Illusions of Control

Consumer Digital Replica Platforms and the Governance of Identity

Makena Binker Cosen*

* Makena Binker Cosen is an Associate at Kirkland & Ellis LLP in New York. She graduated Columbia Law School in 2025.

© 2026 Binker Cosen. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction, provided the original author and source are credited.

To me, the Symposium has become a marker in the passage of time. I began attending before I was even sure that law school was for me, and the enthusiasm in the room played a meaningful role in that decision. Each year, I returned with a clearer grasp of the conversations unfolding around me—eventually participating in discussions I once worked so hard just to follow. As I became familiar with the faces reappearing year after year, the Symposium has come to feel like a reunion of a community I am honored to be part of.

Returning after graduating law school as a speaker, standing alongside the mentors who had shaped my path, carried real personal significance. That is why, when my phone rang in the middle of my presentation on deepfakes, I should have been mortified. Instead, when I answered the call, I was confident knowing that my digital replica would step in for me. Deliberately planned, this moment illustrated a future in which the same technologies often associated with deepfakes are used to create digital replicas that make casual appearances in everyday life. Built with amateur knowledge, on a tight budget, and in a matter of days, consumer digital replica platforms are already ushering in that future.

Digital replicas are becoming ordinary. People record short samples of speech to preserve a voice, upload photographs to generate images they never posed for, or rely on systems trained on their prior writing to draft messages on their behalf. These practices are rarely framed as moments of identity delegation. They are presented instead as conveniences, accessibility tools, or creative efficiencies. Yet each depends on the same underlying move: extracting elements of a real person's identity and embedding them into a persistent digital system capable of acting without the person's ongoing participation. Once created, these systems can speak, appear, or respond in ways meaningfully attributable to an individual, even as control over their operation increasingly rests with the platform that hosts them. Existing legal frameworks have struggled to account for that shift.

This Article argues that consumer digital replicas constitute a distinct category of identity-bearing systems that cannot be adequately understood through doctrines

focused on deception, labor, or isolated misuse. The central problem is not that digital replicas are inherently harmful. Many are beneficial and empowering. The problem is governance: how identity becomes persistent, generative, and platform-mediated over time.

Part I stabilizes "digital replica" as an analytical category. It traces the term's reactive evolution across preservation, engineering, entertainment, labor, and AI governance, and distills shared structural features that cut across those contexts. From this synthesis, the Comment proposes a working definition suited to the consumer era. Part II examines how digital replicas enter everyday life through consumer products. Focusing on voice, visual, and behavioral systems, it shows how modest user inputs generate reusable identity-signifying outputs, even as these tools are framed as neutral utilities rather than identity technologies. Part III turns to governance. It analyzes the limits of revocable consent and the ongoing contest over permissible use shaped by platform design, safeguards, pricing, and portability. Across the lifecycle of a replica, control migrates incrementally from users to platforms. The Article concludes by explaining why existing legal frameworks arrive too late. Without a structural account of digital replicas as persistent identity systems, law intervenes only after control over identity has already been decided.

I. DEFINING DIGITAL REPLICAS

The term "digital replica" has become a familiar part of contemporary debates about artificial intelligence, identity, and personality rights. Yet despite its frequent use, the term does not carry a single, uniform meaning across legal, technical, and cultural contexts. It has been used to describe everything from digital museum objects to CGI "digital doubles" of actors, AI-generated voice models, and synthetic media regulated under deepfake statutes. These uses did not develop together, and they do not rest on a shared understanding of what, exactly, is being replicated—or why it matters.

For legal purposes, this definitional fragmentation is a problem. Treating digital replicas as merely deceptive deepfakes or as a niche labor issue misses a broader and increasingly common phenomenon: the creation of persistent digital systems that reproduce aspects of a real human identity and can operate independently of the person they represent. This section will develop a definition of "digital replica" by examining the term's evolution, isolating its defining features, and distinguishing it from adjacent concepts.

A. THE REACTIVE EVOLUTION OF THE TERM "DIGITAL REPLICA"

The phrase "digital replica" did not begin as a legal concept, nor did it initially refer to people at all. Its earliest uses appeared in technical and academic contexts concerned with faithfully reproducing physical objects in digital form. In cultural heritage scholarship, for example, a digital replica was defined as "a faithful copy of an original artifact in the digital domain, including its appearance, its morphology and how it is

meant to interact."[1] In that setting, replication was about accuracy and preservation. The object being replicated was static, nonhuman, and socially inert. Legal questions about consent, control, or misuse were simply not part of the picture.

A more dynamic version of the term emerged in engineering and industrial systems design, where "digital replica" became closely associated with the idea of a "digital twin." Industry sources described digital replicas as virtual representations of physical assets, systems, or processes, often connected to real-time data and used for monitoring, simulation, or optimization.[2] McKinsey, for instance, defined a digital twin as "a digital replica of a physical object, person, system, or process, contextualized in a digital version of its environment."[3] These definitions introduced persistence and reusability as core features, but only in service of operational goals. Even when people were included as possible subjects, they were treated as systems to be modeled, not as identities to be protected.

The term took on a very different meaning once advances in computer-generated imagery (CGI) and machine learning made it possible to replicate human likeness. By the late 2010s, media outlets began describing "digital replicas" of actors as highly realistic digital doubles capable of performing on screen without the actor's physical presence—including, in some cases, after the actor's death.[4] Legal scholarship soon followed. In the entertainment context, a digital replica was defined as "a computer-generated image that recreates the likeness of a person—their face, body, voice, and movement."[5] At this point, replication was no longer about copying objects or optimizing systems. It was about standing in for a human being.

Labor organizations adopted the term with similar assumptions. SAG-AFTRA defined a "digital replica" as "a replica of your voice and/or likeness that is created using digital technology, such as artificial intelligence," and drew a sharp distinction between replicas created with a performer's participation and those created independently from existing recordings.[6] These definitions were narrow by design. They focused on

1. Fabrizio Ivan Apollonio, Marco Gaiani & Simone Garagnani, *Visualization and Fruition of Cultural Heritage in the Knowledge-Intensive Society*, *in* HANDBOOK OF RESEARCH ON IMPLEMENTING DIGITAL REALITY AND INTERACTIVE TECHNOLOGIES TO ACHIEVE SOCIETY 5.0 471, 495 (Francesca Maria Ugliotti & Anna Osello eds., 2022).

2. *What Is Digital-Twin Technology?*, MCKINSEY & CO. (Aug. 26, 2024), https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-digital-twin-technology [https://web.archive.org/web/20260205210906/https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-digital-twin-technology].

3. *Id.*

4. Sean Cummings, *Q&A: Stanford Engineers Discuss Digital Doubles*, STAN. REP. (Sep. 26, 2023), https://news.stanford.edu/stories/2023/09/qa-stanford-engineers-discuss-digital-doubles [https://web.archive.org/web/20250818120434/https://news.stanford.edu/stories/2023/09/qa-stanford-engineers-discuss-digital-doubles].

5. Alexandra Curren, Note, *Digital Replicas: Harm Caused by Actors' Digital Twins and Hope Provided by the Right of Publicity*, 102 TEX. L. REV. 155, 159 (2023).

6. *Digital Replicas 101*, SAG-AFTRA 1 (2023), https://www.sagaftra.org/sites/default/files/sa_documents/DigitalReplicas.pdf [https://web.archive.org/web/20241003065947/https://www.sagaftra.org/sites/default/files/sa_documents/DigitalReplicas.pdf].

professional performers, third-party creators, and unauthorized reuse. The problem was framed as one of bargaining power and consent in employment relationships—not as a general issue of identity governance.

Legislatures then picked up the term in response to generative AI and deepfakes, narrowing it further. The European Union's Artificial Intelligence Act – hailed as the first comprehensive AI framework of its kind—did not define the concept of a "digital replica." However, it addressed deep fakes as including AI-generated media that "resembles existing persons" and "would falsely appear to a person to be authentic."[7] By contrast, U.S. state laws expressly adopt the language of "digital replicas." For example, New York defines a digital replica as "a digital simulation of the voice or likeness of an individual that so closely resembles the individual's voice or likeness that a layperson would not be able to readily distinguish the digital simulation from the individual's authentic voice or likeness"; California and Illinois likewise emphasize high realism and the likelihood of audience confusion.[8] Federal policy echoes this framing: The U.S. Copyright Office has described digital replicas as digital recordings that "realistically but falsely depict an individual."[9]

These definitions make sense given their regulatory goals. They are aimed at fraud, impersonation, and misinformation, and they therefore focus on realism and deception. However, they also reflect assumptions that are increasingly misaligned with the direction in which digital replica technologies are evolving. While such systems are not yet mainstream, they are becoming progressively more accessible, and are often designed for disclosed, consented, and self-directed use rather than deception or impersonation.

Consumer and accessibility technologies illustrate this shift. Voice banking tools and personal AI systems allow users to create persistent models trained on their own speech or expressive patterns. Apple's Personal Voice feature, for example, enables users to generate a synthetic voice derived from their own recordings for ongoing use.[10]

---

7. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Laying Down Harmonised Rules on Artificial Intelligence, art. 3(60), 2024 O.J. (L 1689) 1; Shiona McCallum, Liv McMahon & Tom Singleton, *MEPs Approve World's First Comprehensive AI Law*, BBC (Mar. 13, 2024), https://www.bbc.com/news/technology-68546450 [https://web.archive.org/web/20260226030228/https://www.bbc.com/news/technology-68546450].

8. N.Y. GEN. OBLIG. LAW § 5-302 (McKinney 2025); CAL. LAB. CODE § 927(c)(1) (West 2025); 765 ILL. COMP. STAT. 1075/5 (West 2025); *see also* Miguel A. Lopez, Brad Kelley & Shreya Mantrala, *New York's Digital Replica Law and Its Impact on Artificial Intelligence and the Entertainment Industry*, LITTLER (Jan. 21, 2025), https://www.littler.com/publication-press/publication/new-yorks-digital-replica-law-and-its-impact-artificial-intelligence [https://web.archive.org/web/20260213181856/https://www.littler.com/news-analysis/asap/new-yorks-digital-replica-law-and-its-impact-artificial-intelligence-and].

9. U.S. COPYRIGHT OFF., COPYRIGHT AND ARTIFICIAL INTELLIGENCE, PART I: DIGITAL REPLICAS 2 (2024), https://copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-1-Digital-Replicas-Report.pdf [https://web.archive.org/web/20260205212401/https://copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-1-Digital-Replicas-Report.pdf].

10. Apple, *Apple Previews Live Speech, Personal Voice, and More New Accessibility Features* (May 16, 2023), https://www.apple.com/ke/newsroom/2023/05/apple-previews-live-speech-personal-voice-and-more-new-accessibility-features

These systems are rarely called "digital replicas," but they share the same underlying features found in other contexts: persistence, reusability, and the ability to generate identity-signifying outputs without live human performance.[11]

Looking across these domains, a pattern becomes clear. "Digital replica" has been defined reactively, with each definition tailored to a specific institutional concern—preservation, optimization, labor protection, or deception. None of these definitions are wrong, but each is incomplete. Together, they obscure a broader trend that cuts across contexts: the externalization of human identity into persistent digital systems that can be stored, modified, licensed, and reused over time.

B. CORE CHARACTERISTICS OF DIGITAL REPLICAS

A more general definition of digital replicas must therefore move away from sector-specific use cases and toward shared structural features. What unites the disparate technologies described above is not how they are deployed, but how they function. These features also clarify why certain adjacent technologies—deepfakes, avatars, filters, and generic or fictional AI systems—fall outside the category, even where they overlap in form or technical sophistication.

First, a digital replica is defined by referential fidelity to a real person. It is tied to an identifiable individual and derives its meaning from that connection, whether the fidelity is visual, vocal, behavioral, or expressive. Perfect realism is not required. What matters is that the system purports to stand in for a particular person, not the medium used or whether the output is deceptive. By contrast, generic AI outputs, fictional characters, and virtual influencers lack a real-world referent and therefore fall outside the category.

Second, digital replicas exhibit technical persistence and reusability. Unlike a single photograph, recording, or live performance, a digital replica is designed to endure: It can be stored, copied, updated, and redeployed across contexts, often without the person's ongoing involvement.[12] This distinguishes replicas from one-off deepfakes

[https://web.archive.org/web/20260205212356/https://www.apple.com/ke/newsroom/2023/05/apple-previews-live-speech-personal-voice-and-more-new-accessibility-features/].

11. Consumer discourse reflects an understanding of these features, with one technology commentator seeking to re-deploy Apple's Personal Voice to screen calls from telemarketers, generate narrations for videos and presentations, and enhance live speech in professional settings on sick days. *See* Brad Morton, *Five Ways I Wish Apple Would Let Me Use My AI Cloned Voice*, HOW-TO GEEK (June 3, 2024), https://www.howtogeek.com/5-ways-i-wish-apple-would-let-me-use-my-ai-cloned-voice/ [https://web.archive.org/web/20260326155141/https://www.howtogeek.com/5-ways-i-wish-apple-would-let-me-use-my-ai-cloned-voice/].

12. *See, e.g.*, *Auto-Updating Knowledge Base*, PICKAXE, https://pickaxe.co/ai-knowledge-base [https://web.archive.org/web/20260326155455/https://pickaxe.co/ai-knowledge-base] (last visited Mar. 26, 2026) (automatically updates AI agent's sources on a daily basis "without any manual work"); *Use Cases: Unlock New Ways to Create with HeyGen*, HEYGEN, https://www.heygen.com/use-cases [https://web.archive.org/web/20260326155641/https://www.heygen.com/use-cases] (last visited Mar. 26, 2026) (showing AI avatar can be used for social media, advertising, training, presentations, and courses); Sydney Bradley, *Death Isn't the End: Meta Patented an AI That Lets You Keep Posting from Beyond the Grave,* BUS. INSIDER (Feb. 11, 2026), https://www.businessinsider.com/meta-granted-patent-for-ai-llm-bot-dead-

and real-time filters, which remain tied to a single use or moment. Persistence matters legally because it allows identity representations to accumulate meaning over time, appear in unanticipated settings, and outlast changes in the individual's preferences—or even their life.

Third, digital replicas involve a separation of identity from contemporaneous human action. Traditionally, speech, performance, and appearance are inseparable from the person producing them in real time. Digital replicas break that link. They allow identity-signifying outputs to be generated autonomously or semi-autonomously-speaking without the person speaking, performing without the person performing.[13] This feature distinguishes replicas from avatars and similar tools that mediate live user action, rather than functioning as independent stand-ins.

Finally, modern digital replicas are capable of recombination and automation. Rather than merely replaying a stored likeness, they apply learned identity traits across new prompts and contexts to generate novel outputs over time.[14] A replica may speak on topics the individual never addressed, respond to questions never anticipated, or combine elements of style, tone, and content the individual never authored. This generative capacity explains why neither realism nor deception alone defines the category. Legally, the significance lies in derivation: a limited act of creation can produce an open-ended stream of identity-signifying outputs, transforming the replica from a finite representation into an ongoing source of identity.

C. A WORKING DEFINITION FOR CONSUMER-ERA DIGITAL REPLICATION

Taken together, these features support the following working definition:

> A digital replica is a persistent, reusable digital system that is referentially anchored to a specific, identifiable human being and capable of generating identity-signifying outputs—such as voice, appearance, behavior, or expressive style—without contemporaneous human performance.

This definition is deliberately forward-looking. It does not hinge on deception, perfect realism, or malicious intent. Instead, it captures the structural features that give rise to recurring legal and governance concerns across consumer, enterprise, and creative contexts.

---

paused-accounts-2026-2 [https://web.archive.org/web/20260326155909/https://www.businessinsider.com/meta-granted-patent-for-ai-llm-bot-dead-paused-accounts-2026-2].

13. *See* Pietro Ruiu et al., *Metaverse & Human Digital Twin: Digital Identity, Biometrics, and Privacy in the Future Virtual Worlds*, 8 MULTIMODAL TECH. & INTERACTION 48, 50–51 (2024).

14. Ajay Bandi, Pydi Venkata Satya Ramesh Adapa & Yudu Eswar Vinay Pratap Kumar Kuchi, *The Power of Generative AI: A Review of Requirements, Models, Input–Output Formats, Evaluation Metrics, and Challenges*, 15 FUTURE INTERNET 260, 268–69 (2023).

II. CONSUMER-FACING DIGITAL REPLICA PRODUCTS

Digital replicas enter everyday life not as a single, unified technology, but through a growing ecosystem of consumer tools that repurpose complex generative systems for familiar tasks. Users are invited to record speech, upload images or video, or provide prior writings in order to generate new content that resembles their own voice, appearance, or manner of response. What distinguishes these products is not merely what they produce, but how they are framed: as tools for accessibility or to facilitate content creation, communication, or continuity.[15] Across voice, visual, and behavioral domains, increasingly sophisticated systems are introduced through interfaces that emphasize ease and speed, while the technical processes that allow identity-related traits to persist and recombine remain largely invisible.

A. VOICE REPLICATION

Voice replication is often the first point at which consumers encounter digital replica technologies, in part because it builds directly on familiar practices.[16] The basic user action is simple: recording speech or uploading existing audio. This can involve reading a short script, providing a few minutes of prior recordings, or recording sample phrases through a web interface. The output is a synthetic voice capable of generating new speech from text, without requiring the user to re-record every time.

In practice, voice replicas are used in a wide range of everyday applications. Creators use them to generate voiceovers for videos or podcasts without repeated recording sessions.[17] Speakers and presenters use them to correct mistakes in recorded talks by typing revised sentences rather than re-recording entire segments. Media producers use voice replicas to dub content into other languages while preserving a familiar vocal identity. For users with speech impairments or degenerative conditions, voice replication can support accessibility by recreating a person's own voice for assistive communication.

Technically, these systems rely on prebuilt speech models that already capture how human voices work in general. The user's recordings are used to adapt that system to sound like a particular person, either almost instantly or through a longer refinement process. From the user's perspective, however, the mechanics are abstracted away. The experience is one of typing text and receiving speech. The resulting voice is not a fixed

15. *See ElevenLabs Impact Program*, ELEVENLABS, https://elevenlabs.io/impact-program [https://web.archive.org/web/20260326160546/https://elevenlabs.io/impact-program] (last visited Mar. 26, 2026); *Voice Cloning: Create a Replica of Your Voice That Sounds Just Like You*, ELEVENLABS, https://elevenlabs.io/voice-cloning [https://web.archive.org/web/20260326160506/https://elevenlabs.io/voice-cloning] (last visited Mar. 26, 2026); Bradley, *supra* note 12.

16. For example, ElevenLabs asks users to upload or record ten to thirty seconds of audio to generate a voice clone. *Voice Cloning*, *supra* note 15.

17. *Id.*

recording, but a reusable system that can speak new words, in new contexts, long after the initial input.[18]

Voice replication also functions as a foundational layer for other forms of digital replication. Many visual replica systems—particularly talking-head videos and avatar presenters—combine synthetic video with synthetic voice, integrating voice models directly into visual pipelines.

B. VISUAL REPLICATION

Visual replication encompasses a broader and more varied set of consumer products than video alone. At the simplest level, users upload a set of photographs—often taken from different angles or with different expressions—to generate new images placing them into novel scenarios.[19] Popular applications include generating portraits in imagined roles or settings, such as historical figures, fantasy characters, or alternate professions. These systems replicate not only facial features but often body shape, posture, clothing, and overall presentation, producing images that depict the user in contexts they never physically occupied.

More advanced workflows build on this foundation. Filmmakers and creators increasingly begin with generative still images—concept art, storyboards, or character portraits—and then use AI tools to animate those images, adding motion, facial expression, and lip synchronization.[20] Earlier consumer tools required a progression from still images to animated portraits to short video clips. Improvements in generative video models are collapsing these steps, making it increasingly possible to generate video directly rather than animating a static image after the fact.[21]

User input scales with output capability. Uploading a single image supports limited animation; providing multiple images or minutes of video footage allows the system to learn how a face and body move over time.[22] Some platforms guide users through

18. *See* Genesis Gregorious Genelza, *A Systematic Literature Review on AI Voice Cloning Generator: A Game Changer or a Threat?*, *4 J. EMERGING TECHS.* 54, 55 (2024).

19. *See Most Realistic AI Image-to-Video Animation Tools*, AKOOL (July 11, 2025), https://akool.com/blog-posts/best-5-ai-video-generators-free-for-turning-images-into-video [https://web.archive.org/web/20260205213915/https://akool.com/blog-posts/best-5-ai-video-generators-free-for-turning-images-into-videos].

20. *See id.*

21. *See, e.g.*, *AI Video Generator*, ADOBE, https://www.adobe.com/products/firefly/features/ai-video-generator.html [https://perma.cc/EA2C-KA4N?type=image] (last visited Mar. 26, 2026); *Text-to-Video: A Step-by-Step Guide to Creating Your First Video Using AI*, KLING AI (Aug. 29, 2025), https://kling.ai/blog/text-to-video-ai-creation-guide [https://perma.cc/MF3U-7RSB]; *Free AI Video Creation*, OPENART, https://openart.ai/suite/create-video/kling-3-omni [https://perma.cc/TYP8-YT8X] (last visited Mar. 26, 2026).

22. *See, e.g.*, David Ephraim, *The Complete Guide to Creating AI Avatars: From FLUX.1 to Video Generation*, ATAK INTERACTIVE (Aug. 1, 2025), https://www.atakinteractive.com/blog/the-complete-guide-to-creating-ai-avatars-from-flux.1-to-video-generation [https://web.archive.org/web/20260326162045/https://www.atakinteractive.com/blog/the-complete-guide-to-creating-ai-avatars-from-flux.1-to-video-generation] (suggesting using fifteen to fifty photos or at least ten videos, each ten seconds long).

capture sessions, asking them to read scripted sentences aloud to a camera so the system can model facial motion during speech. The result can be a reusable digital presenter capable of delivering new scripts on demand.

These visual replicas are widely used for content creation and professional communication. Companies deploy synthetic talking heads for corporate training videos, onboarding materials, and internal communications.[23] Marketers use them to generate personalized or localized video content at scale. Educators and nonprofits use animated historical figures or spokespersons to deliver information in engaging formats. In many of these cases, visual replicas are tightly integrated with voice replication technologies, producing a combined audiovisual stand-in that can be updated without re-recording.

As visual replicas have become more capable, the systems that generate them have also become more mediated. What once required studio-grade CGI, motion capture, or manual compositing can now be accomplished through browser-based tools. At the same time, the rendering, animation, and storage processes are handled entirely by platforms, leaving users with clear inputs and outputs but little visibility into how their likeness is maintained or extended over time.

C. BEHAVIORAL AND PERSONA REPLICATION

Behavioral and persona replication systems operate through a different form of user contribution. Rather than capturing new media, these tools are built from what users have already produced: emails, documents, messages, posts, or transcripts.[24] The system analyzes these materials to learn patterns not only of language, but often of preferences, expertise, or values reflected in prior decisions and responses.

The resulting replicas are used in a variety of roles. Some are designed to draft or respond to communications in a user's typical style. Others function as customer-support or fan-engagement chatbots, answering questions in a way that reflects a particular brand voice, creator persona, or subject-matter expertise.[25] More personal applications include companionship agents or griefbots trained on the communications

23. *See* Roberto Gozalo-Brizuela & Eduardo Garrido-Merchán, *Applications of Video Generation Models in the Journal of Computer Science*, 20 J. COMPUT. SCI. 801, 801–818 (2024).

24. *See How It Works*, TONECLONE https://toneclone.ai/#how-it-works [https://web.archive.org/web/20260326162428/https://toneclone.ai/#how-it-works] (last visited Mar. 26. 2026); Brooke Steinberg, *People Are Making Digital Clones" of Themselves—To Do Their Work for Them*, N.Y. POST (Apr. 15, 2024), https://nypost.com/2024/04/15/lifestyle/people-making-digital-clones-of-themselves-to-do-their-work/ [https://web.archive.org/web/20260216065115/https://nypost.com/2024/04/15/lifestyle/people-making-digital-clones-of-themselves-to-do-their-work/]; Amanda Caswell, *I Made a Digital Twin of Myself in ChatGPT—And It Changed How I Work Every Day*, TOM'S GUIDE (Feb. 17, 2026) https://www.tomsguide.com/ai/i-made-a-digital-twin-of-myself-in-chatgpt-and-it-changed-how-i-work-every-day [https://web.archive.org/web/20260308222028/https://www.tomsguide.com/ai/i-made-a-digital-twin-of-myself-in-chatgpt-and-it-changed-how-i-work-every-day].

25. *See* Heng Gu, Senthil Chandrasegaran & Peter Lloyd, *Synthetic Users: Insights From Designers' Interactions With Persona-Based Chatbots*, 39 A.I. FOR ENG'R, DESIGN, ANALYSIS & MFG. 1, 3 (2025).

of deceased individuals, allowing family members or loved ones to interact with a simulation shaped by that person's prior expressions.

In these systems, the output is not a static artifact but an ongoing interaction. Users prompt the system with questions or scenarios, and the replica generates responses dynamically. Unlike voice or visual replicas, behavioral systems are often constrained not just by how someone speaks, but by what they are understood to know, believe, or prioritize. A replica trained on professional correspondence may answer technical questions; one trained on personal messages may reproduce emotional tone or relational patterns.

From the user's perspective, setup may involve uploading archives, linking accounts, or interacting with the system over time. Once established, the system operates with a degree of autonomy, generating responses that feel personally attributable even as the mechanisms that produce them remain opaque.[26]

D. PATTERNS ACROSS CONSUMER SYSTEMS

Across voice, visual, and behavioral domains, consumer digital replicas share a common structure: modest user inputs yield systems capable of producing new outputs across time and context. As these tools have matured, two features have driven their widespread adoption. First, refinement: Outputs have become more natural, flexible, and convincing. Second, accessibility: Platforms increasingly make these capabilities cheap, fast, and easy to use. Just as importantly, they are introduced and experienced as tools—for editing, creating, communicating, or scaling—not as technologies of identity.

That framing has proven durable even as capabilities expand. Voice and visual replication, once distinct, are already routinely combined in talking-head videos, avatar presenters, and synthetic media workflows. The likely next step is not simply better quality, but broader integration: the addition of behavioral replication to systems that already speak and appear. At that point, digital replicas would not only sound and look like a person, but also respond, prioritize, or interact in ways associated with that person.

Whether such systems will continue to be understood as tools or begin to feel like extensions of the self remains an open question. What is clear is that this evolution is unfolding against a backdrop of normalization. By the time more fully integrated replicas become feasible, the practices that enable them—recording, uploading, delegating—will already feel routine. That context matters. It shapes not only how these technologies are adopted, but how their significance is recognized, or overlooked, as they move from helpful utilities toward something closer to a stand-in.

26. *See, e.g.*, PICKAXE, *supra* note 12 (noting that the tool automatically updates AI agent's sources on a daily basis "without any manual work").

III. GOVERNANCE ISSUES IN CONSUMER DIGITAL REPLICATION

The lifecycle of most consumer digital replicas follows a common pattern. Users consent to creation through onboarding and identity verification; the platform uses the provided inputs to train a replica model; the replica is then used—often repeatedly—within the platform's tools and workflows; its outputs are constrained through safeguards and platform rules; access is mediated through pricing and subscriptions; and, if the user later chooses to leave, the replica may be deleted or disabled under conditions set by the platform.[27] At each stage of this lifecycle, control over identity shifts incrementally from the user to the platform. This section will focus on two governance issues that emerge from that shift: the limits of revocable consent over time, and ongoing contests over the permissible scope of use.

A. THE LIMITS OF REVOCABLE CONSENT

Consumer digital replicas often begin with consent practices that reflect the seriousness of what is being authorized.[28] Because replicating a person's voice, likeness, or expressive style implicates core aspects of identity, platforms frequently require more than passive agreement to terms of service. Users may be asked to verify their identity through one-time codes, record themselves reading acknowledgments aloud, or otherwise confirm, in real time, that they are authorizing the use of specific inputs they are uploading. These measures serve an important threshold function: they help ensure that the person whose identity is being replicated is the one giving consent, and that consent is tied to identifiable materials.

That initial authorization, however, rarely stands alone. In addition to enabling the creation of a particular digital replica on the platform, users are often asked whether their inputs may also be used for broader training purposes, such as improving models or refining system performance. At this point, consent begins to operate across time.

27. *See, e.g.*, Jess Diaz-Gomes, *Disagreeable Content*, SYNESTHESIA, https://help.synthesia.io/en/articles/8330530-disagreeable-content [https://web.archive.org/web/20260407025321/https://help.synthesia.io/en/articles/8330530-disagreeable-content] (last visited Mar. 26, 2026) (moderating what type of content an AI avatar can communicate—e.g., legal advice, news reporting, conversations about religious practices—based on subscription plan and degree of specificity); *Customer Terms of Service*, SYNESTHESIA (effective Feb. 23, 2024), https://www.synthesia.io/legal/customer-terms-of-service [https://web.archive.org/web/20260325014043/https://www.synthesia.io/legal/customer-terms-of-service] (reserving discretion not to retain or update custom avatars after termination and committing to deleting them upon account deletion, subject to legal limits).

28. *See, e.g.*, *Create an Avatar*, SYNESTHESIA, https://docs.synthesia.io/docs/personal-avatars [https://web.archive.org/web/20251230064332/https://docs.synthesia.io/docs/personal-avatars] (last visited Mar. 26, 2026) (requiring live consent recording, with the same person in the personal avatar footage creation).

A choice made at onboarding may authorize downstream uses that persist independently of the user's continued engagement with the service.[29]

Crucially, users may later want to disentangle those authorizations. A person might wish to continue using their digital replica—because it is embedded in workflows, accessibility tools, or creative projects—while retracting consent for broader training or reuse beyond that specific replica. Whether such selective withdrawal is possible depends entirely on platform design. Some systems purport to offer granular controls; others do not. In many cases, users are faced with an all-or-nothing choice: retain the replica along with any associated training permissions, or exit the system entirely.[30]

Deletion brings the temporal stakes into sharper relief. When platforms offer deletion, they typically mean deletion of the particular digital replica model associated with the user's account on that platform—disabling access to that replica and, in some cases, removing stored inputs used to create it. What deletion does not always make clear is whether and to what extent training effects derived from those inputs persist elsewhere in the system. A user may be able to stop using a replica without knowing whether their identity data continues to shape other models or features.

This matters because identity-linked training data is not a transient resource. Unlike storage space or computer processing power, trained representations of a person's voice, likeness, or expressive patterns can remain valuable, reusable, and sensitive long after the user's immediate relationship with the platform has ended. Retained training data may be repurposed as systems evolve, exposed in the event of a breach, or applied in contexts the user did not contemplate when consenting.[31] Even where no misuse occurs, continued retention without ongoing authorization raises concerns about autonomy and control that are not resolved simply because a service was once provided in exchange.

The practical significance of these issues is compounded by limits on verification and enforcement. Users must largely trust platform representations about how training works, what deletion reaches, and whether consent has been meaningfully withdrawn. Those claims are difficult to test.[32] Once inputs have been incorporated into training processes—particularly where influence is diffuse or embedded in model parameters—users lack tools to inspect or audit compliance. Even in litigation, challenges remain. Claims about retention or reuse of identity data may be difficult to substantiate without access to internal systems, and obtaining such access typically requires surviving early

29. *See* Jennifer King et al., *User Privacy and Large Language Models: An Analysis of Frontier Developers' Privacy Policies* 1, 6, ARXIV (Sept. 10, 2025), https://arxiv.org/html/2509.05382v1 [https://web.archive.org/web/20260213041359/https://arxiv.org/html/2509.05382v1].

30. *See* Hannah Ruschemeier, *Generative AI and Data Protection* 11, 13, CAMBRIDGE F. ON AI: LAW & GOVERNANCE (2025).

31. *See* King et al., *supra* note 29, at 8–9.

32. *See, e.g.*, Judy Hanwen Shen et al., *The Limits of AI Data Transparency Policy: Three Disclosure Fallacies* 1, 10–12, ARXIV (Jan. 26, 2026), https://arxiv.org/html/2601.18127v1#S5 [https://web.archive.org/web/20260304134349/https://arxiv.org/html/2601.18127v1#S5] (noting that AI data-transparency rules often depend on self-reported disclosures that are difficult to audit or verify in practice—producing an "enforcement gap" that limits meaningful testing of whether companies actually follow their stated data-handling practices).

motions to dismiss.[33] Absent clear evidence of continued retention or use, proving that consent was exceeded can be a significant hurdle. In this sense, revocable consent operates less as a technical guarantee than as a promise whose credibility depends on platform governance and transparency.

B. CONTESTING THE SCOPE OF USE

If consent governs entry into digital replica systems, control over how replicas function day to day is exercised through platform architecture. Safeguards, portability limitations, pricing structures, and evolving contractual terms collectively shape the scope of permissible use.

Safeguards are the most visible expression of this control. They typically operate at the level of prompts and outputs, rather than at the level of training or persistence.[34] The replica continues to exist, but certain expressions or applications are restricted after the fact. In many systems, these safeguards are enforced primarily through automated monitoring. Automation can produce both underinclusive and overinclusive outcomes. In some situations, prohibited uses may slip through; in others, legitimate or advertised uses may be blocked because the system misinterprets intent or context.

Users often respond by adapting. They rephrase prompts, divide tasks into stages, chain multiple tools together, or move parts of a workflow across platforms. These work-arounds may be benign, questionable, or expressly disallowed, but they are a predictable response to systems that promise broad functionality while enforcing it imperfectly. Whether a restriction can be challenged depends on platform infrastructure. Some platforms provide escalation mechanisms—such as human review or expanded usage licenses that permit broader application—while others rely almost entirely on automated enforcement, leaving little room to contest errors or edge cases.

Economic controls further shape the scope of use. Digital replicas are commonly tied to subscription models, usage caps, or tiered access.[35] Tools initially framed as personal conveniences may become relied upon professionally or commercially. As terms of service evolve, continued access to a trained replica may require higher payments, additional licenses, or acceptance of new limitations.[36] In these settings, pricing does more than compensate for services rendered; it regulates who may rely on scalable identity replication and for what purposes.

---

33. *See* Valerian Stolpe, *AI Without Audit Trails Is Becoming a Legal and Governance Liability*, ZL TECH (Dec. 5, 2025), https://www.zlti.com/blog/ai-without-audit-trails-is-becoming-a-legal-and-governance-liability/ [https://web.archive.org/web/20260213071922/https://www.zlti.com/blog/ai-without-audit-trails-is-becoming-a-legal-and-governance-liability/].

34. *See* Markus Anderljung, Julian Hazell & Moritz von Knebel, *Protecting Society from AI Misuse: When Are Restrictions on Capabilities Warranted?,* 40 AI & SOC'Y 3841 (2024).

35. *See* Rupsa Majumdar, Anjula Gurtoo & Minnu Maileckal, *Developing a Data Pricing Framework for Data Exchange*, 11, 4 FUTURE BUS. J. 1, 7 (2025).

36. *See* Nontokozo Mokoena & Ibidun Christiana Obagbuwa, *An Analysis of Artificial Intelligence Automation in Digital Music Streaming Platforms for Improving Consumer Subscription Responses: A Review* 7 FRONTIERS IN AI 2 (2025).

Portability intensifies this contest. At present, most consumer digital replicas are not interoperable. Users can typically export outputs—audio files, images, video clips—and recombine them through traditional editing workflows, but the underlying trained replica usually cannot move across platforms. Identity remains bound to the system that trained it. At the same time, there are clear indications that this may change. Generative AI tools that once required standalone platforms are now embedded directly within creative software ecosystems, allowing users to access similar capabilities across environments.[37] As voice, visual, and behavioral replication converge, comparable expectations of interoperability are likely to emerge for digital replicas themselves.

Until then, users retain alternatives. They can rebuild replicas on different platforms, stitch together outputs using older tools, or dispense with replicas altogether and rely on their real selves—time and stamina permitting. Still, as digital replicas become normalized tools for scalable communication, accessibility, and creation, these alternatives carry increasing cost. The scope of permissible use is therefore not fixed at creation, but continually contested through safeguards, pricing, portability, and design choices that determine how identity may be expressed, constrained, or monetized over time.

IV. CONCLUSION

Digital replicas are no longer a speculative or marginal technology confined to entertainment studios, research labs, or controversial deepfakes. They are increasingly ordinary consumer tools, encountered through products that promise convenience, accessibility, and creative efficiency. The legal challenge they pose is not limited to deception, labor displacement, or misappropriation. It lies in how these systems externalize human identity into persistent, reusable digital forms that operate beyond the moment of human action.

Existing legal definitions of "digital replica" reflect the narrow contexts in which they developed. They are reactive and purpose-built—designed to address preservation, optimization, performer protection, or fraud. Each responds to a real concern, but none captures the broader structural shift now underway. Consumer digital replicas are identity-bearing systems capable of generating new expressive outputs over time. Their legal significance does not depend on realism or intent, but on derivation and continuity: a limited act of contribution can produce an open-ended stream of outputs that remain meaningfully tied to a particular person.

That understanding clarifies why governance concerns arise even where replicas are disclosed, consented to, and self-created. Across voice, visual, and behavioral domains, these technologies are framed as tools rather than as forms of identity delegation.

37. Generative AI video models like Runway, PikaLabs, and OpenAI's Sora are now available directly in video editing software like Adobe Premiere. *See* Jess Weatherbed, *Adobe Premiere Pro Is Getting Generative AI Video Tools,* THE VERGE (Apr. 15, 2024), https://www.theverge.com/2024/4/15/24130804/adobe-premiere-pro-firefly-video-generative-ai-openai-sora [https://web.archive.org/web/20250723005223/https://www.theverge.com/2024/4/15/24130804/adobe-premiere-pro-firefly-video-generative-ai-openai-sora].

Consent is concentrated at onboarding, while platform design governs everything that follows. Over time, control migrates. Decisions about retention, training, safeguards, pricing, and portability are largely set by platforms, not users. Revocability, where it exists, is often partial and difficult to verify.

This gap exposes the limits of existing legal frameworks. Doctrines built around deception, discrete consent, and observable misuse do not map well onto openly synthetic systems that persist and recombine identity over time. The problem is not use—many consumer digital replicas are beneficial—but governance. As these tools become more integrated and normalized, the costs of treating them as ordinary software features will grow. Without a structural account of digital replicas, law arrives only after control over identity has already been decided.