
DharmaBot: An Intelligent Conversational Agent for Hindu Studies Education Based on Large Language Model and Global Hindu Knowledge Base

I Wayan Sindia Griya Danika¹,

¹Institut Agama Hindu Negeri Tampung Penyang Palangka Raya

dana.danikadas@gmail.com

ABSTRACT

The rapid advancement of artificial intelligence, particularly Large Language Models, has created significant opportunities for transforming academic support systems in religious higher education. Hindu studies education in Indonesia faces persistent challenges in providing accessible and doctrinally accurate consultation beyond classroom hours, particularly in institutions with limited instructional resources such as Institut Agama Hindu Negeri Tampung Penyang Palangka Raya. Students frequently encounter knowledge gaps spanning classical scriptures, philosophical schools, ethical frameworks, and global Hindu practice that instructors cannot address within formal class hours alone. This study presents DharmaBot, an intelligent conversational agent designed to support Hindu Dharma consultation for students in Hindu higher education institutions. Employing a Design Science Research methodology through four stages of problem identification, objective definition, system design, and prototype demonstration, the study produced a system integrating a Large Language Model engine with a Retrieval-Augmented Generation pipeline and a structured Hindu knowledge base covering eight thematic domains including classical philosophy, sacred scriptures, theology, ethics, ritual practice, global Hinduism, and Hindu-science intersections. The system delivers bilingual responses in Indonesian and English with transparent source attribution drawn from verified Hindu scholarly documents. Three demonstration scenarios confirmed the prototype's capacity for accurate, reference-grounded responses to philosophical, theological, and comparative Hindu studies queries. DharmaBot contributes a replicable system design for AI-assisted religious education and offers practical implications for the responsible digital transformation of Hindu higher education institutions in Indonesia and globally.

Keywords : DharmaBot, Large Language Model, Hindu Education, Conversational Agent, Design Science Research

I. Introduction

The integration of artificial intelligence into educational environments has accelerated substantially following the public release of generative AI tools in 2022. Large Language Models now demonstrate capabilities that extend well beyond information retrieval, offering contextual dialogue, explanatory reasoning, and

personalized academic support across diverse disciplines (Yigci et al., 2025). This shift has prompted higher education institutions worldwide to reconsider how AI can serve as a substantive academic support mechanism rather than merely a supplementary tool (Ocen et al., 2025). The global expansion of conversational AI in education has reached disciplines previously considered resistant to technological mediation, including philosophy, ethics, and religious studies.

Religious education presents a distinctive and demanding context for AI integration. Unlike technical disciplines, religious studies require contextual sensitivity, doctrinal accuracy, and culturally grounded interpretation that standard AI systems are not designed to deliver. Hindu studies education encompasses a rich and multifaceted knowledge corpus spanning thousands of years of philosophical development, from the Vedas and Upanishads to the Bhagavad Gita and the diverse traditions of global Hindu practice. Students in Hindu higher education institutions frequently encounter knowledge gaps that extend far beyond what instructors can address within formal class hours, particularly when scriptural consultation and philosophical clarification are needed in real time.

In Indonesia, Hindu higher education is primarily administered through state-recognized institutions known as Institut Agama Hindu Negeri, serving students across Bali, Kalimantan Tengah, and other regions with significant Hindu communities. These institutions operate within constraints that limit the availability of consultative academic support, including restricted faculty availability outside scheduled hours, geographic isolation in several campuses, and a scarcity of accessible scholarly resources in Indonesian-language Hindu studies. While the digital transformation of education in Indonesia has gained institutional momentum, the digital divide documented by Aditya et al. (2022) continues to present barriers to equitable technology adoption across Indonesian higher education, particularly in non-urban campuses.

The emergence of conversational AI agents offers a viable pathway to bridge this consultative gap. Chatbot systems powered by LLMs have demonstrated measurable benefits in student engagement, academic support, and accessible knowledge delivery across multiple higher education contexts (Kuhail et al., 2023).

A student consulting a well-designed AI agent can receive immediate, reference-grounded answers to philosophical questions at any hour, whereas without such a system the same student might wait days for an instructor response or rely on unverified online sources. The integration of Retrieval-Augmented Generation (RAG) techniques further enables AI systems to deliver responses grounded in specific knowledge corpora, a capability particularly valuable when doctrinal accuracy in scriptural interpretation carries institutional significance (Lewis et al., 2020).

This study responds to a documented gap in the application of AI-powered conversational agents specifically for Hindu studies. While AI chatbots have been explored in Islamic religious education (Adeshola & Adepoju, 2024) and biblical studies (Chrostowski & Najda, 2024), the design of dedicated AI consultation systems for Hindu academic institutions remains largely unexplored in scholarly literature. The Gitaverse chatbot study published in Chauhan and Dutt (2026) represents an early initiative applying LLM technology to Bhagavad Gita instruction, yet it does not address the broader scope of Hindu studies education in formal higher education settings, nor does it account for the multilingual and multicultural dimensions of Hindu education in Southeast Asia.

This paper presents DharmaBot, an intelligent conversational agent designed to address the academic consultation needs of Hindu studies students in Indonesian higher education. The system integrates LLM technology with a structured Hindu knowledge base covering classical scriptures, philosophical schools, ethical frameworks, and the global dimensions of Hindu practice. The study employs a Design Science Research framework to systematically design, develop, and demonstrate the system prototype. The primary objectives of this study are: (1) to identify the informational and consultative needs of Hindu studies students in higher education; (2) to design an AI-powered conversational system grounded in verified Hindu knowledge sources; and (3) to demonstrate a working system prototype applicable to the context of Hindu higher education institutions in Indonesia.

Literature Review

1. Large Language Models in Higher Education

Large Language Models represent a transformative development in artificial intelligence, characterized by their capacity to generate contextually coherent, semantically rich natural language responses at scale (Yigci et al., 2025). These models, trained on extensive corpora of human-generated text, demonstrate competencies across reasoning, summarization, translation, and dialogue that have prompted widespread adoption in educational settings globally. Denny et al. (2024) argue that higher education has entered a defining era shaped by generative AI, where LLMs fundamentally reshape instructional design, student assessment, and academic support mechanisms in ways that institutions are only beginning to accommodate.

The application of LLM-based chatbots in higher education has been subject to growing empirical scrutiny. Sørby and Simensen (2024) conducted a scoping review of 24 empirical studies on student perceptions and usage of AI chatbots in higher education between 2022 and 2023, finding that students regard these tools as highly useful for immediate knowledge retrieval, writing assistance, and personalized academic guidance. A notable finding from that review is that student engagement with AI chatbots is particularly pronounced in Asia, suggesting strong contextual relevance for the Indonesian higher education setting. Ansari et al. (2024) extend this landscape analysis, reviewing the emerging field of generative AI chatbots in higher education and identifying both the pedagogical opportunities and the institutional governance challenges associated with their deployment at scale.

Ocen et al. (2025) provide a comprehensive systematic review of AI innovations, opportunities, and challenges in higher education institutions, identifying chatbot systems and intelligent tutoring tools as among the most impactful AI applications in formal education. Their review of 54 documents confirms that AI tools consistently enhance research support, personalized learning delivery, and administrative efficiency. Ethical concerns around academic integrity and data fabrication were identified as the primary institutional risks requiring governance frameworks, a consideration relevant to the design of any educational AI system including DharmaBot.

2. Conversational AI and Chatbot Systems

Conversational AI refers to systems capable of engaging in goal-directed natural language dialogue with human users across extended, contextually coherent exchanges. The field has undergone a generational transformation following the introduction of LLM-powered systems, with rule-based chatbots largely supplanted by generative systems capable of flexible, contextually adaptive responses across an open range of topics (Raaijmakers et al., 2023). Kuhail et al. (2023) conducted a systematic review of educational chatbot interactions spanning 1999 to 2022, documenting that LLM-powered chatbots substantially outperform earlier rule-based systems across the dimensions of user satisfaction, response accuracy, and topic coverage flexibility.

The deployment of conversational agents in educational contexts has produced consistently positive learning outcomes when systems are designed with clear pedagogical objectives and domain-specific knowledge bases. Güler et al. (2026) demonstrate that student acceptance of AI chatbots in educational settings is positively associated with perceived usefulness and interaction ease, findings that underscore the importance of both interface design and knowledge accuracy in determining the educational effectiveness of deployed systems.

The application of conversational AI to religious and sacred-text education has received emerging scholarly attention. Chrostowski and Najda (2024) examine ChatGPT as a tool for biblical instruction in German confessional religious education, identifying the consultative potential of LLMs for scriptural interpretation while cautioning against the ethical risks of unverified doctrinal outputs. Munibi (2025) traces the broader impact of AI on religious practice and education in the digital era, documenting how AI-mediated platforms are transforming the dissemination of sacred knowledge across faith communities. Thatcher and Karimi (2024) examine religion-based AI chatbots through a critical feminist framework, noting that these systems function as conversational agents for sacred text guidance while presenting risks of doctrinal distortion that must be proactively mitigated through knowledge curation and institutional oversight.

3. Retrieval-Augmented Generation

Retrieval-Augmented Generation is a hybrid AI architecture that augments the generative capabilities of LLMs with real-time retrieval from external knowledge bases (Lewis et al., 2020). In a RAG system, an input query is first matched against a curated document corpus to retrieve semantically relevant passages, which are then provided as contextual input to the LLM generator. This architecture enables the model to produce responses grounded in specific, verifiable sources rather than relying solely on parametric knowledge encoded during model pre-training. Lewis et al. (2020) demonstrated that RAG architectures substantially outperform standalone generative models on knowledge-intensive natural language processing tasks, establishing the framework as a foundational approach for domain-specific AI applications requiring factual accuracy.

The application of RAG to educational chatbots addresses a fundamental limitation of general-purpose LLMs: the tendency to produce plausible-sounding but factually unreliable responses in specialized knowledge domains, a phenomenon known as hallucination. For Hindu studies education, where the accuracy of scriptural interpretation carries doctrinal significance, a RAG-grounded architecture provides an essential knowledge verification layer. The retrieval step ensures that system responses to queries about specific philosophical concepts, scriptural passages, or ritual interpretations draw from authentic, curated sources rather than from generalized model outputs that may inadvertently conflate traditions or misattribute teachings. This design principle aligns with the AI application frameworks identified by Terkowsky et al. (2022) in their DSR-based study of digital innovation engineering for conversational AI systems.

4. AI Applications in Religious Studies Education

The intersection of artificial intelligence and religious studies education has emerged as a growing research area with distinct ethical and pedagogical dimensions. Adeshola and Adepoju (2024) analyze the opportunities and challenges of ChatGPT in education broadly, with findings directly applicable to religious contexts where AI response accuracy carries both institutional and doctrinal weight. Their study identifies student trust, perceived accuracy, and interface usability as

the primary determinants of AI chatbot acceptance, all three of which are addressed in the DharmaBot design.

Chauhan et al. (2024) conducted a comparative study of moral learning through the Bhagavad Gita using social robot interaction, demonstrating that technology-mediated engagement with Hindu sacred texts produces measurable improvements in knowledge retention and moral reasoning capacity. The Gitaverse chatbot study published in Chauhan and Dutt (2026) further establishes the viability of LLM-based conversational systems for Bhagavad Gita instruction in formal educational settings, finding that students engaging with AI-powered philosophical dialogue demonstrated higher cognitive engagement than those relying on traditional text reading alone. These findings provide empirical support for the pedagogical premises underlying DharmaBot's design.

Thatcher and Karimi (2024) interreligious AI dialogue research (Corpuz, 2026) collectively identify the emerging institutional norm of AI as a mediator of sacred knowledge, while cautioning that systems lacking verified knowledge grounding risk generating outputs inconsistent with established doctrinal interpretations. Munibi (2025) affirms this analysis by documenting how language and AI technologies are reshaping religious practices and knowledge transmission in ways that demand intentional design rather than ad-hoc deployment.

5. Design Science Research in Information Systems

Design Science Research is a research paradigm that produces and evaluates IT artifacts designed to solve specific organizational and human problems through rigorous, iterative design and demonstration processes (Hevner et al., 2004). The DSR methodology formulated by Peffers et al. (2007) provides a structured six-stage process model: problem identification and motivation, definition of solution objectives, design and development, demonstration, evaluation, and communication. This framework has been extensively adopted in information systems research for the systematic design of AI-powered applications, decision support systems, and educational technology platforms.

Terkowsky et al. (2022) demonstrate the application of DSR methodology to the engineering of digital innovations in information systems, showing that the

framework produces artifacts with both scientific rigor and practical applicability across diverse technology domains. The present study adopts the DSR process model through Stage 4, presenting a verified system design and working prototype without formal user evaluation, consistent with the research scope of a conference proceedings contribution that precedes full-scale implementation. This methodological positioning is well-established in DSR literature, as Peffers et al. (2007) explicitly acknowledge that not all DSR contributions require completion of all six stages, with communication artifacts such as the present paper representing legitimate DSR outputs that advance the design knowledge base.

Method

This study employs the Design Science Research (DSR) methodology as formulated by Peffers et al. (2007), applied through the first four stages within the scope of this contribution. The problem addressed is the limited availability of contextual and doctrinally accurate Hindu Dharma consultation for students in Indonesian Hindu higher education institutions, motivated by the breadth of the Hindu knowledge corpus, the geographic constraints on instructor availability in IAHN campuses, and the digital equity gap documented by Aditya et al. (2022). The design objectives were defined to produce a conversational AI system capable of delivering accurate, reference-grounded responses across eight thematic domains of Hindu studies in bilingual Indonesian and English interaction, with transparent source attribution for every generated response. The system design integrates three primary components: a Large Language Model engine using Claude Sonnet as the generative core; a Retrieval-Augmented Generation pipeline in which verified Hindu knowledge documents are segmented, converted to vector embeddings, indexed in a local ChromaDB vector database, and retrieved through cosine similarity matching against incoming user queries; and a React.js web interface providing topic-filtered dialogue, session history, and inline source citation cards displayed within each bot response. The prototype is demonstrated through three consultation scenarios representing philosophical, theological, and comparative inquiry categories in both Indonesian and English, confirming system feasibility within the consultative

context of Indonesian Hindu higher education prior to formal user evaluation in future research.




Figure 1. DharmaBot system architecture — RAG pipeline and component integration




Figure 2. DharmaBot web interface wireframe — bilingual chat with topic filter and source citation

II. Discussion

1. System Architecture and RAG Pipeline

The DharmaBot architecture follows a retrieval-then-generate pipeline that ensures response accuracy through knowledge grounding prior to LLM generation. The processing sequence begins when a user submits a query through the web interface. The query undergoes vector encoding through an embedding model, after which the system performs a cosine similarity search against the indexed Hindu knowledge corpus stored in ChromaDB. The top three to five most semantically relevant document segments are retrieved and injected into the LLM prompt as structured contextual input. The LLM then synthesizes the retrieved content into a coherent, contextually appropriate response with explicit attribution to the source documents provided in a citation card displayed below the response text.

This architecture directly addresses the doctrinal accuracy risks documented in studies of unverified Hindu AI applications (Thatcher & Karimi, 2024). The retrieval step ensures that responses to queries about specific scriptural passages,

philosophical positions, or ritual interpretations draw from authentic, curated knowledge segments rather than from generalized model inference that may conflate traditions or misattribute teachings from different schools. Every claim in a DharmaBot response is traceable to a specific source document in the knowledge base, providing a layer of scholarly accountability absent from general-purpose AI systems currently used by students for Hindu studies reference.

2. Knowledge Base Architecture

The DharmaBot knowledge base is structured across eight thematic domains that collectively represent the full breadth of Hindu studies as taught in Indonesian Hindu higher education. The Classical Philosophy and Core Concepts domain covers Brahman, Atman, Dharma, Karma, Moksha, Samsara, Maya, Ahimsa, and the four Yoga paths. The Six Philosophical Schools domain encompasses the Sad Darshana: Vedanta, Samkhya, Yoga, Nyaya, Vaisheshika, and Mimamsa. The Sacred Scriptures domain covers the four Vedas, 108 principal Upanishads, Bhagavad Gita, Ramayana, Mahabharata, the 18 Mahapurana, and Manusmriti. The Theology and Divine Traditions domain addresses Trimurti, Dasavatara, Sakti, and Tantra traditions.

The Ethics and Hindu Values domain integrates both universal Hindu ethical frameworks, including Yama and Niyama from Patanjali's Yoga Sutra, Catur Purusartha, Catur Ashrama, and Catur Warna, and specifically Indonesian Hindu ethical frameworks including Tri Hita Karana, Pancasrada, Tri Kaya Parisudha, Sapta Timira, and Dasa Yama Brata. The Global Hinduism domain covers Hindu communities in India, Nepal, Bali, Kalimantan, the diaspora in the United States and Europe, ISKCON as a global movement, and the Hindu communities of Mauritius and Trinidad and Tobago. The Ritual and Ceremonial Practice domain encompasses Puja, Yajna, the 16 Samskara, Vivaha, Antyesti, and major festivals including Nyepi, Galungan, Diwali, Holi, and Navaratri. The Hindu-Science Intersections domain covers Yoga and health as recognized by the World Health Organization, meditation and neuroscience research from Harvard and MIT, Ayurveda, Vedic astronomy, and the Hindu origins of numeral zero and the decimal system.

This taxonomic structure ensures that DharmaBot can respond to the full range of student queries encountered in Hindu higher education curricula, from

introductory philosophical concepts to advanced comparative studies of Hindu traditions across global contexts. The deliberate inclusion of Indonesian and Balinese Hindu ethical frameworks alongside classical and globally-comparative content reflects the specific curricular needs of IAHN-TP Palangka Raya students while positioning the system within the broader global Hindu studies discourse appropriate for an international proceedings contribution.

3. Interface Design and User Experience

The DharmaBot web interface is designed around a structured two-panel layout that supports intuitive academic consultation. The fixed navigation bar displays the DharmaBot identity with a lotus flower icon and provides access to the application pages and a dark-light mode toggle. The left sidebar panel provides a thematic domain filter allowing students to select one of eight knowledge areas before submitting a query, which guides the RAG retrieval toward the most contextually relevant document segments. The right main panel presents the conversational interface, with bot responses displayed in chat bubbles that include inline source citation cards identifying the specific text, chapter, and verse from which each response is drawn.

The welcome screen displays six suggested query prompts drawn from common student consultation patterns in Hindu higher education, covering Dharma definition, Karma and Moksha, Bhagavad Gita inquiry, Yoga and meditation, Hindu global practice, and ritual significance. This guided-entry design follows the principles identified by Sørby and Simensen (2024), who confirm that visible knowledge scope indicators and suggested query prompts significantly improve student confidence in AI-powered academic tools and reduce the cognitive effort required to initiate productive consultations.

A typing indicator displays during LLM response generation, and completed responses include a copy function and a timestamp. The session history panel in the sidebar lists previous queries for reference retrieval within the session. The system supports bilingual input, generating Indonesian-language responses to Indonesian queries and English responses to English queries, consistent with the multilingual academic environment documented in Indonesian Hindu higher education contexts.

4. Prototype Demonstration Scenarios

Three demonstration scenarios illustrate DharmaBot's consultative capability across the primary knowledge domains.

Scenario 1 presents an English-language philosophical inquiry: "What is the concept of Karma in Hinduism and how does it relate to daily life?" The RAG pipeline retrieves segments from the Bhagavad Gita (Chapter 4, Verse 17) and the Yoga Sutra of Patanjali (Pada 2, Sutra 12-14), generating a response that defines Karma as the law of moral causality governing the consequences of intentional action, distinguishes between Sanchita, Prarabdha, and Agami Karma, and explains the practical implications for ethical daily conduct. Both source references appear in citation cards below the response text. This scenario demonstrates the system's capacity to ground philosophical explanation in primary scriptural authority.

Scenario 2 presents an Indonesian-language query: "Jelaskan perbedaan antara Moksha dan Samsara dalam perspektif Advaita Vedanta." The retrieval pipeline draws from the Mandukya Upanishad (Verse 7) and Shankaracharya's Vivekachudamani, generating an Indonesian-language explanation that distinguishes the state of Moksha as the realization of non-dual Brahman consciousness from the cyclical bondage of Samsara driven by Avidya, or ignorance of the true Self. This scenario confirms bilingual response generation and the system's command of classical Vedantic philosophy in the language of Indonesian students.

Scenario 3 presents a comparative global inquiry: "How is Hinduism practiced differently in Bali, India, and the Hindu diaspora?" The knowledge base retrieval draws from the Global Hinduism domain, generating a comparative response that identifies Tri Hita Karana as the organizing principle of Balinese Hindu daily and ceremonial life, contrasts it with the Bhakti and Jnana devotional traditions dominant in mainland India, and describes the organizational structures of ISKCON and temple-centered Hindu diaspora communities in the United States and United Kingdom. This scenario demonstrates the system's capacity for comparative cross-regional Hindu studies consultation, directly relevant to the international scope of the host proceedings.

5. Implications for Hindu Higher Education

DharmaBot addresses a structural gap in the academic support infrastructure of Hindu higher education institutions. The system's integration of conversational AI into Hindu studies education yields four primary implications for institutional practice and future research.

The first implication concerns consultative access. The system extends scholarly consultation availability beyond instructor office hours and geographic constraints, providing 24-hour access to doctrinal and philosophical guidance grounded in verified sources. Students in remote IAHN campuses gain access to the same quality of knowledge consultation available to students in better-resourced urban institutions, partially bridging the digital equity gap identified by Aditya et al. (2022).

The second implication concerns doctrinal accuracy. The RAG architecture directly mitigates the doctrinal accuracy risks associated with general-purpose AI systems applied to Hindu content, a concern documented by Thatcher and Karimi (2024) in their analysis of unverified religion-based chatbots. By anchoring every response in specific, curated knowledge segments from verified Hindu scholarly sources, DharmaBot provides an academically defensible alternative to the unverified AI platforms students are already using for Hindu studies reference.

The third implication concerns pedagogical design. The system's taxonomic knowledge structure, guided-entry interface, and citation transparency align with the pedagogical principles for educational chatbot design identified across the reviewed literature (Kuhail et al., 2023; Sjøby & Simensen, 2024; Güler et al., 2026), suggesting that DharmaBot's design approach is transferable to other religious education contexts where doctrinal accuracy and source attribution are institutional priorities.

The fourth implication concerns the digital transformation of Hindu higher education institutions. Adeshola and Adepoju (2024) and Ocen et al. (2025) both identify the need for AI governance frameworks in higher education institutions that adopt AI academic support tools. DharmaBot's design incorporates governance principles through its knowledge curation protocol and source attribution

mechanism, offering a model for responsible AI integration in Indonesian Hindu higher education that respects both technological opportunity and doctrinal responsibility.

III. Closing

This study set out to address the limited availability of contextual, accurate, and accessible Hindu Dharma consultation for students in Hindu higher education institutions in Indonesia, a problem driven by the breadth of the Hindu knowledge corpus, the geographic and temporal constraints on instructor availability in IAHN campuses, and the growing but unguided use of general-purpose AI tools for Hindu studies reference among students. DharmaBot was designed and demonstrated as a systematic response to this problem through a Design Science Research process that produced an intelligent conversational agent integrating a Large Language Model engine, a Retrieval-Augmented Generation pipeline grounded in verified Hindu scholarly documents, and a bilingual web interface covering eight thematic domains of Hindu studies. The three prototype demonstration scenarios confirmed that the system delivers accurate, source-attributed responses to philosophical, theological, and comparative Hindu studies queries in both Indonesian and English, directly addressing the doctrinal accuracy risks and consultative access gaps identified in the problem statement. Future research should prioritize the full implementation and indexing of Indonesian-language Hindu scriptural texts to complete the knowledge base, formal usability and accuracy evaluation through structured user testing with IAHN-TP Palangka Raya students, and the exploration of adaptive personalization features that adjust response complexity to the academic level of the querying student, collectively advancing DharmaBot from a verified system design toward a fully deployed academic support tool for Hindu higher education institutions in Indonesia and beyond.

Bibliography

- Adeshola, I., & Adepoju, A. P. (2024). The opportunities and challenges of ChatGPT in education. *Interactive Learning Environments*, 32(10), 6159-6172. <https://doi.org/10.1080/10494820.2023.2253858>
- Aditya, B. R., Ferdiana, R., & Kusumawardani, S. S. (2022). Identifying and prioritizing barriers to digital transformation in higher education: A case study

- in Indonesia. *International Journal of Innovation Science*, 14(3-4), 445-460.
<https://doi.org/10.1108/IJIS-11-2020-0262>
- Ansari, J. A. N., Khan, N. A., & Agarwal, S. (2024). Generative AI chatbots in higher education: A review of an emerging research area. *Higher Education*.
<https://doi.org/10.1007/s10734-024-01288-w>
- Chauhan, S., Chaudhary, G., Bhavsar, A., Mombaur, K., & Dutt, V. (2024). Exploring moral learning through Bhagavad Gita: A comparative study of social robot interaction and traditional reading. In *International Conference on Social Robotics*. https://doi.org/10.1007/978-981-99-6809-0_20
- Chrostowski, W., & Najda, A. (2024). ChatGPT as a modern tool for Bible teaching in confessional religious education. *Journal of Religious Education*.
<https://doi.org/10.1007/s40839-024-00244-1>
- Denny, P., Prather, J., Becker, B. A., Finnie-Ansley, J., Hellas, A., Leinonen, J., Luxton-Reilly, A., Reeves, B. N., Santos, E. A., & Sarsa, S. (2024). Computing education in the era of generative AI. *Communications of the ACM*.
<https://doi.org/10.1145/3624720>
- Chauhan, S., & Dutt, V. (2026). Bridging ancient wisdom and cognitive engagement: A comparative study of chatbot-based moral instruction. *Frontiers in Robotics and AI*, 13, 1739259.
<https://doi.org/10.3389/frobt.2026.1739259>
- Güler, M. A., Dağ, S., Kayal, S., Can, M., & Şimşekkaya, S. N. (2026). Physiotherapy students' acceptance of AI-based chatbots including ChatGPT in education: A multi-institutional study from Turkey. *BMC Medical Education*.
<https://doi.org/10.1186/s12909-025-08535-3>
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105.
<https://doi.org/10.2307/25148625>
- Kuhail, M. A., Alturki, N., Alramlawi, S., & Alhejori, K. (2023). Interacting with educational chatbots: A systematic review. *Education and Information Technologies*, 28(1), 973-1018. <https://doi.org/10.1007/s10639-022-11177-3>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474. <https://arxiv.org/abs/2005.11401>
- Munibi, I. (2025). Faith in the digital era: How language and artificial intelligence technology reshape religious practices. *Jurnal SMART (Studi Masyarakat, Religi, dan Tradisi)*. <https://doi.org/10.18784/smart.v11i1.2906>
- Ocen, S., Elasu, J., Aarakit, S. M., & Olupot, C. (2025). Artificial intelligence in higher education institutions: Review of innovations, opportunities and challenges. *Frontiers in Education*, 10, 1530247.
<https://doi.org/10.3389/feduc.2025.1530247>
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77.
<https://doi.org/10.2753/MIS0742-1222240302>

- Raaijmakers, S., Cremers, A., Krahmer, E., & Westera, M. (2023). Editorial: Conversational AI. *Frontiers in Artificial Intelligence*. <https://doi.org/10.3389/frai.2023.1203910>
- Corpuz, J. C. G. (2026). Artificial intelligence and interreligious dialogue: Emerging implications for faith-based organizations. *Religions*, 17(3), 354. <https://doi.org/10.3390/rel17030354>
- Søby, M. A., & Simensen, A. M. (2024). Perceptions and use of AI chatbots among students in higher education: A scoping review of empirical studies. *Education Sciences*, 14(8), 922. <https://doi.org/10.3390/educsci14080922>
- Terkowsky, C., Haertel, T., Ortelt, T. R., & May, D. (2022). Development of a method for the engineering of digital innovation using design science research. *Information*, 13(12), 573. <https://doi.org/10.3390/info13120573>
- Thatcher, P. B., & Karimi, F. (2024). Feminist re-engineering of religion-based AI chatbots. *Philosophies*, 9(1), 20. <https://doi.org/10.3390/philosophies9010020>
- Yigci, D., Eryilmaz, M., Yetisen, A. K., Tasoglu, S., & Ozcan, A. (2025). Large language model-based chatbots in higher education. *Advanced Intelligent Systems*, 7, 2400429. <https://doi.org/10.1002/aisy.202400429>