifferent to whether the resulting teaching is conventional. A novel or heterodox use of AI that is openly presented as such is not counterfeit, however unwelcome some may find it. An entirely orthodox teaching falsely attributed to a living teacher is.

Counterfeit religious authority has four defining features. First, it appropriates an identifiable religious person, institution, scripture, ritual position, or tradition. Second, the appropriation occurs without valid authorization or consent. Third, the synthetic or altered origin of the content is concealed or inadequately disclosed. Fourth, the content is presented in a form likely to be received as an authentic teaching, statement, or religious act.

The distinction between hallucination and counterfeit authority is also important. A disclosed chatbot may produce an inaccurate quotation because of model error. This con-

stitutes an epistemic failure. Counterfeit authority occurs when false content is deliberately or functionally attached to an identity or source that gives it undeserved religious credibility. The harm lies not only in the falsity of the statement but also in the unauthorized transfer of accumulated trust.

Imitation is not in itself counterfeit. Buddhist traditions have long reproduced the voice, image, and words of teachers through recitation, statuary, the voicing of a teacher's words in commentary, and recorded talks. What distinguishes counterfeit authority is the combination of concealed origin, absent authorization, and presentation in a form likely to be received as an authentic act of teaching. Impersonation of a deceased master is one form; fabricated institutional statements, forged passages presented as canonical, and synthetic endorsements attributed to living teachers are others. What they share is not resemblance but the transfer of trust that no one granted.

Unauthorized use of a canonical corpus is a related but distinct matter. It is in the first instance a question of copyright and provenance, addressed in Section 7.4. It becomes counterfeit religious authority only where the resulting system also conceals its origin and presents its outputs as authentic statements of the tradition or teacher whose corpus was taken.

Documented Korean cases show that this pathway is already empirically visible. AI-generated videos have used the images and voices of the late Seon master Seongcheol and the monk and writer Beopjeong to present fabricated teachings as genuine Dharma talks (BTN News 2026; BBS News 2026). Other videos imitated the voice and appearance of Ven. Pomnyun and attributed apocalyptic predictions and folk-religious instructions to him. Pomnyun and the Jungto Society publicly warned audiences that the materials were synthetic and requested that suspected videos be reported for verification and removal (Beopbo News 2026; Jungto Society 2026).

Counterfeit authority must consequently be governed differently from authorized instrumental, relational, or religiously authoritative applications. The primary response is not merely improved doctrinal review but provenance authentication, explicit disclosure, consent protection, rapid reporting, platform removal, preservation of authentic archives, and legal protection of voice, image, name, and institutional identity. These measures are developed in Sections 7.3 and 7.6.

6.5. Role Transitions and Boundary Conditions

The three legitimate roles form an analytical progression, but they do not constitute an inevitable technological-development sequence. AI does not automatically move from instrument to relationship and then to religious authority as it becomes more advanced. Role transitions depend on changes in task structure, interaction, presentation, institutional delegation, and user reception. They may be gradual or abrupt, reversible or contested, and different users may attribute different roles to the same system.

The first transition—from instrumental to relational—occurs when bounded task support becomes continuing interaction. Memory, personalization, proactive communication, affective language, and availability are especially significant. A scripture-retrieval chatbot remains instrumental when it returns passages and citations. It becomes relational when it remembers the practitioner's history, asks about progress, adapts advice, and is treated as a continuing practice companion.

The second transition—from relational to religiously authoritative—occurs when trust and continuity become attached to the system's ability to interpret, direct, validate, or represent. A practice companion approaches religious authority when it tells users which teaching is correct, prescribes a religious discipline as obligatory, evaluates spiritual attainment, or speaks as an authorized representative of a lineage. The transition is strength-

ened when the system bears a religious title, embodies a sacred or monastic form, operates within a temple, or is formally endorsed by a Buddhist institution.

BuddhaBot and BuddhaBot-Plus illustrate how changes in system function can alter role potential. The original retrieval-oriented system was designed to return Buddhist passages, making its primary role instrumental. The addition of generative explanation enables more adaptive dialogue and therefore creates greater potential for relational interaction. Its deployment within a monitored monastic setting may increase legitimacy while simultaneously constraining the system's authority through human oversight (Kyoto University 2025). The same technical system can therefore be relationally engaging yet institutionally maintained as an instrument.

Mindar demonstrates the opposite configuration. Its computational autonomy is limited, but its embodied representation of Kannon, sermon format, and temple setting may create a strongly authorizing encounter. Gabi similarly shows that symbolic and ritual recognition can produce an authoritative role before the technology possesses advanced dialogue or decision-making capabilities. These cases establish that religious authority is not a direct function of model sophistication.

Counseling and practice systems present a different boundary. They may become deeply relational without acquiring legitimate religious authority. A user may disclose distress, seek encouragement, and interact daily with a Buddhist chatbot while still understanding that the system is not competent to interpret doctrine or authorize practice. The critical governance task is therefore to prevent relational trust from being converted silently into unwarranted authority.

The comparison supports six propositions: religious significance is shaped more by performed role and context than by modality; repetition, personalization, disclosure, and responsiveness encourage relational transition; relational trust becomes authority when AI interpretations or directives receive privileged religious weight; technical sophistication and authority intensity are independent; one system may occupy different roles across users and institutions; and counterfeit authority is a divergent pathway based on unauthorized appropriation. These propositions describe possible and reversible transitions, not an inevitable development sequence.

Read against Table 1, coded purpose and performed role are associated. All five cases coded primarily for the preservation of heritage were instrumental, as were six of the eight coded for scholarly inquiry: where the declared aim was to recover, analyze, or evaluate a corpus, the system remained a tool under human correction. Among the five cases coded primarily for propagation, two were religiously authoritative, two relational, and one instrumental. No case coded for commercial distribution was instrumental; two were relational conversational applications and the third an embodied ritual performer whose ritual claim its documented adoption did not support. Because ten of the thirty purpose codes were reconstructed from documented function rather than stated by the initiator, and because the inventory is not representative, these patterns are descriptive rather than causal.

The four cases classified as religiously authoritative reached that classification by three routes: institutional delegation, in which an order or temple conferred a name, a robe, a placement, or a ceremony; commercial substitution, in which a vendor proposed a machine as a lower-cost officiant; and individual experiment, in the composition of a fictional scripture. Of the four, only the two conferred by institutional delegation are recorded in the register as being in continuing operational or ceremonial use; the remaining two are documented as a demonstration and an experiment respectively. Role is therefore not well explained by technical modality alone. It is associated with what people set out to do, and with whether an institution lent its standing to the result.

7. Risks and Governance

Risk arises from the interaction of capability, attributed role, institutional setting, affected users, data, and distribution. The study distinguishes doctrinal and epistemic, authority and authenticity, relational and psychological, privacy, bias and exclusion, copyright and heritage, deception and misinformation, and accountability and safety risks. A conceptual heuristic combining role intensity, autonomy, user vulnerability, system opacity, and scale clarifies why a technically modest system may require strict controls when it influences vulnerable users or carries broad religious authority.

7.1. Doctrinal and Epistemic Risk

Generative systems can fabricate quotations, flatten doctrinal differences, and synthesize incompatible sources, while retrieval systems remain dependent on corpus selection, metadata, ranking, translation, and model use of the retrieved passages (Bender et al. 2021; Lewis et al. 2020; Weidinger et al. 2022). Buddhist corpora are multilingual, historically layered, and internally plural; source citation alone cannot determine which tradition is represented or who may resolve disagreement. Governance should make sources, uncertainty, textual versions, and interpretive scope visible and preserve referral to qualified human teachers or scholars.

7.2. Relational and Psychological Risk

Persistence, personalization, affective language, and availability can expand access while creating dependence, misplaced trust, excessive disclosure, or substitution for human support. These risks are acute in counseling and practice systems used by children, older adults, bereaved persons, or people in crisis. Relational AI should disclose its status and competence, avoid dependency-inducing design, minimize retention of intimate data, and provide reliable transfer to teachers, counselors, caregivers, or emergency services. Evaluation must examine patterns of reliance and failed escalation as well as content accuracy (Laestadius et al. 2024; World Health Organization 2024).

7.3. Authenticity and Counterfeit Authority

Authenticity concerns whether a message, voice, image, ritual performance, or institutional statement is genuine, authorized, modified, or synthetic. Korean deepfake cases appropriating the faces and voices of prominent monks show that counterfeit authority exploits trust accumulated by human teachers rather than developing authority through a disclosed AI relationship. Persistent disclosure, consent with a defined scope, authenticated official channels, provenance records, rapid reporting and removal, public correction, and preservation of genuine archives should be combined. Transparent AI-assisted restoration or translation must be distinguished from deceptive substitution.

7.4. Data, Copyright, and Cultural Heritage

Practice histories, counseling disclosures, donations, ritual requests, and temple attendance may reveal sensitive religious, health, and family information. Data minimization, purpose limitation, meaningful consent, encryption, retention limits, and control of third-party providers are therefore necessary. Copyright and cultural obligations also apply to scriptures, translations, sermons, images, music, and scans. Heritage projects should preserve source files, metadata, textual variants, uncertainty, and reversible editorial decisions, distinguish reconstruction from surviving evidence, and involve relevant custodial communities (UNESCO 2015; World Intellectual Property Organization 2024).

7.5. Embodiment, Ritual Status, and Safety

Robes, sacred iconography, temple placement, voice, gaze, and gesture can imply authority beyond a robot's technical function. Institutions should disclose whether an embodied system is an exhibit, interface, assistant, representative, or ritual participant and separate technical capability from contested religious efficacy. A dual-authorization rule is required: technical authorization establishes physical fitness and safety, while religious authorization defines the symbolic, pedagogical, representative, or ritual role. Site-specific risk assessment, operating boundaries, emergency stops, failure testing, cybersecurity, and trained supervision remain necessary in public and care environments.

7.6. A Role- and Risk-Based Governance Framework

Governance controls escalate with role intensity and contextual risk. Every application requires a defined purpose and accountable owner, documentation of models and data, privacy and security controls, pre-deployment testing, monitoring, incident correction, accessibility, and a suspension or decommissioning mechanism. Table 2 differentiates the additional posture and controls associated with instrumental, relational, religiously authoritative, and counterfeit roles.

For instrumental AI, the default posture is controlled adoption. The system may be used when its task is bounded, its sources and data practices are documented, and a competent human retains final judgment. Required controls include source citation, accuracy checks, privacy protection, copyright review, version control, and correction procedures. The central governance principle is that efficiency must not obscure the location of interpretive responsibility.

For relational AI, the posture is conditional adoption. In addition to instrumental controls, the system must disclose its artificial status and role, limit anthropomorphic and dependency-inducing design, protect intimate disclosures, and provide reliable transition to human support. Institutions should test not only whether answers are accurate but whether patterns of use create overreliance, manipulation, exclusion, or avoidance of human relationships. Relational systems should never silently expand their role from companionship or preliminary support into counseling, doctrinal direction, or crisis management.

For religiously authoritative AI, the posture is exceptional and restrictive. Deployment requires explicit specification of the tradition, institution, and scope represented; formal doctrinal review; community deliberation; identifiable human and institutional responsibility; transparent authorization; logs and auditability; appeal and correction; and strict limits on autonomous judgment. AI should not independently confer monastic or ritual status, declare spiritual attainment, resolve serious ethical disputes, or present generated teachings as binding without accountable human authorization. Institutional approval may legitimate a bounded representative function, but it does not transfer ultimate responsibility to the system.

Counterfeit religious authority is not a high-risk service to be cautiously optimized; it is an unauthorized practice to be prevented and disrupted. Required measures include provenance verification, consent and identity protection, detection, reporting, rapid removal, public correction, preservation of authentic records, and legal or platform remedies. Exceptions for disclosed satire, scholarship, artistic work, or archival reconstruction should be clearly labeled and should not create reasonable confusion about endorsement or authenticity.

Table 2. Role- and risk-based governance framework for Buddhist AI.

Role Tier	Default Posture	Primary Risks	Mandatory Controls	Escalation or Prohibition Trigger
Instrumental AI	Controlled adoption	Doctrinal error; privacy; bias; copyright; operational error	Bounded task; source visibility; accuracy testing; data minimization; copyright review; human final judgment; correction and incident process	Suspend or restrict when outputs cannot be verified, sensitive data cannot be protected, or humans cannot meaningfully review consequences
Relational AI	Conditional adoption	Dependence; manipulation; excessive disclosure; substitution for human care; role drift	All instrumental controls plus AI and role disclosure; limits on anthropomorphic persuasion; human referral and crisis escalation; dependency monitoring; stronger safeguards for vulnerable users	Suspend or redesign when the system encourages exclusivity, obstructs human help, exceeds its declared role, or creates persistent psychological harm
Religiously Authoritative AI	Exceptional and restrictive adoption	Doctrinal misrepresentation; illegitimate delegation; ritual confusion; accountability gaps; concentrated authority	Explicit institutional authorization; defined lineage and scope; doctrinal and communal review; audit logs; appeal and correction; authority limits; identifiable human responsibility	Do not deploy autonomous doctrinal adjudication, spiritual-status determination, or binding ritual and ethical direction without accountable human authorization
Counterfeit Religious Authority	Prohibit and disrupt	Impersonation; fraud; fabricated teaching; reputational harm; erosion of trust	Consent and identity protection; provenance; detection; reporting; rapid removal; public correction; authentic archives; legal and platform remedies	Any unauthorized simulation or appropriation reasonably likely to be understood as genuine religious speech, identity, or endorsement

Governance responsibility is distributed but not diffuse. Developers are responsible for system design, documentation, testing, security, and technical limitations. Deploying Buddhist institutions are responsible for selecting the use case, defining religious boundaries, validating sources, protecting users, and maintaining human oversight. Monastics, scholars, counselors, archivists, and ritual specialists are responsible for domain review within their competence. Platforms are responsible for responding to impersonation, fraud, and harmful distribution. Users retain agency, but user responsibility cannot be used to excuse foreseeable design or institutional failures.

A practical decision process follows six questions. First, what role is the system performing and what role are users likely to attribute to it? Second, what religious, psychological, cultural, or physical harm could result? Third, which users or communities are especially vulnerable or underrepresented? Fourth, can sources, data, authorization, and decisions be inspected and contested? Fifth, who can intervene, correct, or stop the system? Sixth, does the proposed use contribute to the reduction of suffering, wisdom, compassion,

non-attachment, and responsible human practice, or does it primarily expand automation, engagement, and institutional convenience?

The framework is therefore neither a universal authorization scheme nor a claim that all Buddhist traditions will draw identical boundaries. It is a structured method for making disagreements visible and for matching governance intensity to the role and risks of a specific application. The following Discussion examines how this role-based account extends digital-religion research, generative charisma, and Buddhist approaches to skillful means and human responsibility.

8. Discussion

8.1. Theoretical Contributions

The study contributes a multidimensional account of Buddhist AI that links religious domain, technical modality, evidence status, performed role, risk, and governance. Its central theoretical move is to treat role rather than technology as the primary unit for interpreting authority. The same model or robot can occupy different roles across contexts, while technically different systems can perform equivalent religious functions. This avoids equating fluency, embodiment, or autonomy with legitimate religious standing.

Authority in this account is a structural position rather than an ontological property: the position from which speech is received as binding, held by whoever occupies it, on whatever grounds a community accepts. AI may be placed in that position through institutional endorsement, ritual or spatial placement, religious naming, embodiment, repeated and personalized interaction, and user reception. These mechanisms depend on human and institutional decisions, although they are also shaped by the system's technical affordances, and each of them can be withheld.

The framework extends digital-religion research by specifying which authority function is delegated, represented, or attributed. Generative and conversational AI do not merely carry archived content; they produce context-sensitive language within an encounter. Embodied systems add material and spatial cues. Authority is consequently configured through output, design, institutional framing, authorization, and user reception rather than located solely in a person, text, institution, or machine (Campbell 2007, 2021).

The separation of evidentiary maturity from normative acceptability is equally important. Observed status indicates verifiable implementation or evaluation, not effectiveness or desirability. Counterfeit religious authority adds a distinct contribution: unlike negotiated or institutionally mediated authority, it combines false provenance, identity misuse, and unauthorized transfer of charismatic capital. Its governance logic is prevention and disruption rather than cautious adoption.

8.2. Skillful Means and Human Responsibility

Buddhist discourse may describe technologies as skillful means when they expand access or adapt communication to an audience, but efficiency and engagement are not sufficient. Evaluation must include intention, consequences, context, discernment, attachment, non-harm, and responsibility. Interpretation, listening, care, and communal practice may be constitutive forms of ethical and religious formation rather than inefficiencies to automate (Hershock 2021; Compson et al. 2025).

The distribution of agency set out in Section 2.3 has a direct bearing on this evaluation. Because a religious outcome is produced through a chain of decisions rather than by a system acting alone, responsibility is apportioned along that chain in proportion to control and competence, and human oversight must remain plural, documented, contestable, and capable of stopping a system. A modest bounded tool may be skillful when it expands access while preserving sources and human judgment. A technically impres-

sive system may be unskillful when it creates false certainty, dependency, exclusion, or counterfeit authority.

8.3. Limitations and Future Research

The study is based on literature and public documentation rather than direct observation of practitioners, monastics, developers, or institutions. Role assignments are analytical interpretations, not measured accounts of every user's experience. English- and Korean-language materials predominate, and Chinese, Japanese, Tibetan, Pali, Sinhala, Thai, Vietnamese, and other sources may be underrepresented. Public visibility and documentation also vary across countries and organizations.

The evidence base also represents a particular Buddhism. Documentation is concentrated in Korean and Japanese institutional settings, in internationally visible temples, and in well-resourced translation and archival projects; lay, monastic, institutional, and scholarly actors within those settings do not share a single view of what AI may legitimately do, and communities that document their activities less publicly are correspondingly absent. The mapping is accordingly an analytical interpretation of documented material, not a report of how developers, institutions, or practitioners describe their own activity. Where the two might diverge—most plainly in the assignment of purpose and role—the coding records the basis for each judgment so that it can be contested.

The field changes rapidly, and applications may be redesigned, discontinued, or newly deployed. Although full-text verification and source consolidation were completed for the present evidence lock, official and journalistic sources cannot substitute for independent ethnography, technical audit, or outcome evaluation. The 7×7 structure remains an analytical heuristic: modalities overlap, applications serve multiple domains, roles change, and the conceptual risk modifiers are not validated numerical weights.

Future work should test role boundaries and authority cues through interviews, ethnography, scenario experiments, longitudinal studies, and cross-traditional comparison. Priority questions include how repeated personalization affects trust and dependence, how institutional endorsement changes legitimacy, whether source-grounded dialogue and human review improve doctrinal reliability, and whether the proposed governance tiers are usable in real institutional decisions. Counterfeit authority also requires collaborative research on provenance, reporting, correction, and platform response.

9. Conclusions

This study structured a rapidly expanding and uneven field by linking seven domains of Buddhist life with seven AI modalities, producing forty-nine analytical intersections. The matrix was not used to imply that all intersections are equally established. Instead, each was differentiated as observed, emerging, or speculative and connected to a transparent evidence trail. This architecture integrates forms of application that are usually discussed separately—including textual retrieval, translation, education, practice support, counseling, heritage work, administration, transnational communication, generative content, and embodied systems—while preserving differences in purpose, institutional setting, evidentiary maturity, and Buddhist context.

The most important theoretical finding is that the religious significance of AI cannot be inferred from technical type or computational sophistication alone. AI roles are configured through use, personalization, repeated interaction, embodiment, naming, institutional endorsement, user attribution, and the degree of authority delegated to the system. The same technology can remain instrumental, become relational, or approach religious authority. Embodiment may intensify presence but is neither necessary nor sufficient for authority, while a highly capable generative system may remain subordinate

under source constraints and human review. Counterfeit religious authority requires separate treatment because it does not represent a legitimate progression along the role continuum; it appropriates established identity and charisma through concealed or unauthorized synthetic representation.

The case register also shows how human purposes and institutional decisions shaped the roles assigned to these systems. All five preservation cases and six of the eight scholarly-inquiry cases were instrumental. Propagation and commercial distribution produced more varied configurations, including relational and religiously authoritative roles. These patterns are descriptive rather than causal, but they identify the decisions for which developers, institutions, and communities are answerable.

The corresponding governance principle is proportionality between the role performed, the risks created, and the strength of oversight. Instrumental applications can be adopted under requirements for accuracy, source visibility, privacy, copyright compliance, and human final judgment. Relational applications require explicit AI disclosure, clear role boundaries, safe transfer to qualified humans, protection of sensitive data, and safeguards against dependency or substitution for community. Applications that communicate or perform religious authority require doctrinal verification, community deliberation, formal authorization, auditable responsibility, contestability, and strict limits on the decisions and representations that may be delegated. Unauthorized impersonation and fabricated religious authority should be treated as a matter for detection, reporting, removal, correction, and rights protection rather than as an acceptable form of religious innovation. Across all levels, AI should remain a means whose legitimacy depends on accountable human and communal judgment.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/rel17091069/s1. Table S1: master register of Korean and international scholarship; Table S2a: exact international search queries and procedures; Table S2b: screening counts and corpus composition; Table S2c: title/abstract and full-text exclusions; Table S3: verified Buddhist–AI case inventory; Table S4: final evidence, role, risk, and status coding of all intersections; shared codebook: domain, modality, role, risk, status, grade, initiating-actor, and coded-purpose definitions.

Author Contributions: Conceptualization, W.G. and C.S.; methodology, W.G.; formal analysis, W.G.; writing—original draft preparation, W.G.; writing—review and editing, C.S.; supervision, C.S. 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: The literature register, search and screening log, verified case inventory, and final 7 × 7 evidence matrix are included in the Supplementary Materials submitted with this article.

Acknowledgments: During preparation of this manuscript, the authors used ChatGPT (GPT-5; OpenAI, San Francisco, CA, USA) to assist with structural organization, preliminary drafting, language refinement, and visualization planning. The authors reviewed and edited the output, verified the supporting evidence and references, and take full responsibility for the content.

Conflicts of Interest: The authors declare no conflicts of interest.

References

Akrich, Madeleine. 1992. The De-Scription of Technical Objects. In *Shaping Technology/Building Society: Studies in Sociotechnical Change*. Edited by Wiebe E. Bijker and John Law. Cambridge: MIT Press, pp. 205–24.

Baffelli, Erica. 2021. The Android and the Fax: Robots, AI and Buddhism in Japan. In *Itineraries of an Anthropologist: Studies in Honour of Massimo Raveri*. Edited by Giovanni Bulian and Silvia Rivadossi. Venice: Edizioni Ca' Foscari, pp. 249–64. [CrossRef] [PubMed]

BBS News. 2026. AI Mannan Bulgyo 3: Musok Gareuchin Beopjeong Seunim? Panchineun 'Gajja Bulgyo' Makaya AI 만난 불교③: 무속 가르친 법정스님? 판치는 '가짜 불교' 막아야 [AI Meets Buddhism 3: Did Ven. Beopjeong Teach Shamanism? Stopping the Rampant Spread of "Fake Buddhism"]. Available online: https://news.bbsi.co.kr/news/articleView.html?idxno=4078636 (accessed on 5 August 2026). (In Korean)

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Paper presented at the 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event, March 3–10; pp. 610–23. [CrossRef]

Beopbo News. 2026. Beomnyun Seunim Beopmun Gajanghan AI Yeongsang Yupo: Bulgyogye-edo 'Dippeikeu Pihae' Sarye Deungjang 법륜 스님 법문 가장한 AI 영상 유포: 불교계에도 '딥페이크 피해' 사례 등장 [AI Videos Disguised as Ven. Pomnyun's Dharma Talks Spread: Deepfake Damage Reaches the Buddhist Community]. Available online: https://www.beopbo.com/news/articleView.html?idxno=334386 (accessed on 5 August 2026). (In Korean)

BTN News. 2026. "31-man Myeong Soatda": Seongcheol·Beopjeong Seunim AI Sacheong Yeongsang Giseung "31만 명 속았다": 성철·법정스님 AI 사칭 영상 기승 ["310,000 People Deceived": AI Impersonation Videos of Vens. Seongcheol and Beopjeong Proliferate]. Available online: https://www.btnnews.tv/news/articleView.html?idxno=90379 (accessed on 5 August 2026). (In Korean)

Buddhist Digital Resource Center. 2025. BDRC Announces the Release of OCR App for Tibetan. Available online: https://www.bdrc.io/blog/2025/03/14/bdrc-announces-the-release-of-ocr-app-for-tibetan/ (accessed on 5 August 2026).

Campbell, Heidi A. 2007. Who's Got the Power? Religious Authority and the Internet. *Journal of Computer-Mediated Communication* 12: 1043–62. [CrossRef]

Campbell, Heidi A. 2010. *When Religion Meets New Media*. New York and London: Routledge.

Campbell, Heidi A. 2021. *Digital Creatives and the Rethinking of Religious Authority*. Abingdon and New York: Routledge.

Cheong, Pauline Hope. 2020. Religion, Robots and Rectitude: Communicative Affordances for Spiritual Knowledge and Community. *Applied Artificial Intelligence* 34: 412–31. [CrossRef]

Compson, Jane, Mark Graves, Peter D. Hershock, and Nikki Mirghafori. 2025. A Middle Path for AI Ethics? Some Buddhist Reflections. *Theology and Science* 23: 1–5. [CrossRef]

Gil, Wanje, Insoo Shin, Kwanghyun Han, and Hwanyoung Jang. 2025. Ingongjineung-eul Baraboneun Bulgyo-ui Siseon: Seukoping Ribyu 인공지능을 바라보는 불교의 시선: 스코핑 리뷰 [A Buddhist Perspective on Artificial Intelligence: A Scoping Review]. *Won-Buddhist Thought and Religious Culture* 원불교사상과 종교문화 106: 209–35. (In Korean)

Hershock, Peter D. 2021. *Buddhism and Intelligent Technology: Toward a More Humane Future*. London and New York: Bloomsbury Academic.

Hongladarom, Soraj. 2020. *The Ethics of AI and Robotics: A Buddhist Viewpoint*. Lanham: Lexington Books.

Jaakkola, Elina. 2020. Designing Conceptual Articles: Four Approaches. *AMS Review* 10: 18–26. [CrossRef]

Jabareen, Yosef. 2009. Building a Conceptual Framework: Philosophy, Definitions, and Procedure. *International Journal of Qualitative Methods* 8: 49–62. [CrossRef]

Jungto Society. 2026. AI-ga Mandeureonaen Yeongsangmul Eodikkajiga Jinjja-ilkkayo? AI가 만들어낸 영상물 어디까지가 진짜일까요? [How Much of AI-Generated Video Is Real?]. Available online: https://www.jungto.org/pomnyun/view/85690 (accessed on 5 August 2026). (In Korean)

Kim, Myeongjun. 2024. Bulgyo Sangdam Giban-ui Chaetbot Gaebal-eul Wihan Gicho Yeongu: Cheongjeong Ma-eumbot Guhyeon-eul Wihan Gicho Seolgye mit Gaeseonsahang Dochul 불교 상담 기반의 챗봇 개발을 위한 기초 연구: 청정마음봇 구현을 위한 기초 설계 및 개선사항 도출 [A Preliminary Study for Developing a Buddhist Counseling Chatbot: Basic Design and Improvement Requirements for Cheongjeong MaeumBot]. *Korean Journal of Buddhist Studies* 한국불교학 112: 77–119. (In Korean)

Kodaiji Zen Temple. 2019. Press Announcement: The Unveiling of and Delivery of a Buddhist Sermon by Mindar—The Android Kannon (Buddha of Compassion). Available online: https://www.kodaiji.com/mindar/press_data/data01.pdf (accessed on 5 August 2026).

Korea.net. 2026. Robot Buddhist Monk Gabi. Available online: https://www.korea.net/NewsFocus/Korea_in_photos/view?articleId=291865 (accessed on 5 August 2026).

Kyoto University. 2025. BuddhaBot, Buddhist Dialogue AI, Marks Its First Overseas Release in Collaboration with Bhutan's Central Monastic Body. Available online: https://www.kyoto-u.ac.jp/en/news/2025-03-03 (accessed on 5 August 2026).

Laestadius, Linnea, Andrea Bishop, Michael Gonzalez, Diana Illenčík, and Celeste Campos-Castillo. 2024. Too Human and Not Human Enough: A Grounded Theory Analysis of Mental Health Harms from Emotional Dependence on the Social Chatbot Replika. *New Media & Society* 26: 5923–41. [CrossRef]

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and et al. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems* 33: 9459–74.

Li, Denghao, Yuqiao Zeng, Jianzong Wang, Lingwei Kong, Zhangcheng Huang, Ning Cheng, Xiaoyang Qu, and Jing Xiao. 2022. Blur the Linguistic Boundary: Interpreting Chinese Buddhist Sutra in English via Neural Machine Translation. Paper presented at the 2022 IEEE 34th International Conference on Tools with Artificial Intelligence, Macao, China, October 31–November 2; pp. 228–32. [CrossRef]

Lim, Francis Khek Gee. 2026. AI and Generative Charisma in Religious Practices. *Religions* 17: 549. [CrossRef]

Loewen-Colón, Jordan, and Sharday C. Mosurinjohn. 2022. Fabulation, Machine Agents, and Spiritually Authorizing Encounters. *Religions* 13: 333. [CrossRef]

Nehrdich, Sebastian, Marcus Bingenheimer, Justin Brody, and Kurt Keutzer. 2023. MITRA-zh: An Efficient, Open Machine Translation Solution for Buddhist Chinese. In *Paper presented at the Joint 3rd International Conference on Natural Language Processing for Digital Humanities and 8th International Workshop on Computational Linguistics for Uralic Languages, Tokyo, Japan, December 1–3*. Tokyo: Association for Computational Linguistics, pp. 266–77.

Promta, Somparn, and Kenneth Einar Himma. 2008. Artificial Intelligence in Buddhist Perspective. *Journal of Information, Communication and Ethics in Society* 6: 172–87. [CrossRef]

Reuters. 2017. In Japan, Robot-for-Hire Programmed to Perform Buddhist Funeral Rites. Republished by Investing.com. Available online: https://www.investing.com/news/technology-news/in-japan-robotforhire-programed-to-perform-buddhist-funeral-rites-521440 (accessed on 5 August 2026).

Snyder, Hannah. 2019. Literature Review as a Research Methodology: An Overview and Guidelines. *Journal of Business Research* 104: 333–39. [CrossRef]

Torraco, Richard J. 2005. Writing Integrative Literature Reviews: Guidelines and Examples. *Human Resource Development Review* 4: 356–67. [CrossRef]

Travagnin, Stefania. 2020. From Online Buddha Halls to Robot-Monks: New Developments in the Long-Term Interaction between Buddhism, Media and Technology in Contemporary China. *Review of Religion and Chinese Society* 7: 120–48. [CrossRef]

Tsadra Foundation. n.d.AI Tools for Tibetan Buddhist Translation. Available online: https://www.tsadra.org/ai-tools-for-tibetan-buddhist-translation/ (accessed on 5 August 2026).

UNESCO. 2015. *Recommendation Concerning the Preservation of, and Access to, Documentary Heritage Including in Digital Form*. Paris: UNESCO.

Verbeek, Peter-Paul. 2011. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press.

Weidinger, Laura, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, and et al. 2022. Taxonomy of Risks Posed by Language Models. Paper presented at the 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21–24; pp. 214–29. [CrossRef]

Whittemore, Robin, and Kathleen Knafl. 2005. The Integrative Review: Updated Methodology. *Journal of Advanced Nursing* 52: 546–53. [CrossRef] [PubMed]

World Health Organization. 2024. *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*. Geneva: World Health Organization.

World Intellectual Property Organization. 2024. *Generative AI: Navigating Intellectual Property*. Geneva: World Intellectual Property Organization.

Yin, Robert K. 2018. *Case Study Research and Applications: Design and Methods*, 6th ed. Thousand Oaks: Sage.

Yonhap News Agency. 2026. Buddhist Sect Welcomes Humanoid Robot Gabi with Precept Ceremony. Available online: https://en.yna.co.kr/view/AEN20260506006000315 (accessed on 5 August 2026).

Zheng, Yutong. 2024. Buddhist Transformation in the Digital Age: AI (Artificial Intelligence) and Humanistic Buddhism. *Religions* 15: 79. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.