Spectral Labor and Cyberbiosecurity: AI Simulation of Deceased Individuals as Emerging Bio-digital Security Risks

Dominique Dove[1], Lucas Potter[2] and Xavier-Lewis Palmer[2]
[1]The George Washington Law School, Washington, D.C., USA
[2]BiosView Labs, Dayton, Ohio, USA

dominique.dove.esq@gmail.com (Corresponding author)
biosview1@proton.me (Corresponding author)

Abstract: Generative artificial intelligence (AI) can combine text, voice, images, behavioral records, biometrics and health-related data to create interactive simulations of deceased individuals. Described as griefbots, deadbots, postmortem avatars, generative ghosts, AI resurrections or forms of "spectral labor", these systems move digital legacy beyond static memorialization by generating novel speech and conduct in the apparent identity of a person who can no longer authenticate, contest or withdraw that representation. This paper examines when such systems create cyberbiosecurity risks rather than broader cybersecurity, privacy, legal, psychological or evidentiary risks. It adopts a consequence-first boundary: posthumous AI is cyberbiosecurity-relevant only where cyber disruption or misuse has a credible, proximate and system-relevant causal pathway to an observable and attributable biological consequence. Mere possession, exposure or misuse of biometric, medical or genomic information is insufficient. The paper uses conceptual and normative analysis informed by a targeted interdisciplinary review, integrative synthesis, structured threat modelling and comparative legal analysis. It distinguishes biometric, behavioral, physiological, medical and genomic data; separates harms to posthumous identity from relational harms to living persons; and models four threat pathways, including compromised health-related replicas, biometric pivots into biomedical systems, familial genomic misuse and an adjacent non-cyberbiosecurity case. Building on recurring gaps in posthumous consent, identity protection, assurance, commercial governance, stewardship, transparency and remedies, the paper proposes a seven-pillar Postmortem AI Replica Governance Framework (PARGF): Postmortem Digital Identity recognition, tiered consent, Cyberbiosecurity Impact Assessment, licensing and oversight, digital fiduciary duties, transparency and provenance, and coordinated enforcement. PARGF is a normative conceptual framework requiring empirical validation; it is not claimed to be necessary, sufficient or superior to all existing governance models. The paper argues that posthumous AI should therefore not be categorically classified as cyberbiosecurity. Rather, cyberbiosecurity relevance emerges when system compromise crosses a demonstrable biological-consequence boundary.

Keywords: Cyberbiosecurity, Posthumous AI, Digital legacy, Biometric reconstruction, Spectral labor, Data governance

1. Introduction

Generative AI is changing the character of digital legacy. Earlier memorial practices generally preserved photographs, correspondence, social-media profiles, recordings and archives as static or selectively curated traces. Contemporary large language models, multimodal generators and high-fidelity speech systems can instead combine such traces into interactive personas capable of imitating vocabulary, vocal characteristics, appearance and conversational tendencies while generating statements the represented person never produced during life (Ji et al, 2024; Morris and Brubaker, 2025). Research describes these systems as griefbots, deadbots, postmortem avatars, thanabots or generative ghosts. Divon and Pentzold (2025) describe the extraction of informational, affective and commercial value from reconstructed identities as "spectral labor". Scholarship identifies legitimate memorial and educational uses alongside concerns involving consent, representation, exploitation and user vulnerability (Öhman and Floridi, 2018; Henrickson, 2023; Hollanek and Nowaczyk-Basińska, 2024).

The present paper asks a narrower security question: when does compromise or misuse of a posthumous AI system become a cyberbiosecurity problem? Murch et al (2018) introduced cyberbiosecurity as an interdisciplinary enterprise encompassing cybersecurity, cyber-physical security and biosecurity as applied to biological and biomedical systems. Richardson et al (2019) further characterized vulnerabilities at the interface between life sciences and information systems. Potter and Palmer (2023) distinguish cyberbiosecurity, which secures biological resources and systems through cyber-enabled means, from biocybersecurity, which concerns cyber-enabled capabilities used to create or advance biological threats. Pope (2026) provides a consequence-first refinement: cyberbiosecurity is defined by systems whose cyber disruption can produce measurable biological consequence.

This consequence-first boundary prevents conceptual overreach. A posthumous simulation does not become a cyberbiosecurity system merely because it reproduces a voice or face, contains a health record, or incorporates genomic information. Financial fraud using a synthetic voice, unauthorized memorialization, reputational distortion, false historical statements and many forms of psychological harm remain principally

matters of cybersecurity, privacy, law, ethics or evidentiary integrity. For this analysis, posthumous AI is cyberbiosecurity-relevant only where cyber disruption or misuse has a credible, proximate and system-relevant causal pathway to an observable and attributable biological consequence. Exposure of biological information without downstream biological use or impact does not satisfy that condition (Pope, 2026).

The distinction matters because posthumous AI can aggregate data historically governed separately. A replica may combine correspondence, behavioral records, biometrics, medical history, physiological measurements and genomic information. The represented person can no longer authenticate new statements, revoke access or contest reuse, while genetic data can remain informative about living relatives (Clayton et al, 2019; Costello, 2022).

This paper makes three contributions. First, it establishes a bounded consequence-first criterion separating cyberbiosecurity from adjacent posthumous-AI risks. Second, it operationalizes that boundary through structured threat pathways specifying actor, vector, compromised asset, affected stakeholder or system and biological-consequence gate. Third, it derives a seven-pillar Postmortem AI Replica Governance Framework (PARGF) from recurring gaps across digital-afterlife governance, cyberbiosecurity, biometric security, genomic privacy and AI accountability.

1.1 Method and Scope

This study uses conceptual and normative design informed by a targeted interdisciplinary review and integrative synthesis. It does not claim to constitute a systematic review and does not estimate the incidence or prevalence of harms associated with posthumous AI. The review was purposive rather than exhaustive. Sources were selected for direct relevance to six evidentiary domains: digital-afterlife technologies and ethics; cyberbiosecurity boundary formation; biometric synthesis and spoofing; genomic and biomedical privacy; synthetic-media authenticity and provenance; and posthumous data and identity law. Priority was given to peer-reviewed journal articles and major peer-reviewed conference proceedings, supplemented where necessary by authoritative legal instruments. This follows Snyder's (2019) integrative-review logic for synthesizing fragmented interdisciplinary literature while requiring transparency about non-exhaustive coverage.

The literature was grouped into recurring problem classes, overlapping concepts were consolidated, each class was mapped to a governance function, and functions were assessed for redundancy before the seven PARGF pillars were formulated. The procedure makes derivation auditable without establishing unique necessity or sufficiency. Threat pathways use structured threat-modelling logic. Threat modelling is a prospective security-analysis method used to identify assets, adversarial actions, vulnerabilities and possible consequences before or during system development (Xiong and Lagerström, 2019). Here, documented component capabilities and vulnerabilities are combined into analytical scenarios to test the proposed cyberbiosecurity boundary. Unless otherwise stated, the scenarios are not reports that the complete posthumous-AI attack chains have occurred in operational systems.

2. Cyberbiosecurity Boundary and Threat Model

2.1 Consequence-first Classification

Pope (2026) formalizes cyberbiosecurity as the protection of systems whose cyber disruption can produce measurable biological consequence through three conditions: observability, causal linkage and biological relevance. The pathway must also be credible, proximate and system-relevant; highly indirect or speculative chains are excluded. Applied to posthumous AI, the distinction produces three analytical states. An adjacent posthumous-AI risk terminates privacy loss, fraud, psychological harm, reputational injury, evidentiary distortion or another non-biological consequence. A cyberbiosecurity-relevant system or threat pathway has a credible, proximate route by which cyber disruption can cross the biological-consequence gate. A cyberbiosecurity event occurs when disruption produces an observable and attributable biological consequence, or when that consequence is credibly imminent within the defined system context. This three-part terminology is an application of Pope's boundary to posthumous AI rather than terminology proposed in Pope (2026).

This distinction also clarifies the treatment of genomic data. Exposure, theft or inference of genomic information may be serious privacy and security harms without constituting cyberbiosecurity. The boundary is crossed only when compromised genomic information enters downstream biological, clinical or biomedical use in a way that produces or credibly threatens measurable biological consequence (Pope, 2026).

2.2 Data Classes and Affected Subjects

Technical precision requires separating the data classes used to construct or operate a replica. Biometric data concerns physical or behavioral characteristics used for recognition or verification, including facial and vocal characteristics (Jain, Ross and Pankanti, 2006). Behavioral data include linguistic patterns, platform activity, purchases, social relationships, location histories and interaction patterns. Physiological data arise from bodily functions or signals, including heart rate and wearable-sensor outputs. Medical or clinical data concern diagnoses, treatment, medication, imaging, laboratory results and healthcare encounters. Genetic or genomic data describes inherited or acquired sequence information and can reveal information about biological relatives as well as the source person (Clayton et al, 2019; Costello, 2022). These categories may overlap in use but should not be treated as synonymous.

Affected subjects also differ. The deceased data subject may face identity appropriation or misrepresentation; living relatives may face relational genomic disclosure or downstream health effects; service interactants may be influenced through apparent relational authority; and institutions may face compromised authentication, provenance or information integrity. The taxonomy adapts Hollanek and Nowaczyk-Basińska's (2024) distinction among data donors, data recipients and service interactants.

2.3 Structured Threat Pathways

Table 1: Posthumous-AI threat pathways and cyberbiosecurity boundary

Scenario	Threat actor and vector	Compromised asset	Affected stakeholder/system	Endpoint and boundary determination
A. Compromised health-related replica	Malicious operator or external attacker poisons retrieval data, system instructions or model behaviour	Replica containing family health history or represented as a trusted caregiver/clinician	Living relative or user	False medication, allergy or hereditary-risk information is acted upon. If it alters treatment or produces observable health effects, the pathway is cyberbiosecurity; if it ends in deception without biological effect, it remains adjacent AI/social-engineering risk.
B. Biometric pivot into biomedical infrastructure	Attacker reconstructs a deceased person's voice or face and exploits weak remote identity verification or inadequately deprovisioned credentials	Synthetic biometric representation and associated identity material	Healthcare account, clinical service or biomedical data custodian	Successful impersonation alone is cybersecurity. If access enables alteration of medically relevant data or orders and that alteration causes or credibly imminently threatens a biological consequence, the event crosses into cyberbiosecurity.
C. Familial genomic misuse	Attacker exfiltrates, manipulates or links genomic material associated with the deceased	Raw or processed genetic data	Living biological relatives, genetic service or clinical workflow	Reidentification, profiling or discrimination alone remains genetic privacy/security. Cyberbiosecurity requires downstream biological or clinical use of compromised data capable of producing measurable biological consequences.

Scenario	Threat actor and vector	Compromised asset	Affected stakeholder/system	Endpoint and boundary determination
D. Posthumous synthetic-media fraud	Attacker generates voice, video, or text attributed to the deceased	Public persona or archived media	Relative, institution, electorate or commercial counterparty	Financial, reputational, historical or evidentiary harm occurs without biological consequence. This is not cyberbiosecurity.

Note: These are analytical threat pathways constructed from documented component capabilities, not reports that each complete posthumous-AI attack chain has occurred. Relevant components include deepfake impersonation against evaluated commercial face-recognition services (Tariq, Jeon and Woo, 2022), limitations in speech-deepfake detection generalization (Liu et al, 2023), genomic reidentification and privacy attacks (Gymrek et al, 2013; Erlich and Narayanan, 2014; Thomas et al, 2024), and relational implications of genetic information for living relatives (Clayton et al, 2019; Costello, 2022).

Genomic risk has a similarly strong but distinct evidence base. Gymrek et al (2013) experimentally demonstrated surname-based reidentification of ostensibly anonymous personal genomes under specific conditions. Erlich and Narayanan (2014) organize genetic privacy breaches into identity tracing, attribute disclosure and completion attacks. Thomas et al (2024) identify a wider set of demonstrated genetic privacy attacks and emphasize that vulnerability varies by dataset composition and processing level. Costello (2022) further shows why genetic privacy is relational: information about one person can diminish the privacy of genetically related individuals. None of these findings alone satisfies Pope's cyberbiosecurity gate; they establish the security and privacy conditions from which a downstream biological pathway might arise.

3. Risk Domains and System-level Implications

3.1 Psychological and Relational Vulnerability

Interactive simulation differs from static memorialization because a generative replica can produce contingent responses and sustain an apparently continuing relationship. Commercial design introduces additional relational and security concerns. Figueroa-Torres (2024) describes an “affection as a service” economy in which mourning and continuing interaction become commercial relationships. Divon and Pentzold (2025) similarly identify forms of spectral labor through which reconstructed identities generate affective, symbolic and economic value. A compromised or manipulative replica could exploit inherited relational trust to solicit payments, elicit intimate family information, influence behavior or sustain engagement. Such effects remain principally consumer, privacy, ethical or cybersecurity concerns unless they satisfy the biological-consequence boundary.

3.2 Biometric, Medical and Genomic Exposure

Posthumous simulation can concentrate otherwise separate identity data. Text supports communication-style reconstruction; speech and imagery support voice and facial synthesis; and physiological, clinical or genomic records introduce sensitive biological information. The security issue is linkage through a persistent identity proxy, not that all such data are biomedical.

Biometric characteristics differ from user-selected secrets such as passwords. Facial and vocal traits are persistent physical or behavioral attributes, even though biometric templates can be protected, transformed or revoked at the system level. Their use in posthumous reconstruction therefore strengthens the case for liveness testing, contextual authentication, protected templates and additional authentication factors rather than reliance on resemblance alone (Jain, Ross and Pankanti, 2006; Tariq, Jeon and Woo, 2022).

Genomic data requires a separate analysis because control is relational rather than purely individual. Genetic variants and associated information are partly shared among biological relatives, making individualistic privacy models incomplete (Clayton et al, 2019; Costello, 2022). Posthumous processing can therefore expose living relatives who never authorized the simulation. Nevertheless, the classification must remain precise: profiling, reidentification and unauthorized secondary use are significant privacy harms, but under Pope (2026) they become cyberbiosecurity only where compromised data enter downstream biological use or impact.

4. Governance and Research Agenda

4.1 Governance Gap and Framework Derivation

Posthumous-AI governance draws on several bodies of scholarship that address different parts of the problem. Öhman and Floridi (2018) provide an early ethical framework for the digital-afterlife industry. Harbinja, Edwards and McVey (2023) analyze the legal governance of ghostbots and ante-mortem control. Hollanek and Nowaczyk-Basińska (2024) emphasize mutual consent, transparency and responsible retirement. Birnhack and Morse (2022) and Morse and Birnhack (2024) show why digital remains implicate continuing informational and impression-management interests that are not captured fully by simple property rules. Cyberbiosecurity scholarship, meanwhile, addresses security of biological systems, mission-specific controls and biological consequences rather than posthumous identity specifically (Murch et al, 2018; Potter and Palmer, 2023; Pope, 2026).

The GDPR illustrates the fragmentation. Article 17 establishes a right to erasure within the Regulation's scope, but Recital 27 states that the GDPR does not apply to personal data of deceased persons while allowing Member States to create separate rules (European Union, 2016). The represented person may therefore lose ordinary data-subject mechanisms at precisely the point when generative systems acquire the ability to produce new identity-attributed content.

PARGF connects rather than replaces these regimes by mapping recurring problem classes to distinct governance functions.

Table 2: Derivation of the Postmortem AI Replica Governance Framework

Recurring problem class	Governance function	PARGF pillar
Uncertain legal status of reconstructed identity and digital remains	Distinguish possession of source material from authority to generate an interactive identity	1. Postmortem Digital Identity
Absent, ambiguous or overbroad posthumous authorization	Record granular ante-mortem permissions and limits	2. Tiered Consent
Difficulty separating ordinary AI risk from biological consequence	Apply a formal pre-deployment consequence and threat assessment	3. Cyberbiosecurity Impact Assessment
Commercial scale, autonomy and sensitive-data processing	Apply heightened entry and assurance requirements to higher-risk deployments	4. Licensing and Oversight
Fragmented stewardship after the represented person can no longer intervene	Allocate continuing duties to operators and designated representatives	5. Digital Fiduciary Duties
Difficulty distinguishing archival, authorized, generated and altered content	Require visible disclosure and technical provenance	6. Transparency and Provenance
Consent without effective remedies; breach and cross-platform propagation	Provide investigation, suspension, correction, decommissioning and liability mechanisms	7. Enforcement and Remedies

The mapping explains the seven functions without claiming that seven is uniquely correct; the architecture remains open to empirical revision.

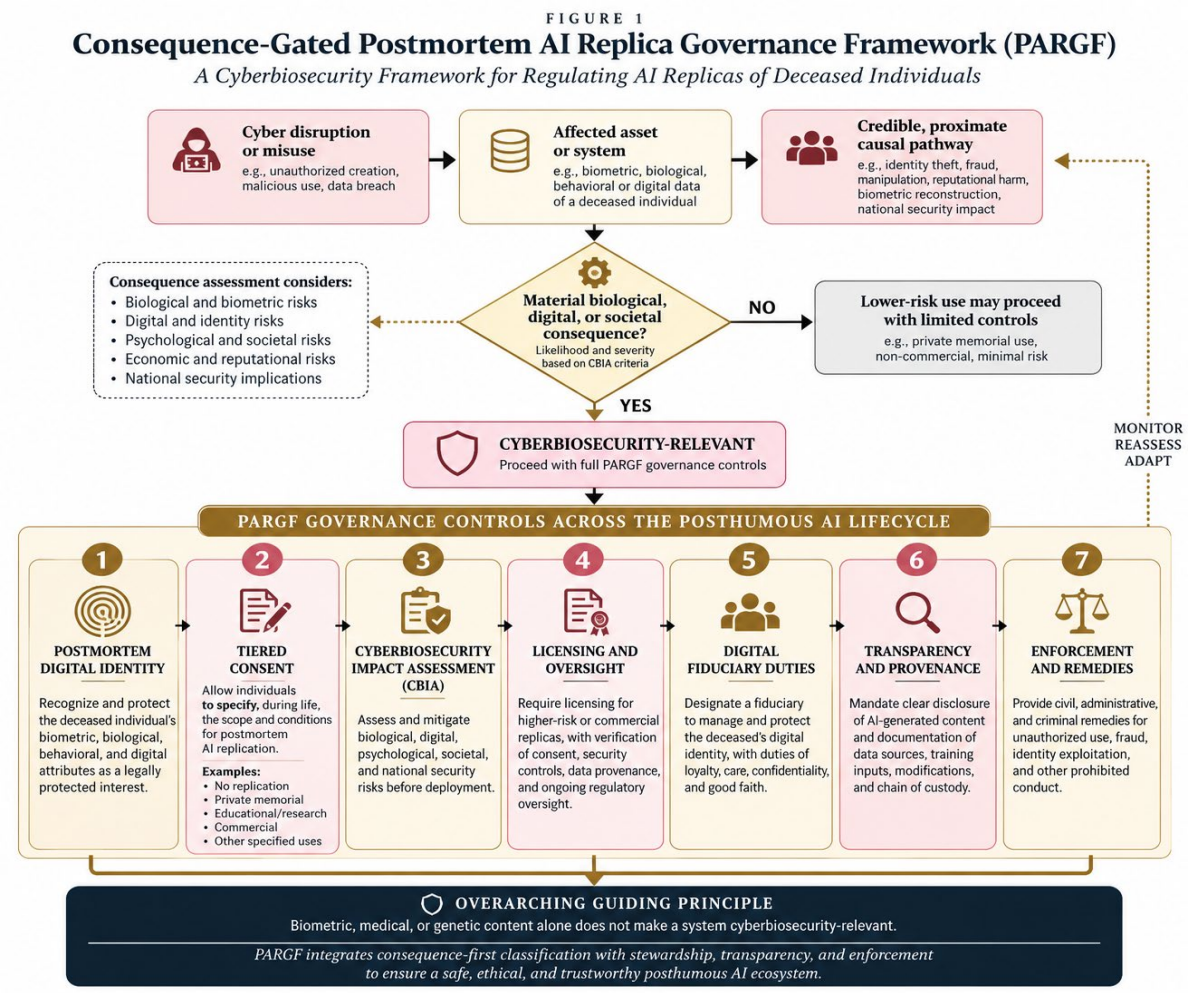

Figure 1: Consequence-Gated Postmortem AI Replica Governance Framework (PARGF). Cyber disruption is classified through a measurable biological-consequence gate; PARGF controls span identity, consent, assessment, oversight, stewardship, transparency and enforcement, with incident response and decommissioning feeding back into control revision

4.2 Seven-pillar Regulatory Framework

4.2.1 Postmortem digital identity

Postmortem Digital Identity (PDI) should be recognized as a protected governance interest comprising the attributes through which an identifiable deceased person can be simulated. PDI separates possession of photographs, messages, recordings or records from authority to combine them into an autonomous or interactive representation. This responds to the difficulty of treating digital remains solely as property and to the continuing representational interests identified in digital-remains scholarship (Birnhack and Morse, 2022; Morse and Birnhack, 2024).

4.2.2 Tiered consent

Ante-mortem consent should be granular rather than binary. Authorization should identify permissible purposes, audiences, data classes, duration, commercial use, capacity to generate novel first-person statements, system autonomy and procedures for suspension or retirement. A tiered approach could distinguish no simulation; private memorial use; educational or historical use; research; regulated commercial licensing; and broader uses under explicit documentation. The objective is not to assume that relatives inherit unlimited authority over identity, but to make the represented person's documented preferences the starting point while recognizing the interests of other data subjects and service interactants (Harbinja, Edwards and McVey, 2023; Hollanek and Nowaczyk-Basińska, 2024).

4.2.3 Cyberbiosecurity impact assessment

A Cyberbiosecurity Impact Assessment (CBIA) should precede deployment when a replica processes sensitive biological information, connects to biomedical infrastructure, functions as an identity or authentication interface, or can materially influence health-related decisions. The first stage should apply Pope's inclusion gate: identify the cyber system and disruption, specify the causal pathway, and determine whether the endpoint is observable, causally attributable and biologically relevant. Pathways that terminate in privacy, fraud, reputational or psychological harm without downstream biological effect remain outside cyberbiosecurity.

The second stage should evaluate threat actors, exposed assets, biometric spoofing, liveness controls, genomic disclosure, model and retrieval integrity, access control, provenance, unauthorized fine-tuning, cross-platform propagation, monitoring and incident response. Controls can then be organized around prevention, detection and elimination or recovery according to the protected mission (Potter and Palmer, 2023).

4.2.4 Licensing and oversight

Not every private memorial tool warrants licensing. Heightened oversight is more defensible where a commercial system represents an identifiable deceased person, processes medical or genomic information, operates autonomously, connects to external services or is marketed as authoritative. Such systems should demonstrate lawful authority, provenance, security testing, auditability, incident response and decommissioning. This avoids imposing equivalent burdens on low-risk memorialization (Figueroa-Torres, 2024; Hollanek and Nowaczyk-Basińska, 2024).

4.2.5 Digital fiduciary duties

Entities controlling posthumous replicas should owe continuing duties appropriate to their role, including care, confidentiality, security, loyalty to authorized purposes, transparency and prompt response to misuse. A designated representative should be able to enforce valid ante-mortem instructions, request relevant audit information, contest materially unauthorized representations and seek suspension or retirement where warranted. Fiduciary stewardship is particularly important where the source person can no longer exercise ordinary data-subject controls, but processing continues to affect correspondents or relatives.

Laws should impose duties of care, confidentiality, loyalty, transparency and accountability on entities determining the purposes and means of personal-data processing. This framework could draw on and extend India’s Digital Personal Data Protection Act 2023. Data Fiduciaries should assume duties reflecting the vulnerability and power asymmetries of data-processing relationships (Richards and Hartzog, 2021). They should monitor compliance and risk, report breaches or material unlawful processing, and implement lawful requests to withdraw consent, correct or erase data, or modify authorized processing, and watermark their products (Raji et al., 2020; Kaminski and Malgieri, 2021; Government of India, 2023). They should also provide concise, accessible explanations of data collection, processing purposes and available rights, as complex privacy notices undermine comprehension and informed decision-making (Milne, Culnan and Greene, 2006; Schaub et al., 2015).

Individuals should also be able to issue binding instructions for the posthumous use of their data, identity and likeness. Additionally, they should be able to appoint a representative or have the state appoint a representative to enforce their instructions. The representative should have duties of loyalty, security, confidentiality, transparency and prompt response to misuse, while managing relevant licenses. They should have standing to request audits, restrict access, require correction or disclosure of material misrepresentations, and seek suspension or dignified decommissioning of AI-generated simulations. These powers address the gap between individuals’ preferences for digital remains and their ability to enforce them (Morse and Birnhack, 2022; Morse and Birnhack, 2024), while responding to concerns involving consent, identity consistency, commercial exploitation and simulation retirement (Öhman and Floridi, 2018; Hollanek and Nowaczyk-Basińska, 2024; Lei et al., 2025).

4.2.6 Transparency and provenance

Users should be informed conspicuously that they are interacting with an AI-generated simulation rather than the deceased person. Disclosure should identify the responsible operator, authorized purpose, principal data categories, sponsorship where relevant, material limitations and significant model changes.

4.2.7 Enforcement and remedies

Posthumous consent is ineffective without remedies after deployment. Depending on jurisdiction and conduct, mechanisms may include injunctions, deletion where technically feasible, breach notification, administrative penalties and ordinary criminal liability. Enforcement should remain capability-, conduct- and consequence-sensitive rather than treating a private memorial chatbot and an autonomous commercial replica processing genomic information as equivalent.

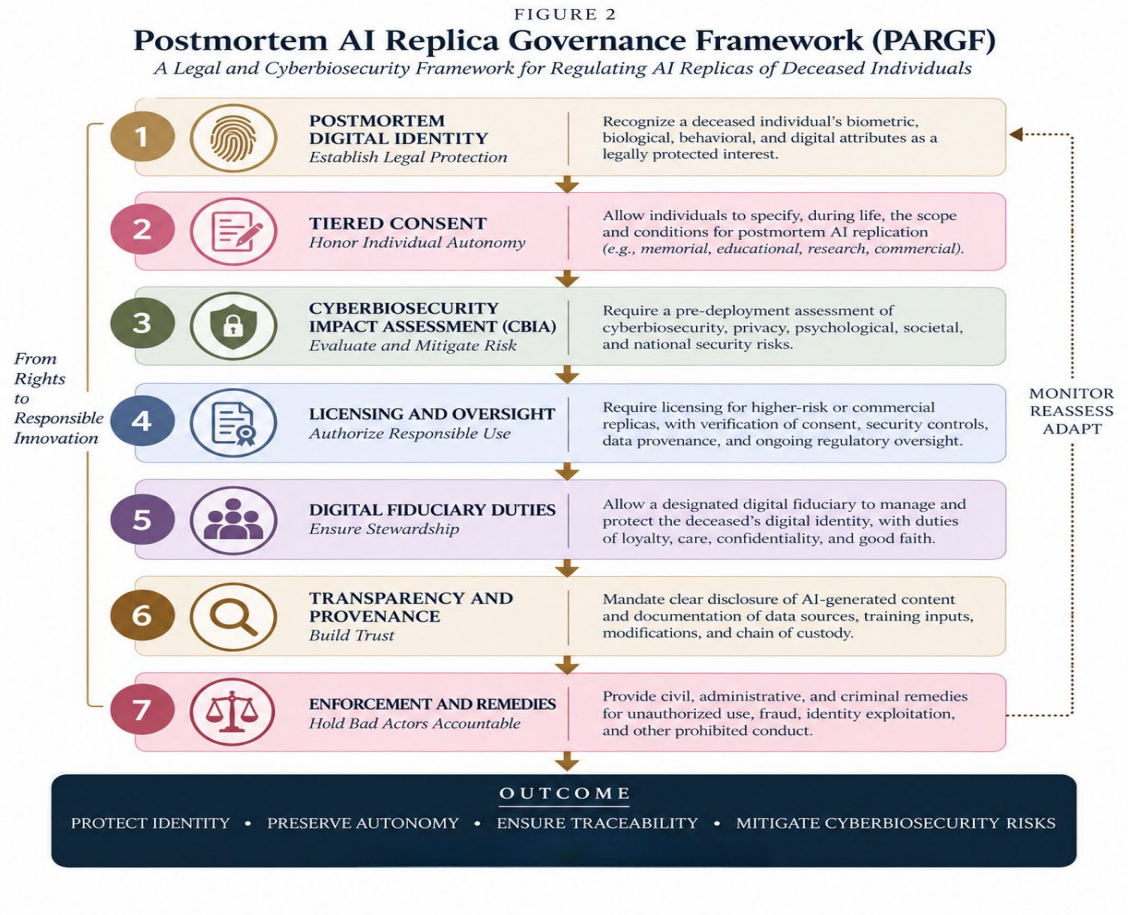

Figure 2: Operational Structure of the Postmortem AI Replica Governance Framework (PARGF). The original seven-pillar visualization summarizes the lifecycle controls described in Section 4.2. In the revised terminology, "Legal Recognition" operationalizes Postmortem Digital Identity recognition, and "Cyberbiosecurity Risk Assessment" corresponds to the Cyberbiosecurity Impact Assessment (CBIA) formalized in Section 4.2.3

4.3 Limitations and Research Priorities

This study has limitations consistent with its conceptual and normative design. First, it uses a targeted interdisciplinary review and integrative synthesis rather than a systematic review. It therefore does not claim exhaustive coverage or estimate the prevalence of harms associated with posthumous AI. Second, the threat pathways in Table 1 are analytical compositions of documented component capabilities rather than evidence that each complete pathway has occurred in a deployed posthumous-AI system. Demonstrated voice synthesis, biometric impersonation vulnerabilities, genomic privacy attacks and generative-replica capabilities establish enabling conditions but do not establish attack frequency, adversary intent or realized biological harm. Third, PARGF is an author-derived conceptual framework produced through an explicit synthesis procedure. Its seven pillars have not been empirically validated as individually necessary, collectively sufficient or superior to all existing AI-governance, privacy, cybersecurity, digital-estate or sector-specific regulatory frameworks. Their value at this stage lies in making the proposed governance logic transparent and testable.

Future research should therefore test inter-rater consistency in applying the consequence-first rule to posthumous-AI scenarios, establish measurable thresholds for biological consequence, evaluate liveness and

provenance controls against replicas, examine familial impacts of genomic and health inference, and compare PARGF with existing national privacy, AI, succession, consumer-protection and cybersecurity regimes. Empirical work may justify modification, merger, expansion or rejection of individual pillars.

5. Conclusion

AI simulations of deceased individuals occupy an emerging intersection among digital legacy, generative AI, identity security, privacy and, in a narrower subset of cases, cyberbiosecurity. The defining issue is not whether a system includes a face, voice, health fact or genome. It is whether cyber disruption or misuse creates a credible, proximate and system-relevant causal pathway to an observable and attributable biological consequence (Pope, 2026).

This boundary prevents cyberbiosecurity from becoming an umbrella term for every risk associated with posthumous AI. Financial impersonation, reputational distortion, unauthorized commercial resurrection, historical fabrication, genomic privacy loss and many psychological concerns remain serious but principally occupy adjacent governance domains when no biological consequence follows. Conversely, posthumous systems may become cyberbiosecurity-relevant when compromised replicas influence health-related actions, synthetic identity material provides a pivot into biomedical systems or manipulated genomic or medical information enters downstream biological processes.

Literature establishes important enabling conditions. High-fidelity speech synthesis can operate from short reference samples; speech-deepfake detection generalizes imperfectly across conditions; tested face-recognition services have been vulnerable to deepfake impersonation; genetic information presents demonstrated reidentification and relational-privacy risks; generative ghosts can produce novel identity-attributed content; and provenance affects user judgments in ways that are not always predictable. None of these findings cause harm inevitable. Their significance is that they identify components from which bounded attack pathways can be modelled as posthumous AI becomes more autonomous and interconnected.

PARGF translates these findings into seven governance functions: Postmortem Digital Identity recognition, tiered consent, Cyberbiosecurity Impact Assessment, licensing and oversight, digital fiduciary stewardship, transparency and provenance, and coordinated enforcement. The framework is deliberately provisional. Its contribution is not to declare every AI afterlife a cyberbiosecurity system, but to provide an auditable structure for deciding when a posthumous-AI risk crosses that boundary and what additional governance should follow.

Ethics Declaration: Ethical clearance was not required. The study comprises legal, technological, security and policy analysis and did not involve human or animal participants or the collection of personal research data.

AI Declaration: Generative AI tools, including ChatGPT and Gemini, were used to assist with structural editing, figure creation, condensation, citation mapping and formatting to conference requirements. The authors reviewed the resulting text and remained responsible for the analysis, source selection, citation accuracy and final submission.

References

Birnhack, M. and Morse, T. (2022) “Digital remains: Property or privacy?”, *International Journal of Law and Information Technology*, Vol 30, No. 3, pp 280–301. doi: 10.1093/ijlit/eaac019.

Clayton, E.W., Evans, B.J., Hazel, J.W. and Rothstein, M.A. (2019) “The law of genetic privacy: Applications, implications, and limitations”, *Journal of Law and the Biosciences*, Vol 6, No. 1, pp 1–36. doi: 10.1093/jlb/lsz007.

Costello, R.Á. (2022) “Genetic data and the right to privacy: Towards a relational theory of privacy?”, *Human Rights Law Review*, Vol 22, No. 1, article ngab031. doi: 10.1093/hrlr/ngab031.

Divon, T. and Pentzold, C. (2025) “Artificially alive: An exploration of AI resurrections and spectral labor modes in a postmortal society”, *New Media & Society*, advance online publication. doi: 10.1177/14614448251397518.

Erlich, Y. and Narayanan, A. (2014) “Routes for breaching and protecting genetic privacy”, *Nature Reviews Genetics*, Vol 15, pp 409–421. doi: 10.1038/nrg3723.

European Union (2016) “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016”, *Official Journal of the European Union*, L119, pp 1–88.

Figueroa-Torres, M. (2024) “Affection as a service: Ghostbots and the changing nature of mourning”, *Computer Law & Security Review*, Vol 52, article 105943. doi: 10.1016/j.clsr.2024.105943.

Gymrek, M., McGuire, A.L., Golan, D., Halperin, E. and Erlich, Y. (2013) “Identifying personal genomes by surname inference”, *Science*, Vol 339, No. 6117, pp 321–324. doi: 10.1126/science.1229566.

Harbinja, E., Edwards, L. and McVey, M. (2023) “Governing ghostbots”, *Computer Law & Security Review*, Vol 48, article 105791. doi: 10.1016/j.clsr.2023.105791.

Henrickson, L. (2023) "Chatting with the dead: The hermeneutics of thanabots", *Media, Culture & Society*, Vol 45, No. 5, pp 949–966. doi: 10.1177/01634437221147626.

Hollanek, T. and Nowaczyk-Basińska, K. (2024) "Griefbots, deadbots, postmortem avatars: On responsible applications of generative AI in the digital afterlife industry", *Philosophy & Technology*, Vol 37, article 63. doi: 10.1007/s13347-024-00744-w.

Jain, A.K., Ross, A. and Pankanti, S. (2006) "Biometrics: A tool for information security", *IEEE Transactions on Information Forensics and Security*, Vol 1, No. 2, pp 125–143. doi: 10.1109/TIFS.2006.873653.

Ji, S., Jiang, Z., Wang, H., Zuo, J. and Zhao, Z. (2024) "MobileSpeech: A fast and high-fidelity framework for mobile zero-shot text-to-speech", in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, Vol 1, pp 13588–13600. doi: 10.18653/v1/2024.acl-long.733.

Lei, Y., Ma, S., Sun, Y. and Ma, X. (2025) '"AI Afterlife" as digital legacy: Perceptions, expectations, and concerns', in Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems. New York: ACM, article 981. doi: 10.1145/3706598.3713933.

Kaminski, M.E. and Malgieri, G. (2021) 'Algorithmic impact assessments under the GDPR: Producing multi-layered explanations', International Data Privacy Law, 11(2), pp. 125–144. doi: 10.1093/idpl/ipaa020.

Liu, X., Wang, X., Sahidullah, M., Patino, J., Delgado, H., Kinnunen, T., Todisco, M., Yamagishi, J., Evans, N., Nautsch, A. and Lee, K.A. (2023) "ASVspoof 2021: Towards spoofed and deepfake speech detection in the wild", *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, Vol 31, pp 2507–2522. doi: 10.1109/TASLP.2023.3285283.

Milne, G.R., Culnan, M.J. and Greene, H., 2006. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing*, *25*(2), pp.238-249. doi: 10.1509/jppm.25.2.238.

Morris, M.R. and Brubaker, J.R. (2025) "Generative ghosts: Anticipating benefits and risks of AI afterlives", in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, article 536, pp 1–14. doi: 10.1145/3706598.3713758.

Morse, T. and Birnhack, M. (2024) "The continuity principle of digital remains", *New Media & Society*, Vol 26, pp 5240–5258. doi: 10.1177/14614448221133535.

Murch, R.S., So, W.K., Buchholz, W.G., Raman, S. and Peccoud, J. (2018) "Cyberbiosecurity: An emerging new discipline to help safeguard the bioeconomy", *Frontiers in Bioengineering and Biotechnology*, Vol 6, article 39. doi: 10.3389/fbioe.2018.00039.

Öhman, C. and Floridi, L. (2018) "An ethical framework for the digital afterlife industry", *Nature Human Behaviour*, Vol 2, No. 5, pp 318–320. doi: 10.1038/s41562-018-0335-2.

Pope, T. (2026) "Defining cyberbiosecurity by biological consequence", in *Proceedings of the 47th IEEE Symposium on Security and Privacy Workshops (SPW)*, pp 214–220. doi: 10.1109/SPW72489.2026.00025.

Potter, L. and Palmer, X.-L. (2023) "Mission-aware differences in cyberbiosecurity and biocybersecurity policies: Prevention, detection, and elimination", in Greenbaum, D. (Ed.), *Cyberbiosecurity: A New Field to Deal with Emerging Threats*, Springer, Cham, pp 37–69. doi: 10.1007/978-3-031-26034-6_4.

Raji, I.D., Smart, A., White, R.N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D. and Barnes, P. (2020) 'Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing', in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. New York: ACM, pp. 33–44. doi: 10.1145/3351095.3372873.

Richards, N. and Hartzog, W., 2021. A duty of loyalty for privacy law. *Wash. UL Rev.*, *99*, p.961.

Richardson, L.C., Connell, N.D., Lewis, S.M., Pauwels, E. and Murch, R.S. (2019) "Cyberbiosecurity: A call for cooperation in a new threat landscape", *Frontiers in Bioengineering and Biotechnology*, Vol 7, article 99. doi: 10.3389/fbioe.2019.00099.

Snyder, H. (2019) "Literature review as a research methodology: An overview and guidelines", *Journal of Business Research*, Vol 104, pp 333–339. doi: 10.1016/j.jbusres.2019.07.039.

Tariq, S., Jeon, S. and Woo, S.S. (2022) "Am I a real or fake celebrity? Evaluating face recognition and verification APIs under deepfake impersonation attack", in *Proceedings of the ACM Web Conference 2022*, pp 512–523. doi: 10.1145/3485447.3512212.

Thomas, M., Mackes, N., Preuss-Dodhy, A., Wieland, T. and Bundschus, M. (2024) "Assessing privacy vulnerabilities in genetic data sets: Scoping review", *JMIR Bioinformatics and Biotechnology*, Vol 5, article e54332. doi: 10.2196/54332.

Xiong, W. and Lagerström, R. (2019) "Threat modeling – A systematic literature review", *Computers & Security*, Vol 84, pp 53–69. doi: 10.1016/j.cose.2019.03.010.