




Journal of Artificial Intelligence

ISSN 1994-5450

science
alert

ANSI*net*
an open access publisher
<http://ansinet.com>

Review Article

AI and the Concept of Afterlife: Philosophical, Technological, Ethical and Cultural Perspectives

Thomas Keller and Mirella Moser

Zurich University of Applied Sciences Institute of Information Systems, Winterthur, Switzerland

Abstract

Recent advances in generative artificial intelligence, conversational agents and multimodal systems have accelerated interest in AI afterlife technologies. These systems aim to simulate, preserve, or extend aspects of a person's identity after death through chatbots, avatars, voice clones and interactive memorial platforms. This article reviews the emerging literature on AI afterlife and examines its philosophical foundations, technological forms, ethical challenges and cultural implications. The review identifies four main system types: Legacy curation tools, memorial recommender systems, conversational griefbots and high-fidelity digital human simulations. Across these approaches, the literature reveals recurring tensions between authenticity and simulation, remembrance and dependency, innovation and exploitation and personal agency and platform control. Central concerns include consent, privacy, identity rights, data governance, commercialization of grief and the psychological effects of ongoing interaction with representations of the deceased. In response, the article offers a holistic perspective on AI afterlife design by linking three dimensions: Technical choices, stakeholder roles and lifecycle governance. The framework links design decisions such as data selection, controllability, disclosure and safety measures with the interests of decedents, relatives, platform providers and regulators across pre-mortem planning, bereavement and long-term stewardship. The article concludes by outlining priorities for future research on evaluation standards, harm monitoring, culturally sensitive design and regulatory approaches that balance innovation with protection of both the dead and the living.

Key words: AI afterlife, digital immortality, griefbots, digital legacy, posthumous data, ethical AI, conversational agents; identity simulation

Citation: Keller, T., M. Moser, 2026. AI and the concept of afterlife: Philosophical, technological, ethical and cultural perspectives. J. Artif. Intel., 19: 38-48.

Corresponding Author: Thomas Keller, Zurich University of Applied Sciences Institute of Information Systems, Winterthur, Switzerland

Copyright: © 2026 Thomas Keller and Mirella Moser. This is an open access article distributed under the terms of the creative commons attribution License, which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

Competing Interest: The authors have declared that no competing interest exists.

Data Availability: All relevant data are within the paper and its supporting information files.

INTRODUCTION

The concept of an AI afterlife refers to socio-technical practices in which artificial intelligence is used to simulate, extend, or curate the posthumous presence of a person through interactive systems such as chatbots, avatars, or other generative agents¹⁻³. The topic has accelerated with the maturation of large language models, speech synthesis and multimodal generation that allow systems to appear conversationally competent and emotionally responsive, giving the impression of continuity with a deceased individual⁴⁻⁶. Scholarship now treats AI afterlife not as speculative fiction but as an emerging domain that blends computer science, philosophy of mind, digital ethics, media studies and cultural anthropology⁷⁻⁹.

The literature is shaped in part by the rapid rise of technical capability. Contemporary systems can ingest personal traces such as messages, recordings, photographs and other digital artifacts; retrieve relevant memories from these materials; and generate fluent responses that resemble a particular person's style of speaking¹⁰⁻¹². Progress in large language models, speech synthesis, multimodal generation and retrieval-based personalization has made posthumous simulation more technically feasible and more convincing than earlier forms of digital memorialization.

At the same time, the social meaning of these systems remains deeply contested. Researchers debate whether AI-mediated remembrance can support mourning, memory and continuing bonds, or whether it risks manipulation, commodification and new forms of harm related to privacy, consent and identity misrepresentation^{6,13,14}. These tensions make AI afterlife a particularly revealing case for broader questions in contemporary AI, including what counts as authenticity, which forms of personhood are socially recognized and how governance should evolve when data about people continues to act after death¹⁴⁻¹⁶.

Accordingly, this review examines AI afterlife as an interdisciplinary socio-technical phenomenon at the intersection of philosophy, AI technology, ethics, law and culture. The review is guided by the following research question: What philosophical, technological, ethical, legal and cultural perspectives on AI afterlife are represented in the current literature and what future research directions can be identified? The scope of the review is limited to scholarly literature addressing AI-enabled posthumous representation, digital legacy, griefbots, digital resurrection, memorial agents and related forms of interactive or generative representation of the deceased.

MATERIALS AND METHODS

This article was conducted as a literature review¹⁷. The literature search was conducted during March 2026 and included Scopus and Web of Science. The review was designed to identify, screen and synthesize academic work addressing philosophical, technological, ethical, legal, cultural and socio-technical perspectives on AI afterlife. Given the conceptual breadth and interdisciplinary nature of the topic, a review design was chosen to ensure transparency, reproducibility and methodological rigor.


Search strategy: The literature search was primarily conducted in Scopus and Web of Science, as these databases provide broad coverage across computer science, social sciences, humanities, law, ethics and interdisciplinary technology studies. The search strategy was developed from the research question and refined iteratively to capture both established and emerging terminology used in the field.

The search string combined terms related to artificial intelligence with terms related to afterlife, posthumous presence, digital legacy and memorial technologies. Boolean operators and truncation were used to maximize recall while maintaining thematic relevance. A representative search structure was as follows.

("artificial intelligence" OR AI OR "generative AI" OR chatbot OR avatar OR "digital human*" OR "virtual agent*" OR "conversational agent*") AND ("afterlife" OR posthumous OR post-mortem OR deceased OR death OR grieving OR mourning OR bereavement OR memorial* OR "digital legacy" OR "digital remains" OR "digital immortality" OR griefbot* OR deathbot*)**.

Search was limited to fields such as title, abstract and keywords where appropriate. Additional backward and forward citation checks were used selectively to identify relevant studies that may not have been retrieved in the database search due to terminological variation.

The database search identified 1679 records in total, including 762 records from Scopus and 917 records from Web of Science. After automatically removing duplicate records and subsequently screening title and abstract with help of AI tools (chatgpt 5.4 in thinking mode) 1602 records were excluded because they were outside the scope of AI afterlife, digital legacy, posthumous representation, griefbots, or related memorial technologies. The full texts of 77 potentially eligible publications were then assessed. Of these, 24 studies were excluded because they did not substantively address



Primary databases: Scopus and Web of Science. Additional backward and forward citation checks were used selectively; no separate counts were reported.

Fig. 1: Screening diagram

AI-enabled posthumous representation, were not scholarly publications, lacked accessible full text, or focused on death, bereavement, or AI without a direct connection to afterlife-related digital representation. A final corpus of 53 studies¹⁸ was included in the qualitative synthesis (Fig. 1).

Inclusion and exclusion criteria: Studies were included if they were peer-reviewed journal articles, conference papers, book chapters, or review articles written in English published after 2006 and made a substantive contribution to understanding AI-related systems, concepts, or debates in connection with death, mourning, posthumous identity, digital legacy, memorialization, or afterlife-related representation. The review considered conceptual, empirical, design-oriented, philosophical, ethical, legal, cultural and technological perspectives to reflect the interdisciplinary nature of the field. Studies were excluded if they were non-scholarly publications, unrelated to AI afterlife or adjacent posthumous digital phenomena, focused on death or bereavement without a meaningful connection to AI or data-driven representation, or addressed AI in healthcare or palliative settings without

relevance to afterlife, memorialization, or posthumous identity. Duplicate records and publications without accessible full text were also excluded.

Screening and selection process: The screening process followed a two-phase approach. In the first phase, records retrieved from Scopus and Web of Science were exported and merged into a single review corpus. Duplicate records were identified and removed. The remaining records were then subjected to an initial automated screening using AI-based tools (ChatGPT 5.4 in thinking mode), which supported the identification of potentially relevant studies based on titles, abstracts and keyword patterns. This phase was used as an efficiency-enhancing step rather than a substitute for scholarly judgment.

In the second phase, all retained records underwent manual screening. Titles and abstracts were reviewed to assess their relevance against the inclusion and exclusion criteria. Where relevance could not be determined from the abstract alone, the full text was examined. Full-text screening was then conducted for all provisionally eligible studies.

Because the topic spans multiple disciplines and uses diverse vocabularies, particular care was taken during manual screening to include conceptually relevant studies even when they did not explicitly use terms such as *AI afterlife*, *deathbot*, or *griefbot*. Borderline cases were resolved through repeated review of the study's central argument, object of analysis and contribution to the research question. This two-step process balanced breadth of retrieval with interpretive precision.

Methodological limitations: Several limitations should be acknowledged. First, the review was restricted to Scopus and Web of Science and to English-language publications, which may have excluded relevant studies indexed elsewhere or published in other languages. Second, the terminology of AI afterlife is still unstable, which means that some relevant studies may use adjacent concepts not fully captured by the search string. Third, the use of AI-assisted screening improved efficiency but required careful manual verification to avoid false exclusions. Finally, because the field is conceptually diverse and methodologically heterogeneous, the review emphasizes thematic interpretation and conceptual synthesis rather than statistical generalization.

DEFINING AI AFTERLIFE

Across the Scopus corpus, AI afterlife functions as an umbrella term that includes digital resurrection, posthumous digital twins, griefbots and interactive memorial technologies. Although implementations differ, they share a common aspiration: to reconstitute patterns of communication, memory and personality from data and then to present these patterns through an interface that invites ongoing interaction^{2,6,8}. Some projects focus on cultural heritage and learning by reconstructing historical figures for museums or education; while not always framed as afterlife, such systems raise adjacent questions about authenticity, authorship and the ethics of speaking as another person¹⁹.

A practical definition for scholarship is therefore functional and relational rather than metaphysical: AI afterlife is any AI-driven system that enables sustained social interaction with a posthumous representation, where the representation is perceived by users as meaningfully connected to the deceased (or to a historically absent person)^{3,5,15} (Fig. 2). This definition foregrounds the social effects of the system, how it shapes grief, memory and relationships while leaving open deeper metaphysical questions about whether consciousness or identity can be preserved in computational form²⁰⁻²².

PHILOSOPHICAL PERSPECTIVES ON AI AFTERLIFE

Philosophical work provides the conceptual vocabulary for evaluating AI afterlife beyond surface realism. The central problem is not whether an avatar can produce convincing language, but what kind of entity the avatar is: A tool for remembrance, an artwork, a social actor, or a counterfeit person^{3,9,15}. Two longstanding debates structure contemporary discussions. One concerns whether computational systems could ever be conscious; the other concerns whether personal identity can persist when the person is reconstructed from data traces rather than embodied life¹⁹⁻²¹.

A recurring analytic distinction is between weak AI, which simulates intelligent behavior without subjective experience and strong AI, which would possess genuine understanding and consciousness. AI afterlife systems are almost always treated as weak AI in practice, yet they are frequently interpreted as if they had strong-AI properties because they speak in the first person and participate in emotionally significant dialogue¹⁹⁻²¹. This interpretive gap between system ontology and user experience drives concerns about authenticity, deception and the moral status of digital persons^{6,8,13}.

Consciousness and identity in AI afterlife: Debates about consciousness in AI afterlife often return to classic arguments in the philosophy of mind. Searle's Chinese Room remains influential because it challenges the inference from syntactic manipulation to semantic understanding; even if a system produces the right outputs, that does not entail it knows what it is saying²⁰. Contemporary commentators revisit this argument in light of generative AI, asking whether scale, embodiment, or new forms of world-modeling alter the conclusion, or whether the Chinese Room diagnosis becomes more urgent as systems become more fluent²¹.

Identity questions are equally persistent. If an AI system uses personal archives to emulate a person's linguistic habits and preferences, does it preserve identity, or merely a curated likeness? Several authors argue that personal identity is not a static bundle of traits but a narrative and relational phenomenon that depends on embodied history, social recognition and continuing agency^{3,15,23}. From this view, AI afterlife can reproduce patterns but not the lived continuity that grounds accountability, change and the first-person perspective. The result is a representation that may be meaningful for mourners yet ontologically discontinuous with the deceased^{6,8,16}.

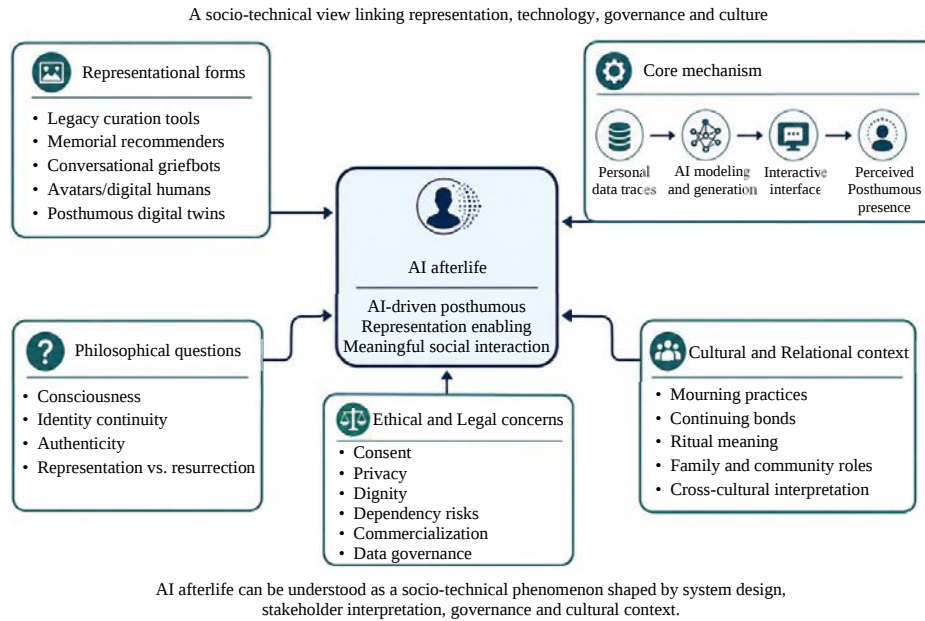


Fig. 2: Conceptual dimensions of AI afterlife

A productive philosophical move is to treat AI afterlife as a spectrum of representational practices rather than an all-or-nothing attempt at resurrection. At one end are archival memorials that support remembrance without claiming agency; at the other are interactive agents that speak, decide and even create new content attributed to the deceased^{2,4,5}. The more agentic the system becomes, the more acute the question of whether it constitutes a new persona built from old data and whether users are entitled to know that distinction explicitly^{6,13,14}.

Ethical theories and digital resurrection: Ethical discussions translate the metaphysical uncertainty of AI afterlife into normative questions about harm, dignity and responsibility. A key concern is emotional impact: Systems that produce comforting dialogue may also prolong grief, encourage dependence, or shape memory in ways that the deceased could not contest^{6,8,13}. Even when users understand that the system is artificial, repeated interaction can generate strong affective bonds, creating a setting in which manipulation, e.g., by platform incentives, advertising, or ideological framing, becomes possible^{13,16,24}.

Another ethical thread concerns commodification. If posthumous representations become products, then death becomes a market category: Companies may sell continued presence, premium intimacy features, or enhanced versions of a person's voice and image^{6,14,15}. Such commodification can entrench inequalities, because affluent families may secure better afterlife services, while others face extractive data

practices or low-quality representations. The literature therefore emphasizes transparency, limits on commercial use and the need for governance institutions that recognize digital remains as ethically weighty rather than merely informational assets^{14,16,24}.

Ethical arguments also grapple with the moral status of the representation itself. If an AI persona becomes socially treated as an agent, does it deserve consideration, or is it solely an instrument for the living? Some work suggests that even if the system lacks consciousness, respectful design and usage norms matter because they shape how communities relate to the dead and to one another^{3,13,25}. This reframes the issue: The ethical locus is not rights of the bot but the integrity of mourning practices, the values encoded in memorialization and the protection of vulnerable users during bereavement^{6,8,9}.

Non-Western philosophical approaches: A notable contribution of recent literature is the expansion beyond predominantly Western individualist frameworks. Buddhist and Confucian traditions, among others, offer ethical lenses that emphasize interdependence, relational identity and ritual practice^{1,26}. In Buddhist ethical analysis of AI, moral evaluation focuses on reducing suffering, cultivating compassion and avoiding attachment that entrenches distress. These themes map directly onto grief technologies: An AI afterlife system could be evaluated by whether it supports healthy remembrance and alleviates suffering, rather than by whether it recreates an individual essence^{1,26,27}.

Confucian-inspired perspectives emphasize roles, family relations and ritual propriety. From this view, the moral quality of an AI afterlife system depends on how it supports communal mourning and respectful remembrance, not merely individual preference satisfaction^{28,29}. Confucian and Daoist analyses in adjacent medical AI contexts also highlight the importance of harmony, moral cultivation and the preservation of human responsibility within socio-technical systems; principles that can guide the design of memorial agents that augment rather than replace human caregiving and community practices^{28,30,31}.

These non-Western approaches also complicate the idea that autonomy and consent are the sole ethical foundations. In relational frameworks, family and community often have legitimate standing in decisions about posthumous representation and moral attention shifts toward negotiated responsibilities, shared memory and ritual continuity^{3,25,29}. This does not weaken individual rights; rather, it motivates governance models that explicitly incorporate pluralistic values and culturally situated practices, especially in cross-border deployments^{16,24,32}.

LIMITS AND POSSIBILITIES OF AI CONSCIOUSNESS

Contemporary philosophy revisits whether advances in generative AI challenge earlier skepticism about machine minds. Some analyses argue that current systems remain fundamentally tool-like because they lack grounded intentionality and biological teleology, even if they display flexible language use¹⁹⁻²¹. Others suggest that new architectures and embodied contexts may gradually blur the line between simulation and cognition, making the strong AI question less straightforward than earlier symbolic AI debates^{19,22,33}.

For AI afterlife, however, the immediate philosophical stakes may be less about proving machine consciousness and more about avoiding category mistakes in social practice. Even if an AI persona is not conscious, people can still treat it as a conversational partner and a repository of meaning. The ethical challenge is therefore to develop norms and interfaces that communicate the system's status, limitations and provenance without undermining its therapeutic or commemorative value^{6,12,15}. Philosophical clarity becomes a design resource: by distinguishing representation from resurrection, systems can support remembrance while reducing deceptive anthropomorphism^{3,23}.

TECHNOLOGICAL PERSPECTIVES

The technical literature on AI afterlife spans human-computer interaction prototypes, system architecture papers and broader reviews of generative AI infrastructure relevant to personalized agents. Practical implementations often combine (a) A data pipeline for collecting and curating personal traces, (b) A modeling layer that generates language, voice, or image outputs and (c) An interaction layer that frames the persona through chat, embodied avatars, or mixed reality^{1,2,12}. While early systems were relatively scripted, recent work leverages foundation models to produce open-ended conversation, enabling a sense of spontaneity that users may read as personality^{4,5}.

At the same time, technical authors increasingly emphasize reliability and controllability. In high-stakes emotional contexts such as bereavement, hallucinated facts, inconsistent persona behavior and security vulnerabilities can cause significant harm. As a result, personalization is no longer framed as merely a user-experience feature; it becomes a safety requirement that demands provenance tracking, guardrails and evaluation methods beyond generic benchmark performance¹⁰.

Current technologies and models: Current AI afterlife systems range from digital memorial platforms³⁴ to interactive avatars presented as digital humans³⁵. Prototype platforms describe pipelines that integrate text chat, speech synthesis and visual embodiment, producing a multimodal surrogate that can respond to user prompts in real time¹. Related work in cultural heritage demonstrates that AI-generated historical characters can be used for learning and engagement, illustrating how speaking agents can influence memory and meaning even when the represented person is not the user's relative¹⁹.

From an engineering standpoint, a core design choice is how closely the system should adhere to available evidence about the person. Some systems are conservative, restricting outputs to retrieved archival materials; others allow generative extrapolation to maintain conversational flow and emotional resonance^{10,12}. This trade-off, fidelity versus flexibility, maps directly to ethical questions about authenticity and misrepresentation^{35,36}. It also affects system evaluation: Fidelity can be assessed through source alignment, while flexibility requires assessing conversational quality, empathy and long-term user outcomes⁶.

Generative AI and adaptive digital personas: Large language models and diffusion-based generators enable adaptive personas that can speak in a consistent style, respond to new contexts and produce multimodal artifacts attributed to the

deceased. These capabilities are technologically impressive but conceptually fraught because they encourage users to treat outputs as expressions of an enduring self rather than as probabilistic completions shaped by training data and prompts^{4,5,23}. Scholars warn that the digital twin metaphor can be misleading: Personalization creates resemblance, not identity and the metaphor can obscure risks related to over-trust and misplaced expectations.

Research also explores personalization techniques such as prompt engineering, few-shot examples, parameter-efficient fine-tuning and memory modules. In afterlife contexts, these methods must contend with sparse or biased personal data, conflicting accounts from relatives and the difficulty of representing ethically sensitive traits (e.g., political views, intimate relationships) responsibly. Technical approaches therefore intersect with governance proposals that define which data can be used, who can authorize its use and how the system should represent uncertainty and contested memories^{14,24}.

Retrieval-augmented generation and knowledge grounding: Retrieval-augmented generation³⁷ (RAG) is frequently proposed to improve factual grounding and reduce hallucinations by conditioning generation on retrieved documents. In AI afterlife, RAG can anchor conversation in a person's verified writings, recordings, or curated biographies, helping the system stay aligned with traceable sources¹⁰. Systematic reviews emphasize that the effectiveness of RAG depends on the quality of the retriever, the structure of the memory store and the prompt that integrates retrieved evidence; poor retrieval can still yield confident but incorrect outputs.

The RAG also changes the ethical profile of afterlife systems. Because retrieved sources can be logged and audited, RAG supports transparency: Users and regulators can inspect whether an output was grounded in a specific artifact or produced by free generation. This makes it easier to implement provenance features, contest incorrect statements and respect do not disclose constraints around sensitive materials^{10,14,24}. However, grounding does not solve the deeper question of persona continuity: Even perfectly sourced quotations do not amount to the agency or inner life of the person, reinforcing the need for careful interface framing^{20,23}.

Explainability, trust and human-AI collaboration: Trust in AI afterlife systems is a multi-layered concept. Users may trust the system's factual accuracy (did it really happen?), its persona faithfulness (would they say that? and its emotional appropriateness (is this supportive?). Explainable AI approaches and human-AI collaboration research argue that

trust should be calibrated rather than maximized: Users need cues about uncertainty, sources and system limits to avoid harmful over-reliance^{6,10,11}. In bereavement contexts, calibrated trust is especially important because emotional vulnerability can increase susceptibility to persuasive or misleading outputs^{13,16}.

Design proposals therefore advocate layered explanations: Brief, non-disruptive indicators in the interface (e.g., source badges, memory citations), with deeper transparency available on demand (e.g., trace views, audit logs). Such approaches align with governance frameworks for human digital remains that treat identity representation as a safety-critical function rather than a purely experiential feature^{14,24}.

FUTURE DIRECTIONS AND CHALLENGES

Looking forward, technical challenges cluster around safety, provenance and efficiency. Persona consistency across long conversations, resistance to prompt injection and secure storage of sensitive personal archives remain active research priorities. Another challenge is evaluation: Standard NLP metrics do not capture the relational and psychological impacts of interacting with a posthumous agent. Researchers therefore call for mixed-method evaluation designs that combine technical audits with longitudinal user studies and cross-cultural validation. Finally, there is growing interest in ethically constrained generation, including rule-based constraints, constitutional prompting and policy-aware decoding. In afterlife systems, these constraints would encode norms about what should not be said in the name of the deceased, how to handle disputed memories and when to redirect users to human support. Such work illustrates a broader shift in AI research: From optimizing raw capability to engineering responsible behavior in context-rich social domains^{16,24}.

ETHICAL AND LEGAL CONSIDERATIONS

Ethical and legal analyses of AI afterlife focus on who has authority over posthumous data and representation, what counts as valid consent and how to protect both the deceased's interests and the living's well-being. Governance discussions increasingly treat digital remains as a distinct category of personal data, requiring explicit rules for storage, access and posthumous use. This framing is motivated by the intuition that data linked to personhood, e.g., voice, private messages, health records, intimate photographs, should not automatically become exploitable simply because the person has died^{14,24}.)

Because AI afterlife often involves transnational platforms, legal analysis also highlights cross-border complexity. Baseline protections may exist in data protection regimes, but afterlife-specific uses expose gaps: many frameworks were not designed for autonomous generation, emotional manipulation risks, or identity simulation. As a result, scholars propose specialized policy instruments and institutional roles such as data trustees and advanced data directives¹⁶.

Informed consent and data privacy frameworks: Consent is foundational but difficult in posthumous contexts. Some proposals introduce digital wills or advance directives that specify which datasets may be used, for which purposes and by whom. Others argue for dynamic consent models that allow ongoing negotiation among stakeholders, e.g., individuals (before death), families and platform operators, because the downstream behavior of generative models is hard to predict in advance^{14,24,38}.

Privacy is not only informational but psychological. When an AI system reconstructs patterns of thought, speech, or preference, it can reveal aspects of a person that were never explicitly documented or intended for public exposure. Governance work therefore, emphasizes mental privacy, restrictions on inferential profiling and controls on secondary use (e.g., marketing, political targeting) of posthumous representations^{16,32}. Technical frameworks for secure data sharing and consent management (including cryptographic and blockchain-based approaches) illustrate how privacy-by-design can be operationalized, although they must be adapted to the unique risks of generative agents³⁸.

Media narratives and public perceptions: Public discourse about AI afterlife is strongly mediated by cultural narratives and media representations. Studies of narrative framing show that AI resurrection is often portrayed through dystopian or sensational tropes, which can amplify moral panic or, conversely, normalize exploitative commercial offerings by presenting them as inevitable^{13,39}. Media stories also shape the moral vocabulary available to users: Whether they interpret an afterlife chatbot as a therapeutic tool, a spiritual medium, a deceptive fake, or a form of technological hubris^{25,32}.

Empirical work on acceptance helps ground these debates. Surveys and behavioral studies identify predictors such as perceived usefulness, religiosity, cultural norms about ancestor remembrance and trust in technology. These findings suggest that ethical governance cannot be one size fits all; instead, regulation and design need to be responsive to plural values and to the distinct vulnerabilities of bereaved populations³².

Governance and regulation: Governance proposals for AI afterlife often combine universal principles with culturally adaptive application. Proposed principles include protecting mental privacy, ensuring faithful identity representation, preventing exploitation and providing mechanisms for contestability and redress when harms occur^{14,16,24}. Comparative approaches recommend harmonized standards and mutual recognition agreements to address cross-border deployment, analogous to how other digital rights regimes coordinate internationally.

A recurring recommendation is to treat posthumous identity simulation as a high-risk AI application. This would trigger requirements for risk assessment, transparency, auditability and oversight. The literature also notes that compliance alone is insufficient: responsible practice requires participatory governance that includes families, religious communities, clinicians, ethicists and marginalized groups whose mourning practices may be poorly served by default platform assumptions^{16,25,32}.

CULTURAL INTERPRETATIONS AND IMPACT

Cultural scholarship emphasizes that AI afterlife is not a single phenomenon but a family of practices whose meanings differ across contexts. The same technical capability, e.g., an avatar that speaks in a familiar voice, can be interpreted as spiritual comfort, disrespectful imitation, or dangerous interference with the boundary between the living and the dead. Cultural difference matters not only for user acceptance but also for which harms are salient: in some contexts, the primary concern may be exploitation and privacy; in others, it may be ritual propriety, spiritual integrity, or communal authority over remembrance.

These insights challenge the implicit universalism in many AI design frameworks. If systems are built around Western notions of individual autonomy and self-expression, they may misfit settings where mourning is primarily communal, where ancestors occupy an ongoing moral role, or where speaking in the voice of the dead carries distinct religious constraints. The literature therefore advocates pluralistic design and cross-cultural governance as core requirements, not optional localization after the fact^{16,25,32}.

Cross-cultural perspectives on AI afterlife: Cross-cultural studies and comparative analyses indicate substantial variation in attitudes toward AI resurrection technologies. Western narratives often combine themes of transcendence and apocalypse, framing digital continuation as either salvation from death or a technologically driven distortion of human finitude³⁹. In contrast, some non-Western contexts

place greater emphasis on relational continuity like maintaining bonds through rituals and remembrance, potentially making interactive memorial technologies more socially intelligible.

Empirical work suggests that acceptance can be higher where cultural norms already support ongoing relationships with the deceased (e.g., through ancestor rituals or memorial practices), whereas contexts that emphasize strict separation between living and dead may experience stronger ontological dissonance when confronted with speaking avatars³². These patterns imply that AI afterlife design should allow for configurable interaction models: Some users may prefer limited, archival interaction; others may value dialogue-like remembrance within established ritual frameworks^{3,25}.

Ritual adaptation and spiritual meaning: Ritual and spiritual meaning emerge as central rather than peripheral issues. Researchers argue that many mourning practices are embodied and communal; AI can mediate certain forms of remembrance but cannot replicate the full social and sensory dimensions of ritual life. Consequently, the ethical legitimacy of AI afterlife may depend on whether systems respect ritual boundaries such as prohibitions on speaking for the dead, requirements for communal authorization, or norms about the timing and context of remembrance.

In religious and spiritual contexts, AI afterlife can also function as a kind of mediating technology that changes how people interpret messages, signs and memory. This creates opportunities for comfort and meaning-making, but it also introduces risks of spiritual manipulation or commercialization of sacred practices. Scholars therefore call for collaboration with faith communities and for design approaches that acknowledge spiritual plurality rather than assuming secular individualist defaults¹⁶⁻²⁵.

Influence of East Asian philosophies: East Asian philosophies provide concrete design principles. Confucian humanism emphasizes benevolence, role-based responsibilities and harmony, encouraging technologies that support family and community rather than replacing human relationships^{28,29}. Daoist perspectives caution against disruptive interventions that violate natural processes, suggesting a preference for gentle augmentation that aligns with existing practices rather than imposing new forms of technological presence^{28,30}.

These perspectives also reinforce a relational ethics of AI: Moral evaluation depends on how the system reshapes relationships and practices over time. Applied to AI afterlife, this implies assessing not only individual user satisfaction but also communal impacts, intergenerational dynamics and the ways platforms may reshape cultural norms about death and remembrance^{25,32}.

Media, generational and social factors: Beyond religious tradition, acceptance is shaped by media exposure, generational experience with digital communication and broader trust in institutions. Younger users accustomed to persistent digital identity and social media memorialization may find AI-mediated remembrance more plausible, while others may view it as an inappropriate extension of platform culture into sacred domains. Media narratives can intensify these divides by portraying AI afterlife either as a hopeful innovation or as a moral transgression¹³.

Because these factors vary within cultures as well as across them, the literature recommends adaptive engagement strategies: Participatory design, transparent communication of system limits and support resources for users who experience distress. This aligns with broader calls for responsible innovation that accounts for social context and long-term outcomes, not only immediate usability^{6,16}.

GAPS AND FUTURE RESEARCH DIRECTIONS

Because of rapid growth, the corpus reveals notable gaps. Philosophically, the debate over AI consciousness remains unresolved and it is unclear whether future technical advances will meaningfully change core objections about understanding and subjective experience. Ethically, there is limited longitudinal evidence about the psychological and social impacts of sustained interaction with posthumous agents, especially across diverse cultural contexts^{6,32}. Technically, while RAG and related methods can improve grounding, evaluation of persona faithfulness, safety and emotional appropriateness remains underdeveloped^{10,11}.

Future research directions therefore require genuinely interdisciplinary programs: Technical work on provenance and controllability; social science work on grief trajectories and community impacts; and philosophical-ethical work that clarifies concepts such as authenticity, dignity and meaning in relation to digital remembrance.

Unresolved philosophical and ethical questions: Unresolved questions cluster around what authenticity should mean in the afterlife context. If authenticity means identity continuity, most systems will fail; if it means faithful representation of a person's publicly shared values and communicative style, then authenticity becomes a governance and design problem rather than a metaphysical one. Another unresolved issue is the meaning of life and death under conditions of persistent digital presence. Some philosophical work on meaning in life and sustainable AI suggests that societies must evaluate whether technological continuation enriches or erodes the finitude that gives life narrative structure⁴⁰.

Ethically, the field lacks consensus on when AI afterlife should be discouraged outright. For example, if a system is used without prior consent, or if it targets vulnerable users with manipulative monetization, many authors argue that harms likely outweigh benefits. Yet there may be therapeutic or educational contexts where carefully constrained systems provide value. Clarifying these boundary conditions is a key task for future normative work^{6,25}.

Technical and governance challenges: Technically, a major challenge is aligning generation with verifiable evidence and declared policy constraints. Research agendas include robust memory management, secure retrieval, adversarial resilience and watermarking or provenance methods that allow users to distinguish sourced recollection from model invention. Governance challenges mirror technical ones: Regulators must decide which audits and documentation are required, how accountability is assigned when outputs cause harm and how to enforce cross-border protections for digital remains^{14,16,24}.

The literature increasingly calls for tiered regulation: Lower-risk archival memorials could face lighter requirements, while interactive agents that claim to speak as a person should be treated as high-risk and subject to stronger oversight, including pre-deployment evaluation, ongoing monitoring and clear processes for withdrawal or deletion upon request^{14,16}.

Cultural inclusion and pluralistic design: Cultural inclusion is an open research frontier. Existing studies show that moral intuitions about death technologies depend on culturally specific notions of personhood, ritual and community authority^{25,32}.

Pluralistic design therefore requires more than translating interfaces or adding cosmetic cultural motifs. It requires participatory methods that allow communities to define what respectful remembrance means, what kinds of interaction are acceptable and what governance structures should mediate conflicts among stakeholders.

A promising direction is modularity: Systems could offer adjustable levels of agency, different ritual modes and context-sensitive safeguards informed by local norms. This would allow AI afterlife to function as a flexible tool for remembrance rather than a single dominant model exported globally. Achieving this vision will require joint work across AI engineering, anthropology, philosophy and law^{16,32}.

SIGNIFICANCE STATEMENT

This review establishes a holistic framework integrating technical design, stakeholder responsibilities and lifecycle governance for AI afterlife technologies. By synthesizing diverse technological mechanisms with Western and non-Western ethical philosophies, it uncovers critical socio-technical tensions regarding identity rights and digital legacy. Ultimately, it provides an essential roadmap for developing culturally adaptive, safe and regulated posthumous digital human systems.

CONCLUSION

The AI afterlife emerges from the convergence of powerful generative technologies with enduring human desires for remembrance, continuity and connection. This review shows that such systems are not simply technical artifacts, but socio-technical forms that reshape how identity, mourning and posthumous presence are understood and enacted. Their promise lies in supporting memory, legacy and emotional connection, yet these same capacities create significant risks involving misrepresentation, dependency, privacy, manipulation and unequal governance of digital remains. Future research should therefore move beyond questions of technical feasibility toward longitudinal evaluation, culturally sensitive design and clearer models of consent, stewardship and accountability. The central task is to ensure that AI afterlife technologies serve as responsible memorial practices that protect both the dignity of the dead and the well-being of the living.

REFERENCES

1. Fink, M.C., S.A. Robinson and B. Ertl, 2024. AI-based avatars are changing the way we learn and teach: Benefits and challenges. *Front. Educ.*, Vol. 9. 10.3389/educ.2024.1416307.
2. Pushkarsky, A.G., 2025. The importance of Kant's philosophy of mind for contemporary research in artificial intelligence. *Rudn J. Philos.*, 29: 473-490.
3. Brubaker, J.R., G.R. Hayes and P. Dourish, 2013. Beyond the grave: Facebook as a site for the expansion of death and mourning. *Inf. Soc.*, 29: 152-163.
4. Bojić, L., I. Stojković and Z.J. Marjanović, 2024. Signs of consciousness in AI: Can GPT-3 tell how smart it really is? *Humanit. Social Sci. Commun.*, Vol. 11. 10.1057/s41599-024-04154-3.
5. Martynenko, N.P., 2025. Cognitive Mechanisms of Large Language Models: Interaction with GigaChat [In Russian]. *Concept Philos. Religion Cult.*, 9: 30-50.

6. Abu-elezz, I., A. Hassan, A. Nazeemudeen, M. Househ and A. Abd-alrazaq, 2020. The benefits and threats of blockchain technology in healthcare: A scoping review. *Int. J. Med. Inf.*, Vol. 142. 10.1016/j.ijmedinf.2020.104246.
7. Adams, R., 2023. Artificial Intelligence Elsewhere: The Case of the *Ogbanje*. In: *Imagining AI: How the World Sees Intelligent Machines*, Cave, S. and K. Dihal (Eds.), Oxford University Press, Oxford, England, ISBN: 9780192865366, pp: 261-274.
8. Park, G., J. Chung and S. Lee, 2024. Human vs. machine-like representation in chatbot mental health counseling: The serial mediation of psychological distance and trust on compliance intention. *Curr. Psychol.*, 43: 4352-4363.
9. Chemero, A., 2023. LLMs differ from human cognition because they are not embodied. *Nat. Hum. Behav.*, 7: 1828-1829.
10. Ali Bin Afarin, M., M. Adham Bin Isa, S.Z.M. Hashim, H.N. Abdull Hamed and A. Matthew, 2026. Mapping the intellectual landscape of retrieval-augmented generation (RAG): A bibliometric analysis. *KSII Trans. Internet Inf.Syst.*, 20: 284-307.
11. Dournaee, B.H., 2010. Comments on "The replication of the hard problem of consciousness in AI and bio-AI". *Mines Mach.*, 20: 303-309.
12. Cunneen, M., R. AnandFinn, R. Friel, P. Tennent and S. Brandt, 2026. From bones to bytes: Anticipating and addressing the governance challenges of human digital remains and posthumous digital human twins. *AI Soc.*, 41: 2021-2040.
13. Kumar, K., 2025. Fundamentals of Generative AI. In: *Generative AI for Photonic Sensing*, Vasimalla, Y. and S. Kumar, Springer Nature, Singapore, ISBN: 978-981-95-1561-5, pp: 33-77.
14. Ouariach, F.Z. and S. Ahmed, 2025. Ethical and Legal Dimensions of AI in Healthcare. In: *Computational Intelligence in Healthcare Law*, Singh, B., C. Kaunert, B. Balusamy and R.K. Dhanaraj (Eds.), Springer Nature, Singapore, ISBN:9781003538462, pp: 86-98.
15. Zhu, Q., 2023. Confucian Ethics of Technology. In: *Technology Ethics: A Philosophical Introduction and Readings*, Robson, G.J. and J.Y. Tsou (Eds.), Routledge, New York, USA, ISBN: 9781003189466, pp: 93-101.
16. Denia, E., 2025. AI narratives model: Social perception of artificial intelligence. *Technovation*, Vol. 146. 10.1016/j.technovation.2025.103266.
17. Snyder, H., 2019. Literature review as a research methodology: An overview and guidelines. *J. Bus. Res.*, 104: 333-339.
18. Cumpston, M., T. Li, M.J. Page, J. Chandler, V.A. Welch, J.P.T. Higgins and J. Thomas, 2019. Updated guidance for trusted systematic reviews: A new edition of the Cochrane handbook for systematic reviews of interventions. *Cochrane Database Syst. Rev.*, Vol. 10. 10.1002/14651858.ED000142.
19. DaCosta, B., 2025. Speaking with the past: Constructing AI-generated historical characters for cultural heritage and learning. *Heritage*, Vol. 8. 10.3390/heritage8090387.
20. Searle, J., 2006. Chinese Room Argument, The. In: *Encyclopedia of Cognitive Science*, Nadel, L. (Ed.), Wiley, New Jersey, ISBN: 9780470018866.
21. Candiottto, K.B.B., 2025. The Chinese room argument revisited: Considerations from generative AI [In Portuguese]. *Aurora Philos. Mag.*, Vol. 37. 10.1590/2965-1557.037.e202532317.
22. Kawashima, D., T. Kempe, Y. Koga, S. Kubota and Y. Urata *et al.*, 2025. Attitudes toward digital immortality and related factors among older adults: An exploratory study from a developmental perspective. *OMEGA J. Death Dying*, 10.1177/00302228251336498.
23. Brown, A., M. Roman and B. Devereux, 2025. A systematic literature review of retrieval-augmented generation: Techniques, metrics, and challenges. *Big Data Cognitive Comput.*, Vol. 9. 10.3390/bdcc9120320.
24. González-Arias, C. and X. López-García, 2025. Technopessimism in the Spanish print media in the public debate on the social impact of artificial intelligence [In Portuguese]. *Austral Comunicación*, Vol. 14. 10.26422/aucom.2025.1401.gon.
25. Sciolla, J.C. and U. Pagallo, 2025. No peace after death? The impact of AI-driven memorial chatbots on privacy and data protection. *Information*, Vol. 16. 10.3390/info16060426.
26. Besser, A., T. Morse and V. Zeigler-Hill, 2023. Who wants to (digitally) live forever? The connections that narcissism has with motives for digital immortality and the desire for digital avatars. *Int. J. Environ. Res. Public Health*, Vol. 20. 10.3390/ijerph20176632.
27. Kavathatzopoulos, I., 2023. Artificial Intelligence and the Sustainability of Thinking. In: *Ethics and Sustainability in Digital Cultures*, Lennerfors, T.T. and K. Murata (Eds.), Routledge, London, United Kingdom, ISBN: 9781003367451, pp: 19-30.
28. Obadia, L., 2026. Deadbots as New Fetish of Postmortal Societies: A "Digital Decolonization" of Death? In: *Decolonising Death Studies*, Pentaris, P., S. Pitsillides and H. Ghorbani (Eds.), Routledge, London, United Kingdom, ISBN: 9781003544104, pp: 74-87.
29. Adil, A.M., G. Singh and M.K. Singh, 2026. Fundamentals of Artificial Intelligence. In: *Nanotechnology for Nucleic Acid Delivery*, Sharma, P., S. Jain and V. Jain (Eds.), Apple Academic Press, New York, USA, ISBN: 9781998511853, pp: 167-198.
30. Hou, Y., 2026. Natural intelligence, not artificial: A confucian reframing of generative AI in higher education. *Ethics Educ.*, 21: 73-91.
31. Lu, H., 2024. Exploring the predictors of public acceptance of artificial intelligence-based resurrection technologies. *Technol. Soc.*, Vol. 78. 10.1016/j.techsoc.2024.102657.
32. Pelea, C.I., 2022. Mirroring cultural fear, anxiety and dystopia in American cinematography: The movie "A.I." (2001). *Colloq. Humanistica*, Vol. 11. 10.11649/ch.2631.

33. Lahoz-Beltra, R., 2026. Can generative AI co-evolve with human guidance and display non-utilitarian moral behavior? *Computation*, Vol. 14. 10.3390/computation14020040.
34. Akinyemi, C. and A. Hassett, 2021. "He's still there": How facebook facilitates continuing bonds with the deceased. *Omega J. Death Dying*, 88: 620-637.
35. Nowak, K.L. and J. Fox, 2018. Avatars and computer-mediated communication: A review of the definitions, uses, and effects of digital representations on communication. *Rev. Commun. Res.*, 6: 30-53.
36. Yang, N.Y. and G.J. Khanna, 2025. AI and technology in grief support: Clinical implications and ethical considerations. *Couns. Psychologist*, 53: 579-604.
37. Amugongo, L.M., P. Mascheroni, S. Brooks, S. Doering and J. Seidel, 2025. Retrieval augmented generation for large language models in healthcare: A systematic review. *PLOS Digital Health*, Vol. 4. 10.1371/journal.pdig.0000877.
38. Floridi, L., J. Cows, M. Beltrametti, R. Chatila and P. Chazerand *et al.*, 2018. AI4People-an ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds Mach.*, 28: 689-707.
39. Edwards, D.J., 2025. Further *N*-Frame networking dynamics of conscious observer-self agents via a functional contextual interface: Predictive coding, double-slit quantum mechanical experiment, and decision-making fallacy modeling as applied to the measurement problem in humans and AI. *Front. Comput. Neurosci.*, Vol. 19. 10.3389/fncom.2025.1551960.
40. R  ther, M., 2024. Why care about sustainable AI? Some thoughts from the debate on meaning in life. *Philos. Technol.*, Vol. 37. 10.1007/s13347-024-00717-z.