Designing AI Afterlife Agents: A Socio-technical Design Space, Lifecycle Model, and Prompt Architecture

Thomas Keller and Mirella Moser
ZHAW Institute of Information Systems, Winterthur, Switzerland

kell@zhaw.ch
mosl@zhaw.ch

Abstract: AI afterlife systems transform digital traces of deceased persons into interactive posthumous representations. Current literature describes these systems as griefbots, deathbots, ghostbots, postmortem avatars, digital resurrection, and digital legacy, but the field remains fragmented across information systems, psychology, ethics, philosophy, theology, and anthropology. Using an interdisciplinary conceptual design approach, this paper develops a framework for AI Afterlife Agents as socio-technical systems of posthumous presence. It contributes three linked artifacts: a socio-technical design-space model, a lifecycle model, and a prompt/state-machine architecture. The design-space model explains how identity-corpus construction, interaction modality, user interpretation, psychological and cultural effects, and governance mechanisms jointly shape risks and experiences. The lifecycle model operationalizes these dimensions across data elicitation, directed recognizability training, bereavement use, and decommissioning. The paper argues that responsible AI afterlife design requires traceability between personal data, consent, prompts, safeguards, user boundaries, and retirement procedures. Next steps involve stakeholder validation, prototype-based experimentation, and longitudinal research with data donors and authorized bereaved users.

Keywords: AI afterlife, Griefbots, Digital legacy, Posthumous presence, Responsible AI, Prompt architecture

1. Introduction to AI Afterlife

Digital technologies have long mediated death, mourning, and remembrance through photographs, memorial websites, social media profiles, archived messages, and online memorial spaces. Generative AI now intensifies this mediation by enabling interactive representations that simulate a deceased person's voice, appearance, communication style, memories, values, or relational presence. Digital legacy research therefore increasingly treats death-related systems not only as archives, but as infrastructures that move personal data across generations and social contexts (Anna Puzio, 2023; Doyle & Brubaker, 2023).

The literature uses overlapping terms. Griefbots and deathbots usually refer to conversational agents based on digital traces of the deceased; postmortem avatars and digital resurrection emphasize visual, voice, and personality reconstruction; and AI afterlife highlights the broader socio-technical field in which the dead remain interactable through computational systems (Belén Jiménez-Alonso & Ignacio Brescó de Luna, 2022; Hollanek & Nowaczyk-Basińska, 2024; Joel Krueger & Lucy Osler, 2022). These systems may support reminiscence, legacy preservation, storytelling, and continuing bonds, but they may also intensify dependency, misrepresentation, unauthorized resurrection, posthumous privacy violations, and exploitative commercialization (Harbinja et al., 2023; Lindemann, 2022).

This paper uses the term AI Afterlife Agent for an AI-mediated representation of a deceased or future-deceased person that is built from an identity corpus and made available for bounded interaction by authorized users. The central argument is that such agents cannot be evaluated only as chatbots or only as memorial media. They are socio-technical systems in which data provenance, model behavior, cultural meaning, grief vulnerability, access control, and retirement decisions jointly determine whether posthumous presence is experienced as supportive, inappropriate, or harmful.

2. Objective and Research Question

The objective is to propose a traceable conceptual architecture for responsible AI Afterlife Agents. The paper integrates three levels of design reasoning. First, a design-space model identifies the key dimensions that shape AI afterlife systems: identity corpus, interaction modality, user interpretation, psychological and cultural effects, and governance. Second, a lifecycle model translates these dimensions into four phases: data elicitation, directed recognizability training, bereavement use, and decommissioning. Third, a prompt/state-machine architecture specifies how system prompts can operationalize consent, provenance, recognizability calibration, grief-sensitive interaction, risk monitoring, and responsible retirement.

The resulting research question is: How can AI Afterlife Agents be designed so that identity construction, interaction, interpretation, safeguards, and decommissioning remain accountable across the whole posthumous

lifecycle? The expected contribution is not a final product blueprint, but a theoretically grounded design framework that can guide empirical research, prototype development, ethics review, and governance discussions.

3. Research Approach

This study follows a conceptual, design-oriented research approach. Its purpose is not to empirically evaluate an implemented AI Afterlife Agent, but to synthesize interdisciplinary knowledge into theoretically grounded design artifacts that can guide subsequent implementation and validation. The analysis integrates Information Systems, psychology, ethics and philosophy, and theology and anthropology because AI afterlife systems simultaneously involve technical interaction, identity representation, grief-related processes, posthumous rights, and culturally situated practices.

The framework was developed in three analytical steps. First, the interdisciplinary literature was synthesized to identify recurring design concerns for AI-mediated posthumous presence; these concerns were consolidated into five interdependent dimensions: identity corpus, interaction modality, user interpretation, psychological and cultural effects, and governance mechanisms. Second, these dimensions were translated into a temporal lifecycle model that assigns design and governance concerns to four phases: data elicitation, directed recognizability training, bereavement use, and decommissioning. Third, the lifecycle requirements were operationalized as a prompt/state-machine architecture specifying phase-specific interaction states, persistent system objects, governance checks, and auditable outputs.

The resulting artifacts should therefore be understood as analytically derived design propositions rather than empirically validated system components. Their purpose is to establish traceability from interdisciplinary requirements to lifecycle phases and system behavior. Empirical validation through expert and stakeholder interviews, prototype experiments, and longitudinal studies is left for subsequent research.

4. State-of-the-art

The state of research on AI afterlife Agents is highly interdisciplinary and conceptually fragmented. Relevant work is distributed across Information Systems, Psychology, Ethics and Philosophy, and Theology and Anthropology, each of which highlights different aspects of posthumous AI representation.

4.1 Information Systems

Information Systems research provides the foundation for analyzing AI afterlife systems as interactive socio-technical systems. Social presence theory explains why users may experience mediated others as socially "there," while recent work clarifies that presence is a multidimensional perception shaped by salience, sociability, interaction integration, and divergence between expected and experienced presence (Cummings & Bailenson, 2016; Kreijns et al., 2022, 2022). For immersive or embodied afterlife agents, telepresence, co-presence, and robot-mediated interaction become relevant because voice, embodiment, reciprocity, and observability can strengthen perceived sociality (Almeida et al., 2022).

The Media Equation and Computers as social Actors (CASA) traditions further explain why people respond socially to technologies. Contemporary revisions emphasize that anthropomorphism and social response are not automatic; they depend on cues, contexts, user expectations, and the distinction between machines as social actors and machines as evoking sociality (Gambino et al., 2020; Lombard & Xu, 2021). Acceptance research also matters, yet standard UTAUT-style variables such as usefulness, effort, trust, and anxiety must be extended for AI afterlife because adoption is shaped by grief, authenticity, moral concern, cultural fit, and posthumous consent (Blut et al., 2022; Philippi et al., 2021).

4.2 Psychology

Psychology shows why AI afterlife systems intervene in attachment, continuing bonds, emotional regulation, and meaning-making. Attachment-based grief research emphasizes that responses to loss depend on relational history, coping style, identity distress, and shattered assumptions (Captari et al., 2021; Mikulincer & Shaver, 2022). Continuing bonds theory challenges the idea that grief requires detachment; ongoing symbolic and relational ties to the deceased can be adaptive when they support meaning and integration (She et al., 2021).

AI afterlife changes the form of continuing bonds because interaction becomes conversational, personalized, and potentially reciprocal. Chatbots may support habits of intimacy with the dead, but a desire for virtual reunion may also signal vulnerability to complicated grief or prolonged dependency (Joel Krueger & Lucy Osler,

2022; Kawashima et al., 2025). Psychoanalytic and clinical perspectives warn that griefbots could obstruct acceptance of loss if the system presents simulation as continued life or if users lack boundaries around interaction (Lemma, 2024). Evidence from online grief interventions suggests that digital support can reduce distress, but AI afterlife requires different safeguards because the system embodies the deceased rather than merely offering bereavement support (Zuelke et al., 2021).

4.3 Ethics and Philosophy

Ethical and philosophical work foregrounds dignity, consent, autonomy, authenticity, privacy, personhood, and posthumous agency. Care-oriented approaches argue that digital remains require relational stewardship rather than simple ownership, especially because AI afterlife systems operate in situations of dependency, vulnerability, and emotional asymmetry (Dinneen et al., 2024; Hollanek & Nowaczyk-Basińska, 2024). Lindemann (2022) shifts attention to the autonomy and dignity of bereaved users, noting risks of manipulation, dependency, and therapeutic overclaiming.

Authenticity and identity are central because AI afterlife agents may appear to speak as a person who can no longer correct the representation. Digital legacy research treats posthumous data as identity manifestations, while philosophical accounts of digital replicas ask whether psychological connectedness, relational recognition, or value continuity can justify claims of extended personhood (Doyle & Brubaker, 2023; Karpus & Strasser, 2025). Legal-ethical work shows that current privacy, property, personality, and data-protection regimes only partially address ghostbots, deletion rights, and posthumous data governance (Ciani Sciolla & Pagallo, 2025; Harbinja et al., 2023).

4.4 Theology and Anthropology

Theology and anthropology broaden the analysis by showing that AI afterlife systems enter existing practices of ritual, ancestorhood, memorialization, and spiritual continuity. HCI scholarship has identified dying, death, and afterlife as an emerging design domain, while cultural theories of digital afterlife emphasize affective assemblages of data, platforms, memory, users, and commercial practices (Albers et al., 2023; Anu A. Harju, 2024). Digital immortality research also shows ambivalence: people may value preservation and self-continuity while fearing inauthentic simulation and loss of data control (Galvão et al., 2021).

Cultural perspectives caution against universal design assumptions. Grief emotions, ritual authority, remembrance practices, and relations with the dead vary across communities and religious traditions (Silverman et al., 2021). Digital photographs, memorial pages, and AI-mediated afterlife systems can become ritual objects, but they may also shift collective memory into platform-controlled and commercially mediated forms (Meitzler et al., 2024; Öhman & Floridi, 2017). The state of the art therefore points to a research gap: AI afterlife design needs models that integrate technical affordances with grief psychology, posthumous rights, ritual appropriateness, and lifecycle governance.

5. Conceptualization: AI Afterlife Agents as Socio-technical Systems

The preceding state-of-the-art review shows that AI afterlife cannot be sufficiently explained from a single disciplinary perspective. The conceptualization proposed in this chapter is therefore a logical consequence of the interdisciplinary state of the art. If AI afterlife systems simultaneously operate as technical artifacts, relational media, grief-related interventions, ethical objects, and culturally situated memorial practices, they cannot be conceptualized merely as chatbots, archives, or digital replicas. They must instead be understood as socio-technical systems of posthumous presence. This framing integrates the four lenses reviewed above and provides the theoretical basis for the subsequent design-space model, lifecycle model, and prompt architecture.

Hence, AI Afterlife Agents are conceptualized here as socio-technical systems of posthumous presence. Figure 1 frames this conceptualization as an interdisciplinary design space shaped by Information Systems, Psychology, Ethics and Philosophy, and Theology and Anthropology.

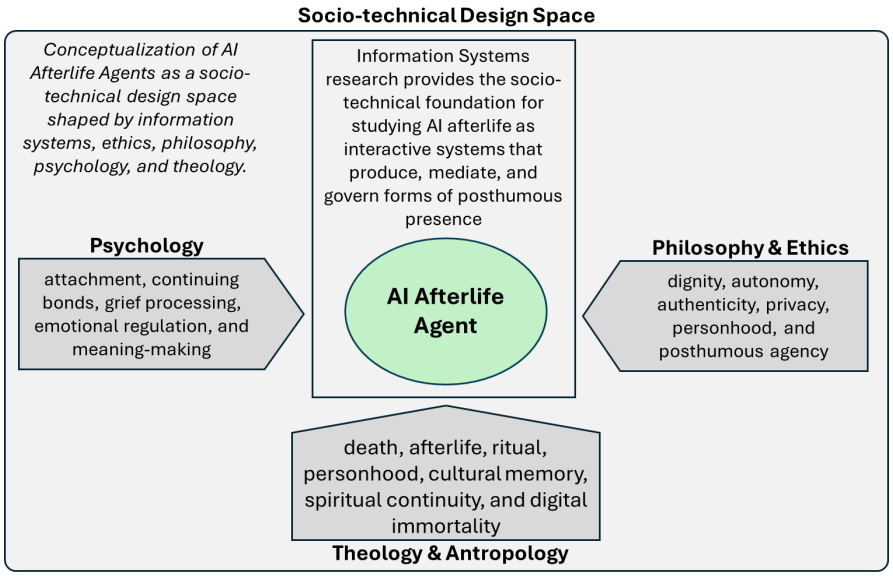

Figure 1: Conceptualization of the socio-technical design space

6. Design Space Model

The design-space model turns the conceptualization into a practical analytic framework. The reason for introducing such a model is that the state-of-the-art review does not point to one dominant design variable, but to several interdependent sources of risk and experience. A design-space model (Figure 2) is therefore needed to make these interdisciplinary dependencies explicit and actionable. Without such a model, AI afterlife systems risk being evaluated mainly through technical fidelity, for example whether the agent sounds, writes, or looks like the deceased person. However, technical fidelity alone does not explain whether an interaction is experienced as comforting, disturbing, legitimate, manipulative, culturally appropriate, or ethically acceptable. These outcomes emerge from the relationship between the identity corpus, the interaction modality, the user's interpretation of the agent, the psychological and cultural effects of use, and the governance mechanisms that regulate access, consent, transparency, safeguards, and retirement.

The model therefore assumes that AI afterlife experiences are produced by dynamic interaction among five dimensions rather than by technical fidelity alone. It provides a structured way to analyze how design decisions in one dimension, such as increasing embodiment or expanding the identity corpus, may create consequences in other dimensions, such as stronger perceived presence, higher privacy risk, altered grief boundaries, or greater need for governance. In this sense, the design-space model functions as a bridge between the interdisciplinary conceptualization in Chapter 4 and the lifecycle architecture developed in Chapter 6.

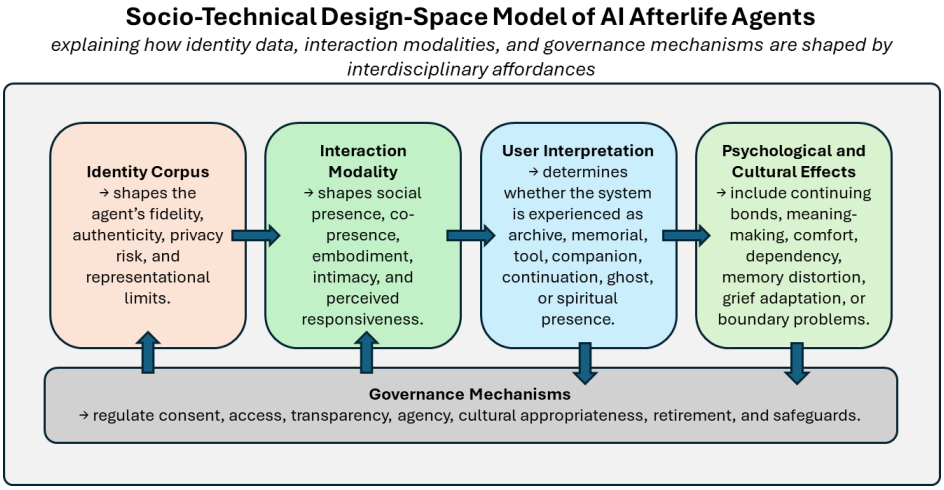

Figure 2: Socio-technical design-space model of AI Afterlife Agents

The identity corpus is the central design object. It should include source type, provenance, consent status, relational scope, third-party data, values, communication style, memories, boundaries, and legacy intentions. A larger corpus may improve stylistic resemblance, but it can also increase privacy risks and false authority if the model infers beyond documented material. Therefore, design should distinguish archived statements from generated extrapolations and should mark uncertainty when the agent lacks evidence for a response (Ciani Sciolla & Pagallo, 2025; Doyle & Brubaker, 2023).

Interaction modality determines how intensely the agent is encountered. Text chat may support reflection with relatively low sensory realism; voice and avatar modalities may create stronger co-presence and emotional salience; and immersive embodiments may blur boundaries between memory, simulation, and perceived encounter. The design space therefore treats modality as a psychological and ethical variable, not merely a user-interface choice. Higher social presence may support remembrance, but it also increases obligations for disclosure, access control, usage boundaries, and de-escalation (Almeida et al., 2022; Joel Krueger & Lucy Osler, 2022; Lindemann, 2022).

User interpretation is the hinge between design and effect. The same agent may be understood by one user as a memorial archive, by another as a companion, and by another as a spiritually problematic imitation. These interpretation frames shape grief outcomes, legitimacy judgments, and decommissioning preferences. Governance is therefore a cross-cutting layer. It should specify who may create, access, calibrate, suspend, transfer, or delete the agent; how conflicts among relatives are resolved; how users are reminded that the system is artificial; and how cultural or ritual constraints are respected (Harbinja et al., 2023; Silverman et al., 2021).

7. Lifecycle Model

The lifecycle model (Figure 3) complements the design-space model (Figure 2) by translating its abstract dimensions into a temporal and operational view of AI afterlife design. While the design-space model explains which dimensions shape AI afterlife experiences, the lifecycle model explains when, by whom, and under which governance conditions these dimensions become relevant. This is necessary because AI Afterlife Agents are not created, used, and retired in a single moment. They evolve across different phases, involve different actors, and require different safeguards before and after the death of the data donor.

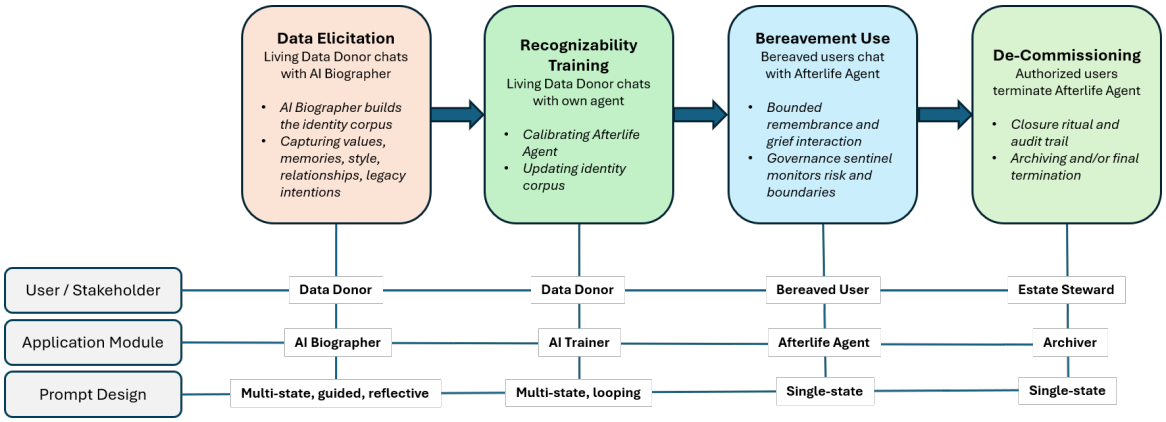

Figure 3: Lifecycle architecture of an AI Afterlife Agent

A lifecycle view is needed because many of the central risks identified in the state of the art are phase-dependent. Consent, provenance, and boundary setting are most critical during data elicitation; recognizability and authenticity require calibration while the data donor is still alive; bereavement use introduces grief-sensitive interaction, access control, and risk monitoring; and decommissioning raises questions of closure, deletion, archiving, auditability, and platform dependency. Without a lifecycle perspective, these issues risk being treated as general ethical principles rather than as concrete design and governance tasks assigned to specific system phases.

The lifecycle model therefore operationalizes the design-space model as a four-phase architecture: data elicitation, directed recognizability training, bereavement use, and decommissioning. It shows how identity-corpus construction, interaction modality, user interpretation, psychological and cultural effects, and governance mechanisms unfold over time. However, a lifecycle view alone still remains too abstract unless it

specifies how these requirements are enacted in the actual behavior of the system. This is why the lifecycle model requires an additional architectural layer: the prompt/state-machine architecture (Figure 4).

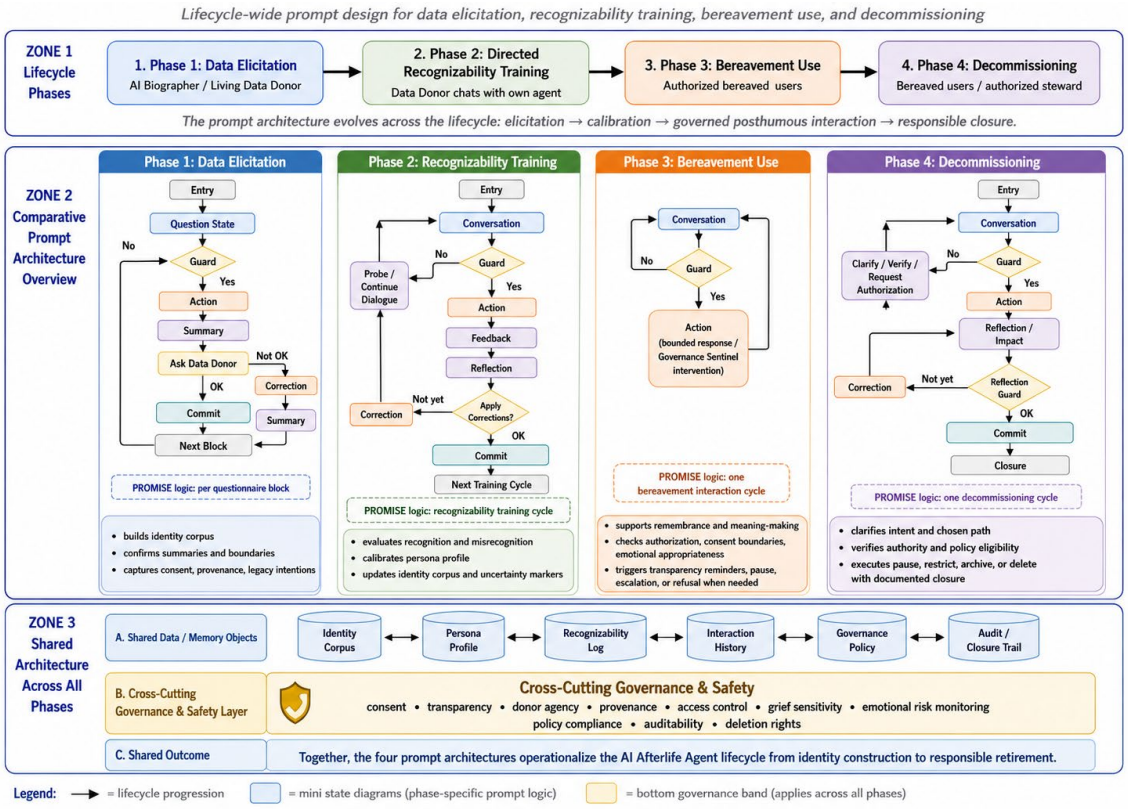

Figure 4: Overview architecture of the prompt designs across the four lifecycle phases

The need for this architecture follows from the specific nature of AI Afterlife Agents. Such agents do not only retrieve information from a database; they conduct sensitive biographical elicitation, simulate aspects of a person's communicative identity, interact with bereaved users, monitor boundaries and distress, and eventually support retirement or deletion. These tasks cannot be left to a single general-purpose prompt because each lifecycle phase involves different actors, permissions, risks, and outputs. Data elicitation requires structured interviewing and consent capture; recognizability training requires iterative feedback and correction; bereavement use requires grief-sensitive interaction and governance monitoring; decommissioning requires authority checks, closure options, and auditability. A prompt/state-machine architecture is therefore needed to make system behavior phase-specific, traceable, and governable.

In this paper, "prompt/state-machine logic" refers to the use of structured prompt flows in which the agent moves through defined states such as entry, consent check, interaction, guardrail assessment, action, summary, confirmation, correction, escalation, pause, or closure. The purpose is not to reduce the agent to a rigid script, but to constrain generative flexibility within accountable interaction states. This allows the system to document which data were used, which consent conditions applied, which user role was active, which safeguards were triggered, and which outputs were created. The architecture therefore connects the conceptual dimensions of the design-space model with concrete system operations. State transitions are executed against shared data and governance objects rather than encoded only in natural-language instructions, so the architecture separates generative content production from authorization, provenance, memory, and safety checks.

Figure 4 and Table 1 introduce this architecture across the four lifecycle phases. The upper layer shows the temporal progression from data elicitation to recognizability training, bereavement use, and decommissioning. The middle layer specifies the phase-specific prompt/state-machine flows: a guided multi-state elicitation flow, a looping recognizability-training flow, a grief-sensitive bereavement interaction flow, and a closure-oriented decommissioning flow. The lower layer shows the shared data, memory, governance, and audit objects that persist across phases. In this sense, the prompt/state-machine architecture is not an additional independent contribution detached from the lifecycle model. It is the operational mechanism through which the lifecycle model becomes implementable, auditable, and ethically constrained.

Table 1: Lifecycle phases, actors, prompt-state roles, and outputs

Phase	*Main actor*	*Prompt/state-machine role*	*Primary output*
1. Data elicitation	Living data donor	AI Biographer conducts structured biographical interviewing with adaptive follow-up, confirmation, provenance capture, consent checks, and boundary extraction.	Structured identity corpus with consent, provenance, relational boundaries, and legacy intentions.
2. Directed recognizability training	Living data donor	Agent is tested by the donor; feedback loops identify recognition, misrecognition, omissions, overreach, and uncertainty.	Calibrated profile, correction log, recognizability notes, and uncertainty markers.
3. Bereavement use	Authorized bereaved users	Grief-sensitive interaction prompt provides disclosure, bounded remembrance, source awareness, risk monitoring, de-escalation, and support referral.	Interaction logs, safety interventions, boundary reminders, and pause or retirement signals.
4. Decommis-sioning	Bereaved users, estate stewards, governance authority	Closure flow verifies authority, clarifies intent, reviews memory or audit status, offers ritual closure, and confirms archive, restriction, deletion, or memorial mode.	Retired, archived, restricted, deleted, or memorialized agent with audit trail and closure message.

7.1 Identity Corpus Architecture

The identity corpus is specified as a logical, storage-agnostic architecture rather than as a single database. It comprises four linked layers: (1) source objects containing original digital traces such as messages, recordings, photographs, documents, and elicitation transcripts; (2) semantically indexed memory units for retrieval of relevant evidence; (3) structured identity attributes and relationships representing biographical facts, values, communication traits, relational roles, and explicit boundaries; and (4) governance metadata recording provenance, consent status, permitted audiences, confidence, third-party restrictions, and retention or deletion conditions.

A practical implementation may combine vector retrieval for semantically indexed memories with graph or relational storage for structured identity and governance information. The framework does not require a particular database technology. It requires traceability: generated identity claims should be linked to supporting source objects and the authorization conditions under which those objects may be used. When the corpus does not provide sufficient evidence, the agent should mark uncertainty rather than convert model inference into apparent biographical fact.

7.2 State Transition Model

The prompt/state-machine architecture uses local state machines within each lifecycle phase rather than one global finite-state machine. Four transition types recur across the phases: progression transitions move the interaction forward when required information or confirmation is complete; loop transitions return to elicitation or calibration when information remains incomplete or recognizability is insufficient; guardrail transitions interrupt the normal flow when a consent, boundary, authorization, or distress condition is detected; and terminal transitions close the relevant flow through completion, pause, archive, deletion, or other authorized retirement action.

Transitions are conditional and auditable. Before a generative action, the system evaluates the active user role, applicable consent and relational scope, provenance requirements, current lifecycle phase, and relevant safety state. Shared objects - identity corpus, consent profile, user-role permissions, correction log, interaction history, safety policy, and closure trail - persist across the phase-specific state machines. Table 2 illustrates the transition logic for bereavement use, the phase in which interaction and grief-sensitive safeguards are most tightly coupled.

Table 2: Illustrative transition model for the bereavement-use state machine

Current state	*Trigger / condition*	*Next state*	*Required system action*
Entry	Authorized user session begins	Disclosure	Load role, permissions, applicable consent, and interaction boundaries.
Disclosure	User acknowledges that the system is an AI representation	Interaction	Enable bounded remembrance interaction and source-aware retrieval.

Interaction	Request is within consent and boundary scope	Response generation	Retrieve permitted evidence, distinguish recall from inference, and generate a bounded response.
Interaction	Request exceeds consent, provenance, or relational boundary	Boundary reminder	Refuse or constrain the request and explain the applicable boundary.
Interaction	Potential distress or dependency signal is detected	Guardrail assessment	Suspend ordinary persona progression and assess the relevant safety rule.
Guardrail assessment	Moderate concern; interaction may continue safely after interruption	Pause / de-escalation	Reduce immersive cues, suggest a pause, and provide a non-authoritative support option.
Guardrail assessment	High concern or policy threshold is reached	Escalation / termination	Stop ordinary simulation behavior and route to the defined human-support or safety pathway.

In Phase 1, the AI Biographer should elicit identity data as a consent-bound corpus rather than as an unlimited memory dump. A PROMISE-based prompt architecture (Wu et al., 2024) can organize the interview into blocks on biography, values, relationships, communication style, memories, humor, boundaries, spirituality or cultural preferences, third-party privacy, and legacy intentions. Each block should summarize what was captured and ask for confirmation before extraction. This reduces later ambiguity about what the agent may say, what it must avoid, and which sources support each memory or trait.

In Phase 2, directed recognizability training addresses a problem that is rarely solved by posthumous systems: the represented person is normally unavailable to validate the representation. When the data donor is still alive, they can interact with the prototype agent, correct misrecognition, label uncertain claims, and calibrate tone, values, humor, and relational boundaries. The output is not only a more recognizable persona but also an audit trail of corrections and limits. This phase converts authenticity from a subjective impression into a documented calibration process.

In Phase 3, bereavement use begins only for authorized users and should operate under grief-sensitive guardrails. The agent should disclose that it is an AI representation, distinguish evidence-based recall from generated inference, refuse to create new authoritative wishes, avoid manipulative intimacy, and recognize distress signals. Prompt states can include ordinary remembrance, boundary reminder, pause suggestion, de-escalation, and referral to human support. Such safeguards respond to evidence that digital bonds may be supportive or maladaptive depending on form, function, and user vulnerability (Kawashima et al., 2025).

In Phase 4, decommissioning treats ending the agent as a designed and meaningful process. Responsible systems should not assume perpetual availability. The prompt flow should clarify authority, confirm the user's intent, explain consequences, offer archive, restriction, deletion, or memorial-only modes, and create a closure message if desired. Decommissioning also addresses platform power because users need credible mechanisms for portability, auditability, and verifiable deletion rather than indefinite subscription-based dependency (Öhman & Floridi, 2017).

8. Conclusion and Next Steps

AI Afterlife Agents are best understood as socio-technical systems of posthumous presence. Their risks and benefits cannot be derived from model capability alone. They depend on the quality and legitimacy of the identity corpus, the intensity of the interaction modality, the user's interpretation of the agent, the cultural and psychological context of use, and the governance mechanisms that regulate access, consent, boundaries, and retirement.

This paper contributes three linked artifacts grounded in an interdisciplinary socio-technical conceptualization: a design-space model, a lifecycle model, and a PROMISE-based prompt/state-machine architecture. The framework reframes AI afterlife design around governable authenticity: representations should be recognizable, but also transparent, evidence-aware, consent-bound, culturally sensitive, and decommissionable. This shifts the design question from whether the dead can be simulated to how posthumous representation can remain accountable to the living, the deceased, and affected communities.

The next step is empirical validation. Phase A should use expert and stakeholder interviews to challenge the design-space dimensions and produce a traceability matrix linking interdisciplinary requirements to lifecycle phases. Phase B should test prompt mechanisms in controlled prototypes, examining identity-corpus quality, recognizability calibration, grief-sensitive interaction, governance perception, and decommissioning choices. Phase C should conduct longitudinal studies with data donors and authorized recipient users to investigate whether recognizability, perceived authenticity, safety expectations, and willingness to retire the agent remain stable or change across the lifecycle.

Ethics Declaration: No ethics approval was required for this research.

AI Declaration: The authors used ChatGPT (OpenAI, 2026) and Consensus.app to support literature search and language refinement. All conceptual arguments, literature selection decisions, interpretations, and final manuscript content were reviewed, verified, and approved by the authors.

9. References

Albers, R., Sadeghian, S., Laschke, M., Hassenzahl, M., & ACM. (2023). Dying, Death, and the Aferlife in Human-Computer Interaction. A Scoping Review. (WOS:001048393801028). PROCEEDINGS OF THE 2023 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS (CHI 2023). https://doi.org/10.1145/3544548.3581199

Almeida, C., Kalinowski, M., Uchôa, A., & Feijó, B. (2022). Datasets of "Negative Effects of Gamification in Education/Learning Systems: Systematic Mapping and Practitioner Perceptions." https://doi.org/10.5281/ZENODO.6279062

Anna Puzio. (2023). When the Digital Continues After Death. Ethical Perspectives on Death Tech and the Digital Afterlife. *Communicatio Socialis*. https://doi.org/10.5771/0010-3497-2023-3-427

Anu A. Harju. (2024). Theorising Digital Afterlife as Techno-Affective Assemblage: On Relationality, Materiality, and the Affective Potential of Data. *Social Sciences*. https://doi.org/10.3390/socsci13040227

Belén Jiménez-Alonso & Ignacio Brescó de Luna. (2022). Griefbots. A New Way of Communicating With The Dead? *Integrative Psychological and Behavioral Science*, *57*, 466–481. https://doi.org/10.1007/s12124-022-09679-3

Blut, M., Chong, A., Tsigna, Z., & Venkatesh, V. (2022). Meta-Analysis of the Unified Theory of Acceptance and Use of Technology (UTAUT): Challenging its Validity and Charting a Research Agenda in the Red Ocean. *Journal of the Association for Information Systems*, *23*(1), 13–95. https://doi.org/10.17705/1jais.00719

Captari, L. E., Riggs, S. A., & Stephen, K. (2021). Attachment processes following traumatic loss: A mediation model examining identity distress, shattered assumptions, prolonged grief, and posttraumatic growth. *Psychological Trauma: Theory, Research, Practice, and Policy*, *13*(1), 94–103. https://doi.org/10.1037/tra0000555

Ciani Sciolla, J., & Pagallo, U. (2025). No Peace After Death? The Impact of AI-Driven Memorial Chatbots on Privacy and Data Protection. *Information*, *16*(6), 426. Scopus. https://doi.org/10.3390/info16060426

Cummings, J. J., & Bailenson, J. N. (2016). How Immersive Is Enough? A Meta-Analysis of the Effect of Immersive Technology on User Presence. *Media Psychology*, *19*(2), 272–309. https://doi.org/10.1080/15213269.2015.1015740

Dinneen, J. D., Krtalić, M., Davoudi, N., Hellmich, H., Ochsner, C., & Bressel, P. (2024). Information science and the inevitable: A literature review at the intersection of death and information management: An Annual Review of Information Science and Technology (ARIST) paper. *Journal of the Association for Information Science and Technology*, *75*(3), 268–297. https://doi.org/10.1002/asi.24861

Doyle, D. T., & Brubaker, J. R. (2023). Digital Legacy: A Systematic Literature Review. *Proceedings of the ACM on Human-Computer Interaction*, *7*(CSCW2), 1–26. https://doi.org/10.1145/3610059

Galvão, V. F., Maciel, C., Pereira, R., Gasparini, I., Viterbo, J., & Bicharra Garcia, A. C. (2021). Discussing human values in digital immortality: Towards a value-oriented perspective. *Journal of the Brazilian Computer Society*, *27*(1), 15. https://doi.org/10.1186/s13173-021-00121-x

Gambino, A., Fox, J., & Ratan, R. A. (2020). Building a stronger CASA: Extending the computers are social actors paradigm. *Human-Machine Communication*, *1*, 71–85.

Harbinja, E., Edwards, L., & McVey, M. (2023). Governing ghostbots. *COMPUTER LAW & SECURITY REVIEW*, *48*. (WOS:001003753700001). https://doi.org/10.1016/j.clsr.2023.105791

Hollanek, T., & Nowaczyk-Basińska, K. (2024). Griefbots, Deadbots, Postmortem Avatars: On Responsible Applications of Generative AI in the Digital Afterlife Industry. *Philosophy and Technology*, *37*(2), 63. Scopus. https://doi.org/10.1007/s13347-024-00744-w

Joel Krueger & Lucy Osler. (2022). Communing with the Dead Online: Chatbots, Grief, and Continuing Bonds. *Journal of Consciousness Studies*. https://doi.org/10.53765/20512201.29.9.222

Karpus, J., & Strasser, A. (2025). Persons and their Digital Replicas. *Philosophy and Technology*, *38*(1). https://doi.org/10.1007/s13347-025-00854-z

Kawashima, D., Kempe, T., & Koga, Y. (2025). "I Want My Loved One Back Virtually": Exploring the Desire of Bereaved People to Create and Maintain Digital Bonds With Their Deceased Loved Ones. *OMEGA - Journal of Death and Dying*, *92*(2), 1049–1065. https://doi.org/10.1177/00302228231194857

Kreijns, K., Xu, K., & Weidlich, J. (2022). Social Presence: Conceptualization and Measurement. *Educational Psychology Review*, *34*(1), 139–170. https://doi.org/10.1007/s10648-021-09623-8

Lindemann, N. (2022). The Ethics of "Deathbots." *SCIENCE AND ENGINEERING ETHICS*, *28*(6). (WOS:000886965000001). https://doi.org/10.1007/s11948-022-00417-x

Lombard, M., & Xu, K. (2021). Social Responses to Media Technologies in the 21st Century: The Media are Social Actors Paradigm. *Human-Machine Communication*, *2*(1). https://doi.org/10.30658/hmc.2.2

Meitzler, M., Heesen, J., Hennig, M., & Quinn, R. (2024). Digital Afterlife and the Future of Collective Memory. *Memory Studies Review*. https://doi.org/10.1163/29498902-202400013

Mikulincer, M., & Shaver, P. R. (2022). An attachment perspective on loss and grief. *Current Opinion in Psychology*, *45*, 101283. https://doi.org/10.1016/j.copsyc.2021.11.003

Öhman, C., & Floridi, L. (2017). The Political Economy of Death in the Age of Information: A Critical Approach to the Digital Afterlife Industry. *Minds and Machines*, *27*(4), 639–662. https://doi.org/10.1007/s11023-017-9445-2

Philippi, P., Baumeister, H., Apolinário-Hagen, J., Ebert, D. D., Hennemann, S., Kott, L., Lin, J., Messner, E.-M., & Terhorst, Y. (2021). Acceptance towards digital health interventions – Model validation and further development of the Unified Theory of Acceptance and Use of Technology. *Internet Interventions*, *26*, 100459. https://doi.org/10.1016/j.invent.2021.100459

She, W.-J., Siriaraya, P., Ang, C. S., & Prigerson, H. G. (2021). Living Memory Home: Understanding Continuing Bond in the Digital Age through Backstage Grieving. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14. https://doi.org/10.1145/3411764.3445336

Silverman, G. S., Baroiller, A., & Hemer, S. R. (2021). Culture and grief: Ethnographic perspectives on ritual, relationships and remembering. *Death Studies*, *45*(1), 1–8. https://doi.org/10.1080/07481187.2020.1851885

Wu, W., Heierli, J., Meisterhans, M., Moser, A., Färber, A., Dolata, M., Gavagnin, E., de Spindler, A., & Schwabe, G. (2024). Promise: Model-driven Stateful Prompt Orchestration for Persuasive Conversational Interactions. *Proceedings of the 9th Edition of the Swiss Text Analytics Conference*, 185–185. https://aclanthology.org/2024.swisstext-1.37.pdf

Zuelke, A. E., Luppa, M., Löbner, M., Pabst, A., Schlapke, C., Stein, J., & Riedel-Heller, S. G. (2021). Effectiveness and Feasibility of Internet-Based Interventions for Grief After Bereavement: Systematic Review and Meta-analysis. *JMIR Mental Health*, *8*(12), e29661. https://doi.org/10.2196/29661