

Echoes of the Prior: A Computational Phenomenology of Forgetting

ANONYMOUS AUTHOR(S)

SUBMISSION ID: ARTPL_110S1

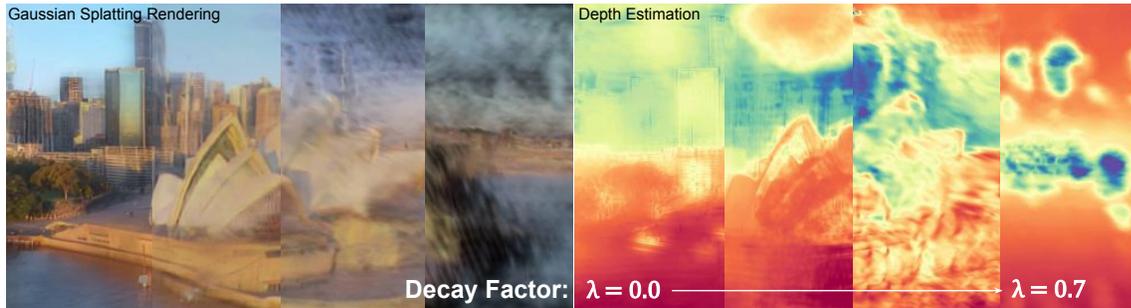


Fig. 1. *Echoes of the Prior*. A computational simulation of the fragility of memory in the face of entropic loss.

Memory is not merely the storage of data; it is the scaffolding of reality. When biological memory fades, the world does not simply turn black; it regresses into an unrecognizable chaos. *Echoes of the Prior* is an interactive installation that attempts to visualize this subjective phenomenology of *forgetting*. By inducing controlled synaptic decay within a Feed-Forward 3D Reconstruction model, we simulate the erosion of the brain's predictive priors. We position the Neural Network not as a tool for engineering, but as a cognitive proxy - a *silicon brain* whose structural degeneration offers us a glimpse into the disorienting, poetic, and terrifying experience of losing one's grip on the world. Ultimately, we offer this framework as a catalyst, inviting the wider community to explore the uncharted potential of neuromorphic aesthetics in visualizing the fragility of intelligence. Interactive demo: <https://decart-4d.github.io/>.

CCS Concepts: • **Human-centered computing** → **Interaction design theory, concepts and paradigms**; • **Computing methodologies** → **Non-photorealistic rendering**; • **Applied computing** → **Media arts**.

Additional Key Words and Phrases: Computational Phenomenology, Glitch Aesthetics, Interactive Installation, 3D Gaussian Splatting, Neuromorphic Art, Real-time Performance System

ACM Reference Format:

Anonymous Author(s). 2026. Echoes of the Prior: A Computational Phenomenology of Forgetting. In *ACM SIGGRAPH 2026 Art Papers*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 Introduction: The Invisible Process

“What I cannot create, I do not understand” is a famous quote by physicist Richard Feynman. While neuroscience can explain the *mechanism* of Alzheimer's – the plaques, the tangles, the synaptic pruning – it cannot show us the experience. We know that the patient forgets, but we do not know *what the world looks like* to them as it fades.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

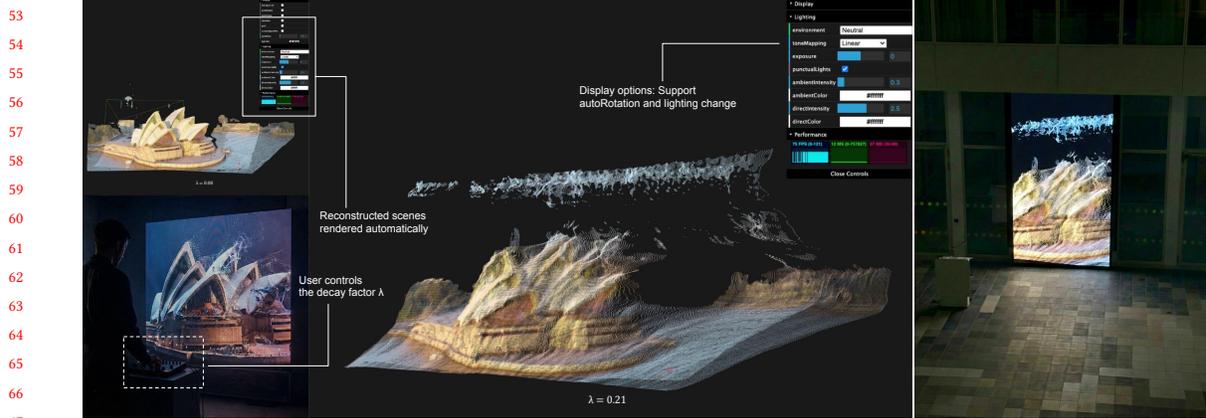


Fig. 2. **Installation Design of *Echoes of the Prior***. Snapshot of the real-time visual feedback ($\lambda = 0.21$). The viewport displays the immediate structural dissolution of the 3D scene as the user adjusts the decay factor.

Modern Vision Transformers share structural principles with biological vision – distributed representation, hierarchical feature extraction, and predictive coding [Yamins and DiCarlo 2016]. While these correspondences are productive analogies rather than mechanistic equivalences [Bowers et al. 2023; Lindsay 2021], they invite a speculative proposition: if these models can approximate aspects of human visual cognition, then their pathology may serve as an *artistic analogy* for the pathology of the mind.

Echoes of the Prior operates on this speculative premise. We treat a state-of-the-art 3D Gaussian Splatting (3DGS) model not as a renderer, but as a proxy observer – a *silicon brain* whose components we map, by analogy, onto cognitive functions. By surgically introducing noise into its semantic priors (analogous to *long-term memory*) and image encoder (analogous to the *sensory cortex*), we force the machine to reconstruct reality through a damaged architecture. The resulting imagery – melting sky, ghostly architectures, and dissolving forms – serves as a visual translation of fading mind, making the invisible internal experience of forgetting visible, visceral, and shared.

We title our installation *DecArt* (a portmanteau of *Decay* and *Art*). Pronounced like the French philosopher *Descartes*, we elaborate on the philosophical implications of this choice in Sec. 5.1.

2 Related Work

Positioning *Echoes of the Prior* requires navigating the intersection of four distinct disciplines: computational aesthetics, artistic engagement with neurodegeneration, volumetric computer vision, and theoretical neuroscience.

2.1 From Machine Vision to Machine Hallucination

In the domain of AI art, early works focused on visualizing the *constructive* capacity of neural networks. **DeepDream** [Mordvintsev et al. 2015] and **Neural Style Transfer** [Gatys et al. 2016] exposed the internal representations of CNNs by maximizing feature activation, essentially asking the machine to “see more” of what it knows. Contemporary media artists like **Refik Anadol** [Anadol 2021] have further expanded this by treating massive datasets as fluid memories, visualizing the *learning* process of the machine.

Our work inverts this paradigm. Instead of visualizing the *genesis* of intelligence (learning), we visualize its *entropy* (forgetting). We align with the tradition of **Glitch Art** [Menkman 2011], which views technological failure not as a bug,

105 but as a revelation of the underlying medium. However, unlike traditional glitches that operate on the surface pixel
106 level (e.g., datamoshing), our system performs **structural pathology**: we erode the internal cognitive weights and
107 predictive priors themselves. This creates a form of “Deep Glitch” that emulates neurodegeneration rather than mere
108 signal transmission error.
109

111 2.2 Art, Memory, and Neurodegeneration

113 Our work also situates itself within a broader artistic tradition engaging with memory loss and cognitive decline.
114 One well-known precedent is the work of William Utermohlen, an artist diagnosed with Alzheimer’s disease in 1995,
115 whose self-portraits over subsequent years document a gradual shift from figurative realism toward abstraction, often
116 interpreted as a first-person record of representational breakdown [Crutch et al. 2001]. Similarly, the late paintings of
117 Willem de Kooning, produced during a period of cognitive decline, have sparked debate about whether their luminous
118 simplicity reflects genuine artistic evolution or the influence of disease [Espinel 1996]. These cases demonstrate that
119 neurodegeneration can produce visually compelling imagery while simultaneously raising complex ethical questions
120 about authorship and artistic intent (see Sec. 5.2).
121

123 In the domain of interactive and new media art, Memo Akten’s installation *Learning to See* [Akten et al. 2019]
124 provides a related example: a neural network constrained by its training data hallucinates misreadings of live camera
125 input, revealing how machine perception can become structurally distorted. Laurie Anderson and Hsin-Chien Huang’s
126 VR installation *Chalkroom* [Anderson and Huang 2017] similarly immerses viewers in rooms of dissolving words and
127 narratives, evoking the fragility of linguistic memory. Collaborative art – science initiatives such as *Created Out of*
128 *Mind* [Crutch et al. 2019] at the Wellcome Collection have further brought together artists, neuroscientists, and people
129 living with dementia to co-create works that reshape public understanding of these conditions.
130

132 Our contribution differs from these precedents in a key respect: rather than depicting memory loss through narrative,
133 metaphor, or the artist’s own declining hand, we engineer it directly within the cognitive architecture of the machine
134 itself, creating what we term a “Deep Glitch” – a structural pathology of the model’s internal representations.
135

137 2.3 Spatializing the Past: 3D Reconstruction as a Cognitive Medium

138 While photography freezes a specific moment (a temporal slice), 3D reconstruction captures a specific environment
139 (a spatial slice). We argue that **volumetric 3D representation** is intrinsically more akin to the phenomenology of
140 biological memory than flat 2D imagery. Neuroscientific evidence suggests that memory encoding is deeply intertwined
141 with spatial navigation (e.g., in the hippocampus) [O’Keefe and Dostrovsky 1971]. To remember is not just to view an
142 image; it is to inhabit a space.
143

145 The evolution of 3D reconstruction, from explicit photogrammetry meshes to implicit **Neural Radiance Fields**
146 (**NeRF**) [Mildenhall et al. 2020] and now **3D Gaussian Splatting (3DGS)** [Kerbl et al. 2023], reflects a shift towards
147 more organic representations of reality. Unlike polygonal meshes which are defined by rigid surfaces, volumetric
148 approaches represent the world as a probabilistic field of density and color.
149

150 In this work, we utilize 3DGS as our implementation carrier not merely for its real-time rendering capabilities, but for
151 its unique aesthetic properties. Unlike generative video models that hallucinate temporal continuity, 3DGS reconstructs
152 the *spatial continuity* of a scene. By subjecting this volumetric representation to decay, we visualize the disintegration
153 of the spatial context itself, simulating how a patient might lose their grip on the *where* and *how* of a memory, rather
154 than just the *what*.
155

2.4 Computational Psychiatry and Predictive Coding

Theoretically, our decay functions are grounded in the **Free Energy Principle** [Friston 2010], which posits that the brain minimizes surprise by constantly predicting sensory input based on internal priors. In this framework, psychopathology is often modeled as a failure of this predictive machinery [Corlett et al. 2019].

Our implementation of the Sensory Decay ($\mathcal{G}_{\text{sense}}$) draws directly from the **Neural Gain Theory of Aging** [Li et al. 2001], which hypothesizes that cognitive decline is driven by a reduction in the signal-to-noise ratio (SNR) of neuromodulators (like dopamine) in the cortex. By injecting noise into the weights of our vision encoder, we create a functional analogy of this biological process [Cichy and Kaiser 2019]. Furthermore, our simulation of hallucinations aligns with models of **Computational Psychiatry**, which suggest that when sensory precision is low (as in our Spectral Amnesia module), the brain over-relies on top-down priors, leading to the perception of phantom stimuli [Powers et al. 2017].

Complementing this, our Memory Decay (\mathcal{F}_{mem}) simulates the breakdown of top-down priors, grounded in the theory of **Attractor Dynamics** [Hopfield 1982]. In biological networks, memories are stored as stable states (attractors) within a high-dimensional manifold. Our geometric rotation of the latent space effectively destabilizes these attractors, simulating the phenomenology of **Semantic Dementia** [Warrington 1975], a condition where the structural links between concepts erode. While the machine retains the raw sensory data, it loses the “semantic glue” that holds the geometry together, mirroring how patients with associative agnosia perceive form without meaning.

3 Methodology: Simulating Cognitive Decay

Modern foundational models [Oquab et al. 2024; Siméoni et al. 2025] do not store pixels; they store semantic concepts in a high-dimensional latent space. To simulate the phenomenology of forgetting, we do not merely degrade the image pixels; we perform surgical interventions on the cognitive architecture of the machine.

Our approach creates a structural analogy between a state-of-the-art 3D reconstruction model [Lin et al. 2025] based on feed-forward 3D Gaussian Splatting (FF-3DGS), and the Predictive Coding framework [Friston 2010] of the human brain, synthesizing reality through the fusion of two distinct neural pathways: *memory* (prior) and *evidence* (sensation).

3.1 The Silicon Proxy: Two Streams Mapping Mind to Code

Can a machine possess memory and cognition? If so, what does it look like?

The act of perceiving the world is a negotiation between what our eyes see – the evidence from *Sensory Cortex*, and what our brains expect – what neuroscience terms the *Engram*, the physical trace of memory. We draw a functional analogy between this biological duality and the specific components of our computational architecture (see Sec. 6.1 for a discussion of the boundaries of this mapping):

3.1.1 The Sensory Stream: Processor Weights. Drawing a functional analogy to the biological retina and V1 cortex, our model employs a shallow CNN to process raw pixel inputs, capturing the texture and local geometry of the *Now*. This sensory signal will serve as a reinforcement to the memory, and together with memory, will be injected into a fusion block for spatial perception and reconstruction [Ranftl et al. 2021].

3.1.2 The Memory Stream: Semantic Features. In an analogy to how biological memory stores concepts rather than raw data, our system utilizes vision foundational models to encode high-dimensional semantic priors.

209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260

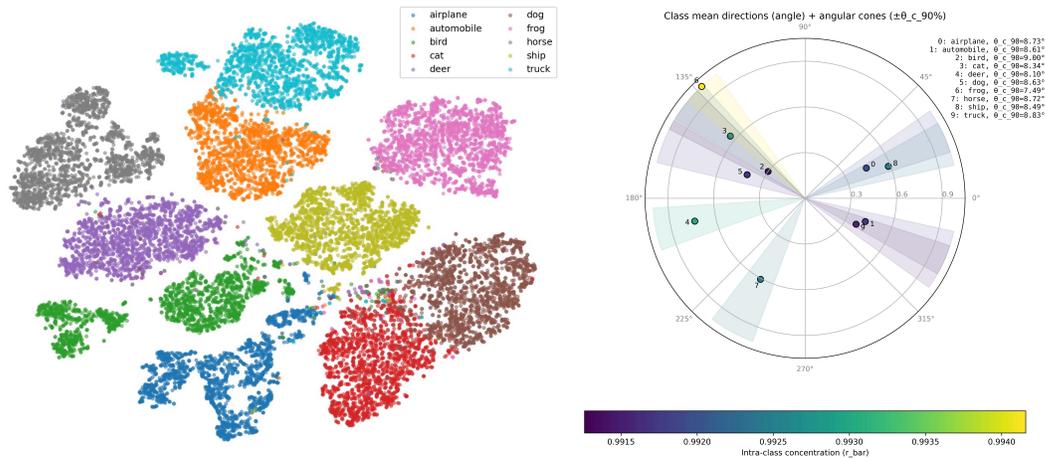


Fig. 3. t-SNE visualization of the DINOv2 latent space. (Left) Semantic concepts cluster into distinct manifolds, representing the structured *healthy* memory baseline. (Right) Cosine similarity metrics confirm that intra-class features are more closely aligned than inter-class features, validating the semantic coherence of the latent space.

Specifically, we use *DINOv2* [Oquab et al. 2024], a state-of-the-art visual understanding model as the *prior* producer in our feed-forward 3D reconstruction model. To give an intuitive sense of this prior, we visualize its latent representation space. As shown in Fig. 3 on the left, in this high-dimensional (1024D) latent space, for instance, visual features of the same semantic concept cluster together and form a distinct manifold, diverging significantly from others in terms of geometric similarity metrics. We employ t-SNE and PCA to project these 1024D latent vectors into 2D space for better visualization. These latent representations serve as the *healthy baseline*: the long-term memory of a healthy silicon brain – structured, dense, and semantically coherent.

In a healthy state, perception is the successful integration of these two streams. Our installation visualizes what happens when these streams decay asynchronously.

3.2 The Mechanics of Forgetting: Two Decay Functions Simulating Biological Entropy

Human aging is characterized by two distinct processes: sensory degradation and cognitive decline. To simulate both processes, we introduce two distinct entropy operators, $\mathcal{G}_{\text{sense}}$ and \mathcal{F}_{mem} , grounded in the *Neural Gain Theory of Aging* [Li et al. 2001], which posits that both types of degradation are driven by a reduction in the signal-to-noise ratio within specific neuromodulatory systems.

3.2.1 Sensory Decay: The Erosion of Visual Primitives $\mathcal{G}_{\text{sense}}$. While the semantic stream deals with concepts, the sensory stream deals with *visual primitives* – edges, textures, and local gradients. A naive simulation of sensory decay would simply degrade the input signal (i.e., pixel-level corruptions), simulating external visual impairments. However, biological aging is a neurological process where the processing circuitry itself degrades.

To simulate this, we perform a perturbation on the CNN weights that acts as the machine’s primary visual cortex (as described in Sec. 3.1.1). We implement a compound decay function $\mathcal{G}_{\text{sense}}$ that creates a diseased clone of the sensory encoder, modeling three distinct biological failures:

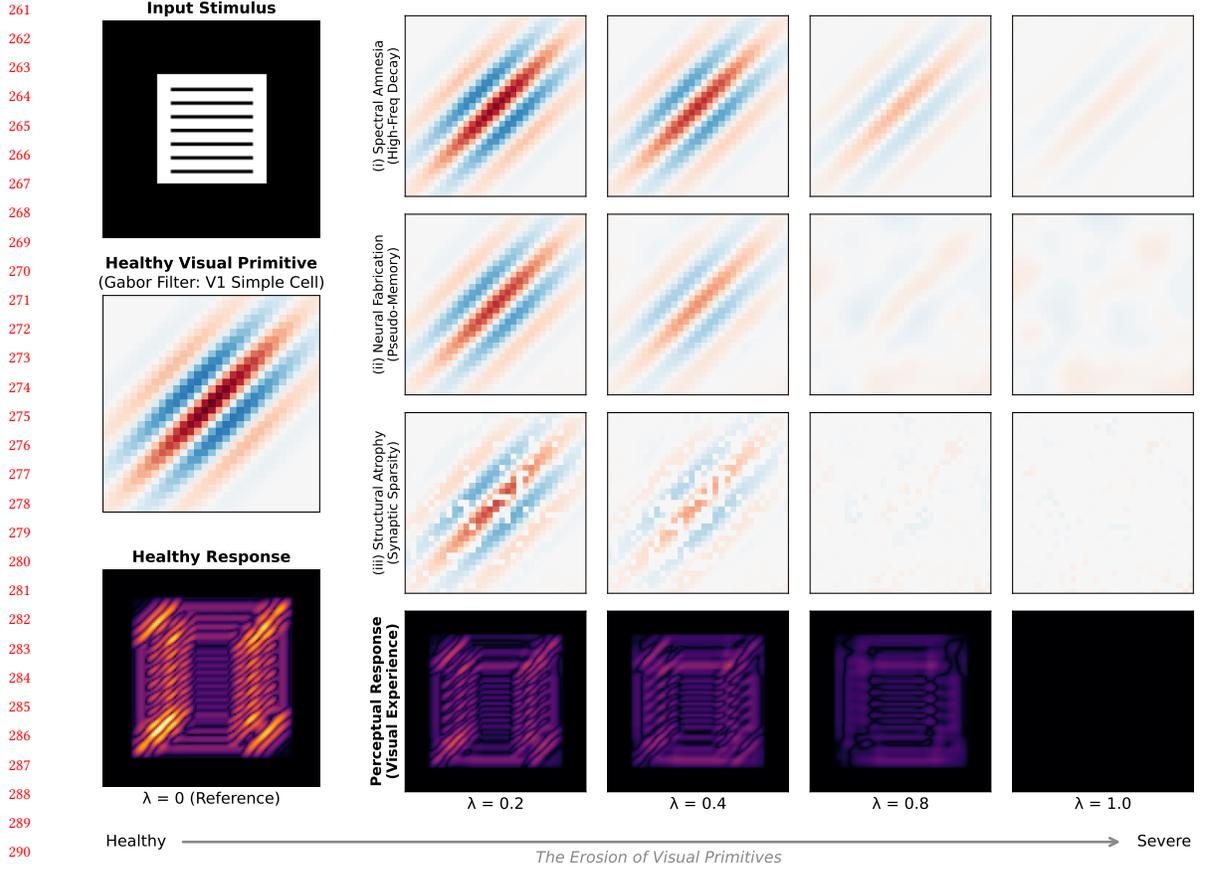


Fig. 4. **Simulating Sensory Decay $\mathcal{G}_{\text{sense}}$.** (Left) **Healthy Baseline ($\lambda = 0$):** Response of a standard Gabor filter (V1 primitive) to a sharp stimulus. (Right) **Pathology:** As λ increases, the filter degrades via three mechanisms: (i) **Spectral Amnesia** (loss of high-frequency acuity/sharpness); (ii) **Neural Fabrication** (injection of coherent noise/hallucinations); and (iii) **Structural Atrophy** (synaptic sparsity leading to signal loss).

(i) **Spectral Amnesia (High-Frequency Domain Decay):** Biological aging often begins with the loss of high-frequency acuity [Owsley et al. 1983]. To simulate this, we transform the convolutional kernels into the frequency domain via Fast Fourier Transform (FFT) and then apply a radial mask that selectively dampens high-frequency components relative to the decay factor λ ¹. This operation is grounded in the functional architecture of the primary visual cortex (V1), where simple cells act as spatiotemporal frequency filters [De Valois et al. 1982]. By dampening high frequencies in the kernel weights directly, rather than the input image, we simulate the degradation of the neurons' receptive fields themselves, rendering them incapable of resonating with fine details. **The Effect:** As shown in Fig. 4, the machine does not see a blurry image; it simply forgets how to detect sharpness. The neural filters responsible for recognizing crisp edges and fine textures (modeled as Gabor-like simple cells [Hubel and Wiesel 1962]) are chemically suppressed, simulating the onset of cortical cataracts.

¹The decay factor in $\mathcal{G}_{\text{sense}}$ can be distinct from the one in memory decay \mathcal{F}_{mem} for asynchronous control, and we keep the same symbols here to avoid complicating the notation.

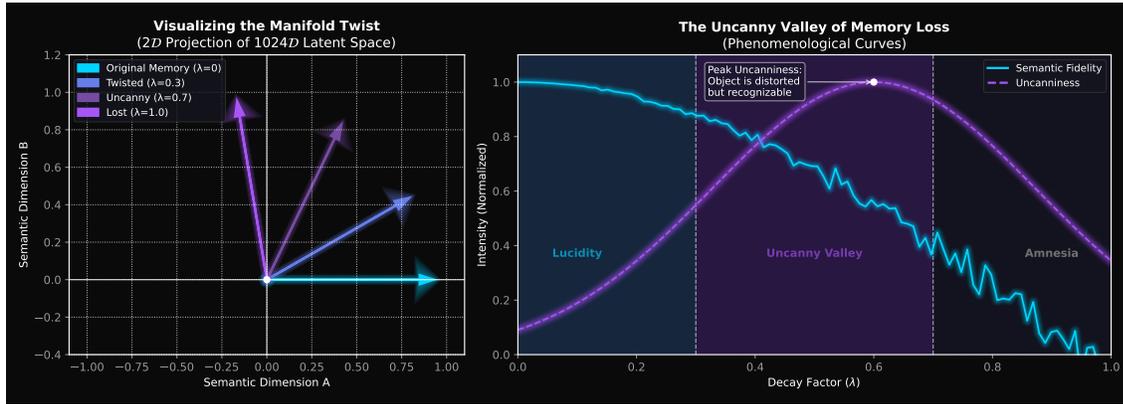


Fig. 5. **Topology of Memory Decay.** (a) **Manifold Twist:** Visualization of the orthogonal rotation $R(\lambda)$. As λ increases (blue \rightarrow purple), the feature vector rotates away from its semantic alignment while preserving its norm. (b) **Computational Uncanny Valley:** A computationally defined model of visual *Uncanniness* (purple), defined as the product of structural integrity (norm) and semantic distortion ($1 - \text{cosine similarity}$). The peak ($0.3 < \lambda < 0.7$) corresponds to the zone where the representation is structurally strong yet semantically twisted. Perceptual research predicts that such feature-inconsistent stimuli are maximally disturbing [Moore 2012; Seyama and Nagayama 2007]; formal viewer validation remains future work.

(ii) **Neural Fabrication (Pseudo-Memory Injection):** The brain abhors a sensory vacuum. When visual inputs degrade, the visual cortex often generates spontaneous activity to fill the void (e.g., Charles Bonnet Syndrome [Reichert et al. 2013]). We model this by injecting spatially smoothed Gaussian noise into the weights. This approximates the “correlated noise” observed in neural circuits when inhibitory control weakens [Shadlen and Newsome 1998]. Unlike white noise, this smooth noise creates coherent but fictitious patterns in the convolution filters. **The Effect:** This causes the machine to *hallucinate* textures that do not exist, overlaying the real world with a dream-like, phantom grain.

(iii) **Structural Atrophy (Synaptic Sparsity):** Finally, we simulate the physical death of neurons. We apply a stochastic mask to the *Weight* matrix, zeroing out connections based on λ . Simultaneously, the *Bias* terms in the convolutional layers are decayed towards zero, representing a loss of neural “confidence” or activation potential. **The Effect:** This creates a *sparse* perception and the image becomes *ghostly*. As the machine loses the confidence to assert the existence of matter, the reconstructed 3D world becomes translucent and ethereal, creating a visual metaphor for a consciousness that is slowly detaching from physical reality.

3.2.2 *Memory Decay: Manifold Distortion and Cognitive Entropy \mathcal{F}_{mem} .* Biological forgetting is often a structural collapse of meaning, where the world does not just fade into black but turns disturbingly wrong. To simulate the phenomenology of *Semantic Dementia* [Warrington 1975] and *Associative Agnosia* [Farah 2004], we model the machine memory as a high-dimensional vector field (as introduced in Sec 3.1.2), where the semantic identity of an object is defined mainly by the direction of its **feature vector** z [Wang and Isola 2020], and introduce a *geometric perturbation* within the high-dimensional memory space.

The Distortion: We model *cognitive decline* acts as a force that *twists* the memory manifold. Using a continuous orthogonal rotation $R(\lambda)$, we misalign semantic vectors while preserving their norm:

$$R(\lambda) = (1 - \lambda) \cdot I + \lambda \cdot Q, \quad (1)$$

where \mathbf{Q} is a random orthogonal matrix generated via QR decomposition [Mezzadri 2007], \mathbf{I} is the identity matrix. As the decay factor λ increases, the semantic feature vectors \mathbf{z} are smoothly rotated away from their ground truth alignment.

$$\mathbf{z}_{\text{rot}} = \mathbf{z} \cdot \mathbf{R}(\lambda)^T. \quad (2)$$

This simulates a brain state where neural firing rates (signal strength) remain high, but the semantic connectivity (signal direction) is twisted – suffering from “cognitive confusion”. The machine continues to reconstruct complex, high-confidence geometries, but because the vectors are misaligned with the original priors, familiar objects are rendered as *hallucinated* topological variants. For instance, as shown in Fig. 1, a construction might be reconstructed with similar texture but the geometry of a liquid.

The Hybrid Entropy: Furthermore, to capture the full spectrum of memory decay, we combine this geometric twisting with a secondary entropy term, reflecting the irreversible loss of information:

$$\mathcal{F}_{\text{mem}}(\mathbf{z}, \lambda, \gamma) = \underbrace{(1 - \gamma) \cdot \mathbf{z}_{\text{rot}}}_{\text{Distortion}} + \underbrace{\gamma \cdot \mathcal{N}(\mathbf{z})}_{\text{Entropy}} \quad (3)$$

where $\mathbf{z}, \mathbf{z}_{\text{rot}} \in \mathbb{R}^D$ are feature vectors (here $D = 1024$), function $\mathcal{N}(\mathbf{z})$ generates a stochastic noise vector, $\mathcal{N}(\mathbf{z}) \approx \epsilon \cdot \|\mathbf{z}\|$, where $\epsilon \sim \mathcal{N}(0, I)$. By fusing manifold rotation ($1 - \gamma = 0.7$) with stochastic noise ($\gamma = 0.3$), the system visualizes a specific trajectory of forgetting: as shown in Fig. 5, memory first becomes confused (twisted geometry) before it eventually becomes lost (dissolved form). This algorithmic choice allows us to visualize what we term a *computational uncanny valley* of memory loss – a regime ($0.3 < \lambda < 0.7$) where the product of preserved structural integrity and semantic distortion peaks. We note that this “valley” is defined computationally, not through measured viewer responses; however, it is consistent with perceptual research showing that stimuli exhibiting high structural fidelity but semantic inconsistency produce maximal discomfort [Moore 2012; Seyama and Nagayama 2007] – an effect Freud [Freud 1919] described as *das Unheimliche*, the disturbing quality of the familiar-yet-wrong. This zone conceptually mirrors the phenomenology of Associative Agnosia [Farah 2004] and Semantic Dementia [Warrington 1975], where the patient perceives form without meaning.

3.3 The Selective Forgetting: Object-Oriented Amnesia

While global decay simulates a systemic cognitive decline (e.g., dementia), human forgetting is often highly selective [Warrington 1975]. We frequently lose the semantic grasp of specific entities – a face, a name, or an object – while the surrounding reality remains intact. To simulate this *dissociative amnesia*, we introduce a targeted decay mechanism that allows the system to surgically erode the memory of specific semantic categories.

3.3.1 Semantic Grounding and Masking. We employ a two-stage pipeline to isolate the *memory trace* of a specific object. First, we utilize an open-vocabulary detection model, **Grounded-SAM** [Ren et al. 2024], to generate a pixel-level binary mask $M_{\text{pixel}} \in \{0, 1\}$ based on a textual prompt (e.g., “woman in red”). This mask represents the object’s footprint in the visual field. As shown in Fig. 6, users can obtain the mask for specific entities by providing a text prompt or drawing bounding boxes. A naive approach would be to apply the mask to the input image directly (pixel-level inpainting), but this would result in a superficial “black hole” rather than a cognitive loss. Instead, we perform the intervention within the latent feature space of the memory.

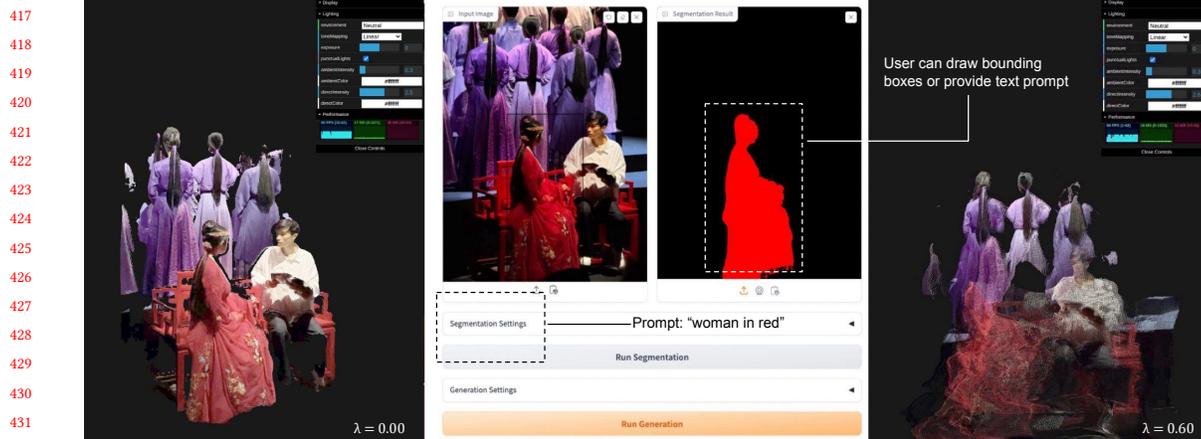


Fig. 6. **Visualizing Selective Forgetting.** Targeted entropy injection using a text prompt (“woman in red”). **(Left)** Original reconstruction. **(Right)** Decayed result ($\lambda = 0.6$). The targeted entity undergoes topological dissolution, effectively ghosting out of reality while the surrounding context remains intact.

3.3.2 *Latent Space Surgery.* To extract the prior about the input images, the vision foundation model [Oquab et al. 2024] first processes images as a sequence of patch tokens. To align with the spatial shape of these patch tokens, we down-sample the pixel mask M_{pixel} to create a discretized token-wise mask $M_{\text{token}} \in \{0, 1\}$.

We then modify the forward pass of our system: we perform decay function Eq. (3) only into the tokens corresponding to the target object:

$$\mathbf{z}' = M_{\text{token}} \cdot \mathcal{F}_{\text{mem}}(\mathbf{z}, \lambda, \gamma) + (1 - M_{\text{token}}) \cdot \mathbf{z}, \quad (4)$$

where \mathbf{z} is the original clear memory.

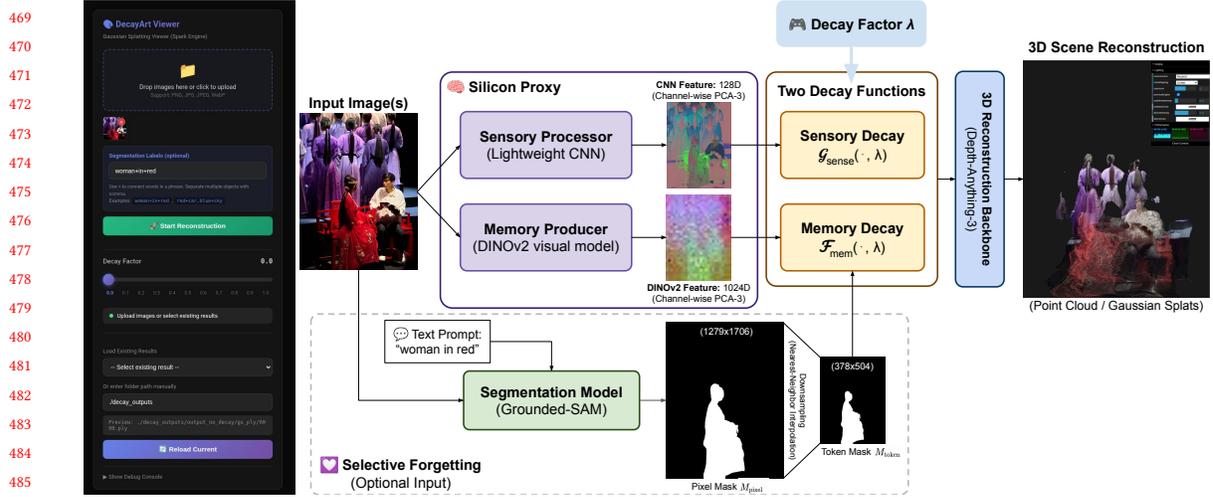
By perturbing only the latent tokens of the target, we achieve a form of *selective agnosia*. As shown in Fig. 6 on the right, the reconstructed 3D point cloud of the surrounding environment (e.g., other actors) remains structurally sound. However, the targeted entity (e.g., “woman in red”) begins to exhibit a topological failure: its geometry melts, its texture turns into noise, and it effectively ghosts out of physical existence. This visualizes the psychological state where a patient may perceive the context of a room perfectly but fails to resolve the identity or form of a specific object within it – a computational emulation of the brain’s refusal to process a traumatic memory. We give more discussion in Sec 6.2.

4 Real-time Installation and Implementation Details

4.1 User Interface and Integration Formula

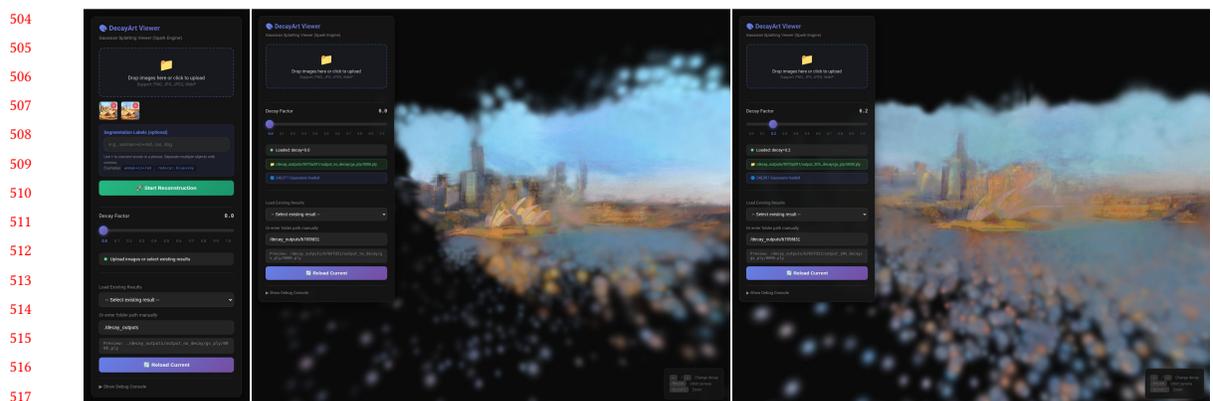
We designed *Echoes of the Prior* as a real-time interactive installation. The core of the installation is a closed-loop reconstruction system that processes visual stimuli into decaying 3D memories.

As shown in Fig. 7, the system accepts visual input (either from a live camera feed capturing the viewer’s environment or pre-loaded image dataset) and processes it through a pre-trained 3D reconstruction backbone, *Depth Anything 3* [Lin et al. 2025], for memory construction. Unlike standard inference, the user-defined decay factor λ **intercepts** this reconstruction process by two parallel decay functions: **The Sensory Decay Path (Top)**, which acts as a spectral filter (see Sec. 3.2.1); **The Memory Decay Path (Bottom)**, which optionally integrates a **Selective Forgetting** module. This module utilizes *Grounded-SAM* [Ren et al. 2024] to generate semantic masks based on user prompts. Entropy

Fig. 7. The System Workflow of *Echoes of the Prior*.

is **surgically injected** only into the tokens representing specific objects, disrupting their structural integrity while preserving the context. Finally, these corrupted latents from both paths are fused together and then decoded by the 3D reconstruction backbone into a fractured 3D scene represented by Gaussian Splats or point cloud, which is rendered in real-time via a WebGL-based viewer to provide immediate visual feedback of the decay. The system supports 60FPS interaction: as the user slides the fader, the world melts instantly, creating a visceral connection between their action and the visual destruction.

Designed with minimalist semantics, the interface centers on a single slider ($\lambda \in [0, 1]$). This transforms an abstract algorithm into a proprioceptive experience: as the user physically drags the fader, they witness the immediate entropic dissolution of the scene. Unlike static video loops, the rendered scene supports full orbit controls. Users can rotate and zoom into the decaying structures, inspecting the glitch artifacts up close.

Fig. 8. **The Interactive Interface.** The web-based GUI enables real-time interaction with the decay process.

4.2 Image Acquisition Strategy

The system supports dual inputs: a **Curated Archive** of semantic archetypes for immediate engagement, and a **Personal Injection** feature allowing users to upload personal photos or video, transforming the installation into a digital elegy.

While the paper figures illustrate the core algorithm, we developed an advanced **Interactive Streaming Interface** (shown in the demo video, a representative frame shown in Fig. 9) to facilitate real-time experimentation. This system integrates Grounded-SAM for prompt-based masking and implements a batch pre-computation pipeline, allowing users to seamlessly visualize the continuous decay process in 4D space without inference latency.

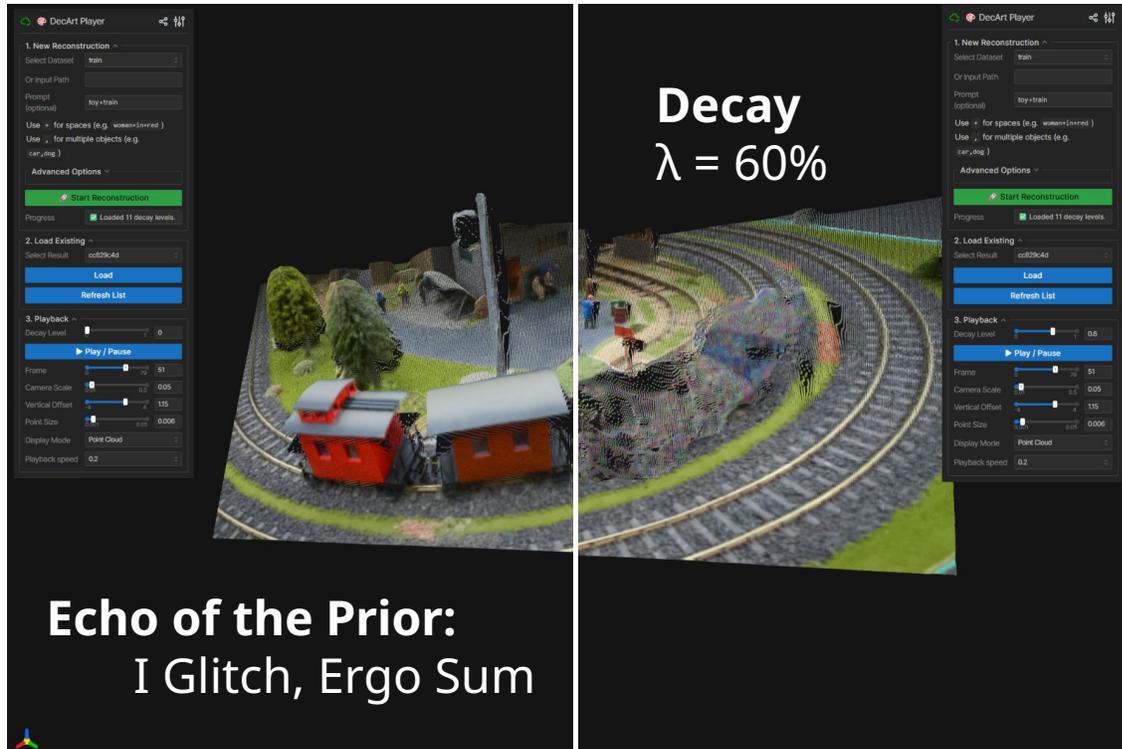


Fig. 9. “I Glitch, Ergo Sum” (I glitch, therefore I am). A visual manifesto generated by our real-time streaming system *DecArt*.

4.3 The Experience Journey

We describe the intended audience encounter to articulate the experiential arc of the installation.

Encounter and Orientation. The viewer enters a darkened space and approaches a large display showing a photorealistic 3D scene, with a tablet serving as the control console. At $\lambda = 0$, the scene appears pristine – a faithful reconstruction that can be orbited freely. The scene may be drawn from a *Curated Archive* or, in a more intimate *Bring Your Own Memory* mode, from the viewer’s own personal video uploaded via AirDrop or QR code.

Two Modes of Interaction. The viewer controls the decay through two complementary modes. In *Global Decay*, a slider governs λ directly: as it increases, edges soften, textures hallucinate, and geometry warps – progressing from subtle shimmer ($\lambda \approx 0.1$) through an unsettling intermediate zone of topological distortion ($0.3 < \lambda < 0.7$) to total

semantic dissolution ($\lambda > 0.8$). The slider is bidirectional, letting the viewer retreat from oblivion at any point. In *Selective Forgetting*, the viewer names or taps a specific element (e.g., “the woman in red”); the system segments the target via Grounded-SAM [Ren et al. 2024] and restricts the entropy injection to its latent tokens alone. The result is a world perfectly remembered except for one conspicuous absence – a ghost-shaped void. When operating on a personal video, selecting and dissolving a specific figure becomes a deeply personal gesture.

Sonic Dimension. A generative soundscape co-evolves with the visual decay: spectral filtering strips harmonics, granular synthesis introduces stochastic micro-events, and amplitude attenuates toward silence – all parametrically coupled to λ , creating a coherent multimodal experience of entropic dissolution.

Aesthetic Character: The Deep Glitch. The resulting visual language – which we term *semantic liquefaction* – is distinct from both traditional Glitch Art (broken data) and generative AI aesthetics (fabricated coherence). Because the decay operates on cognitive representations rather than pixels, objects retain their texture while their geometry undergoes topological metamorphosis: a building does not pixelate but *melts*, retaining just enough identity to be recognized as a ruin of itself. The volumetric medium further enriches this effect – orbiting the scene reveals that corruption is non-uniform, with some vantage points exposing catastrophic voids behind near-coherent facades, giving the work a sculptural rather than pictorial character.

4.4 Negative Justification: The Ontology of the Shell

In the development of *Echoes of the Prior*, we evaluated two distinct paradigms for 3D scene reconstruction: (1) generative multi-view diffusion, represented by MIDI-3D [Huang et al. 2025], and (2) monocular depth reconstruction, represented by Depth Anything V3 (DA3) [Lin et al. 2025].

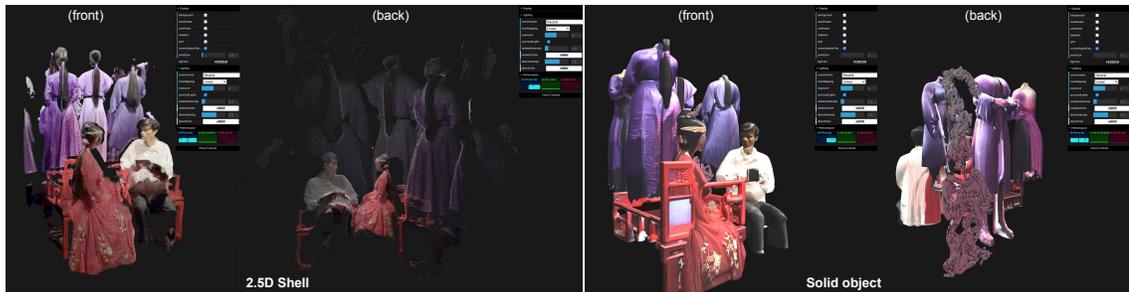


Fig. 10. **Memory as a Potemkin Village: Hallucination vs. Reconstruction.** While 3D generative models (**Right**) fabricate a complete, watertight object by inventing unseen details, our system (**Left**) preserves the epistemological honesty of the input. It projects a thin, fragile film, like a *facade* of memory, that disintegrates when the viewer attempts to look behind the veil. This hollowness is not a technical failure, but a deliberate aesthetic feature representing the void of forgotten context.

While generative models offer watertight meshes, as shown in Fig. 10, we find such completeness antithetical to memory. Memory is a surface, not a solid [Merleau-Ponty 1945]. We therefore choose DA3’s monocular reconstruction to preserve the *Potemkin village effect*: a fragile film that reveals its void upon rotation, metaphorically representing the emptiness of forgotten contexts.

5 Discussion

5.1 Design Insights: Glitcho, Ergo Sum

The homophone *DecArt* invites a reflection on *Cartesian doubt* in the age of silicon. Descartes famously stripped away all sensory perception as potentially deceptive, arriving at the only undeniable truth: *Cogito, Ergo Sum* (I think, therefore I am). Our installation poses a parallel question for the machine: when the sensory weights rot and the semantic priors twist, does the “ghost in the shell” still exist? Or, in our case: *Glitcho, Ergo Sum* (I glitch, therefore I am).

This suggests a profound irony: a perfectly optimized neural network is merely a sterile photocopier – accurate, yet alien. The “humanity” of the machine emerges not in its computation, but in its poetic failure. When the AI struggles to reconstruct reality from corrupted memory, it ceases to be a tool and becomes a mirror, forcing the audience to confront the inevitability of their own cognitive entropy. During early presentations of the work, viewers were invited to interact with the system by adjusting the decay parameter and observing the progressive dissolution of the reconstructed scene. Viewers spent extended periods exploring intermediate decay levels of the system, where objects remained partially recognizable yet semantically distorted. A recurring reaction – especially among viewers familiar with computer vision or 3D reconstruction – was surprise at this reframing: 3D reconstruction is typically evaluated purely in terms of fidelity, whereas the intentional and structured degradation of reconstruction quality was perceived as opening an unexpected aesthetic and conceptual design space. These informal observations suggest that the installation can prompt audiences to reinterpret reconstruction models not only as tools for accurate recovery, but also as media for exploring perceptual and cognitive fragility. These observations are qualitative and are not intended as a formal user study. Ultimately, we argue that generative AI can be more than an engine for productivity; it can be a medium for *radical empathy*.

5.2 Ethical Dimensions: On Aestheticizing Neurodegeneration

Any artistic engagement with neurodegeneration must confront an ethical tension: the risk that aestheticizing cognitive decline reduces lived suffering to visual spectacle [Sontag 2003]. Cultural representations of dementia – as “living death” or “loss of self” – can reinforce stigma and strip dignity from those living with these conditions [Zeilig 2013]. We acknowledge this tension directly. Our intent is not to aestheticize suffering but to cultivate *empathetic understanding* – to let the viewer briefly inhabit a perceptual world that is structurally disorienting, fostering compassion rather than voyeurism. In this, we align with person-centered frameworks [Kitwood 1997] that emphasize the retained capacities and subjective experience of people with dementia, rather than framing the condition solely as loss [Kontos 2005]. We also acknowledge that the current project was developed without direct engagement with communities affected by neurodegeneration. Future iterations should involve collaboration with patients, caregivers, and clinicians to ensure the work serves as a bridge toward understanding rather than an appropriation of experience [Basting 2009].

Early conversations around the work also extended beyond the immediate visual experience. Viewers from humanities and literary backgrounds often interpreted the installation as a provocation about the limits of machine subjectivity. If contemporary generative systems can simulate the appearance of memory, confusion, or loss without possessing first-person sensation, in what sense can such outputs be understood as artistic expression rather than secondary descriptions of human experience? Rather than resolving this question, the project intentionally leaves it open, positioning the installation as a space for reflection on the relationship between artificial generation, perception, and authorship.

6 Limitations and Future Work

6.1 The Boundaries of the Silicon Proxy

We wish to be explicit about the epistemological status of our work. The mappings we draw between computational components and biological processes – CNN weights to V1 simple cells, DINOv2 features to memory engrams – are *productive artistic analogies*, not claims of mechanistic equivalence. Deep neural networks can predict neural responses in the visual cortex [Yamins and DiCarlo 2016], but this correspondence is statistical rather than mechanistic: architecturally diverse models achieve similar predictive performance [Storrs et al. 2021], and DNNs diverge from biological vision in fundamental ways [Bowers et al. 2023; Lindsay 2021]. The mapping between any specific architecture and the brain remains underdetermined [Jonas and Kording 2017].

Our system is therefore best understood as a *speculative computational phenomenology* [Varela 1996]: it uses the structural analogy between neural network pathology and cognitive decay as an artistic lens [Cichy and Kaiser 2019], aligned with the tradition of using computational systems as “metapictures” [Offert and Bell 2021; Zylinska 2020] rather than literal scientific models. Moreover, while we successfully simulate the visual phenomenology of memory loss (the *look* of forgetting), we cannot simulate the affective sensation (the *feeling* of sadness). The machine does not miss the objects it forgets; it simply processes noise. This highlights the *Hard Problem of Consciousness* [Chalmers 1995]. We do not claim to model how the brain actually degenerates; rather, we ask: *if* a machine’s cognition were organized analogously to ours, *what would its forgetting look like?* Future work could integrate *Affective Computing* modules to modulate the decay based on a simulated *emotional state* of the AI.

6.2 Artifacts in Selective Forgetting: The Impossibility of Clean Forgetting

Our current implementation of object-aware forgetting relies on mapping pixel-level segmentation masks to the lower-resolution latent space of the memory. The necessary down-sampling and interpolation of the binary masks result in imperfect alignment at object boundaries, leading to a halo effect where the decay noise spills over into adjacent pixels. While technically a precision error, this limitation may align with the project’s conceptual framework. It visualizes the impossibility of isolating a single memory trace without affecting its connected reality – the impossibility of *clean forgetting*. Just as a fading memory often blurs the details of its spatial context, our installation’s inability to perform a clean cut reflects the entangled nature of visual perception.

6.3 Future Work: Towards a Romantic Neuromorphic Art

While standard Artificial Neural Networks (ANNs) constitute the backbone of current AGI, they represent a static “snapshot” of thought. To simulate *living* memory, we aim to transition to Spiking Neural Networks (SNNs) [Maass 1997]. Though currently confined to neuromorphic robotics, we repurpose this event-driven technology to visualize the *silencing of time*, where forgetting is not just noise, but the cessation of firing.

References

- Memo Akten, Rebecca Fiebrink, and Mick Grierson. 2019. Learning to see: you are what you see. In *ACM SIGGRAPH 2019 Art Gallery*. 1–6.
- Refik Anadol. 2021. Machine Hallucinations. Exhibition at Pilevneli Gallery, Istanbul.
- Laurie Anderson and Hsin-Chien Huang. 2017. Chalkroom. Virtual reality installation. Available at <https://laurieanderson.com/?portfolio=chalkroom>.
- Anne Davis Basting. 2009. *Forget Memory: Creating Better Lives for People with Dementia*. Johns Hopkins University Press.
- Jeffrey S. Bowers, Gaurav Malhotra, Marin Dujmović, Milton Llera Montero, Christian Tsvetkov, Valerio Biscione, Guillermo Puebla, Federico Adolfi, John E. Hummel, Rachel F. Heaton, and et al. 2023. Deep problems with neural network models of human vision. *Behavioral and Brain Sciences* 46 (2023), e385.

- 729 David J Chalmers. 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2, 3 (1995), 200–219.
- 730 Radoslaw Martin Cichy and Daniel Kaiser. 2019. Deep neural networks as scientific models. *Trends in Cognitive Sciences* 23, 4 (2019), 305–317.
- 731 Philip R Corlett, Guillermo Horga, Paul C Fletcher, Ben Alderson-Day, Katharina Schmack, and Albert R Powers. 2019. Hallucinations and strong priors. *Trends in cognitive sciences* 23, 2 (2019), 114–127.
- 732 Charles R. Crutch, Sebastian J. and Harrison, Emilie V. Brotherhood, Paul M. Camic, Brian Day, and Anthony J. Woods. 2019. Created Out of Mind: Shaping Perceptions of Dementia Through Art and Science. In *Mobile Brain-Body Imaging and the Neuroscience of Art, Innovation and Creativity*. Springer International Publishing, 135–143.
- 733 Sebastian J Crutch, Ron Isaacs, and Martin N Rossor. 2001. Some workmen can blame their tools: artistic change in an individual with Alzheimer’s disease. *The Lancet* 357, 9274 (07 2001), 2129–2133.
- 734 Russell L De Valois, Duane G Albrecht, and Lisa G Thorell. 1982. Spatial frequency selectivity of cells in macaque visual cortex. *Vision research* 22, 5 (1982), 545–559.
- 735 C.H Espinel. 1996. de Kooning’s late colours and forms: dementia, creativity, and the healing power of art. *The Lancet* 347, 9008 (1996), 1096–1098.
- 736 Martha J Farah. 2004. *Visual Agnosia*. MIT Press.
- 737 Sigmund Freud. 1919. The Uncanny. In *Imago*.
- 738 Karl Friston. 2010. The free-energy principle: a unified brain theory? *Nature reviews neuroscience* 11, 2 (2010), 127–138.
- 739 Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*. 2414–2423.
- 740 John J Hopfield. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences* 79, 8 (1982), 2554–2558.
- 741 Zehuan Huang, Yuan-Chen Guo, Xingqiao An, Yunhan Yang, Yangguang Li, Zi-Xin Zou, Ding Liang, Xihui Liu, Yan-Pei Cao, and Lu Sheng. 2025. Midi: Multi-instance diffusion for single image to 3d scene generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 23646–23657.
- 742 David H Hubel and Torsten N Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of Physiology* 160, 1 (1962), 106–154.
- 743 Eric Jonas and Konrad Paul Kording. 2017. Could a neuroscientist understand a microprocessor? *PLoS Computational Biology* 13, 1 (2017), 1–24.
- 744 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–14. SIGGRAPH 2023.
- 745 Tom Kitwood. 1997. *Dementia Reconsidered: The Person Comes First*. Open University Press.
- 746 Pia Kontos. 2005. Embodied selfhood in Alzheimer’s disease: Rethinking person-centred care. *Dementia: The International Journal of Social Research and Practice* 4 (2005), 553–570.
- 747 Shu-Chen Li, Ulman Lindenberger, and Sverker Sikström. 2001. Aging cognition: from neuromodulation to representation. *Trends in cognitive sciences* 5, 11 (2001), 479–486.
- 748 Haotong Lin, Sili Chen, Junhao Liew, Donny Y. Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. 2025. Depth Anything 3: Recovering the Visual Space from Any Views. arXiv:2511.10647 [cs.CV]
- 749 Grace W. Lindsay. 2021. Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of Cognitive Neuroscience* 33, 10 (2021), 2017–2031.
- 750 Wolfgang Maass. 1997. Networks of spiking neurons: The third generation of neural network models. *Neural networks* 10, 9 (1997), 1659–1671.
- 751 Rosa Menkman. 2011. *The glitch moment(um)*. Institute of Network Cultures.
- 752 Maurice Merleau-Ponty. 1945. *Phenomenology of Perception*. Gallimard.
- 753 Francesco Mezzadri. 2007. How to generate random matrices from the classical compact groups. *Notices of the AMS* 54, 5 (2007), 592–604.
- 754 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision (ECCV)*. 405–421.
- 755 Roger K. Moore. 2012. A Bayesian explanation of the ‘Uncanny Valley’ effect and related psychological phenomena. *Scientific Reports* 2 (2012), 864.
- 756 Alexander Mordvintsev, Christopher Olah, and Mike Tyka. 2015. Inceptionism: Going deeper into neural networks. *Google Research Blog* (2015).
- 757 Fabian Offert and Peter Bell. 2021. Perceptual bias and technical metapictures: Critical machine vision as a humanities challenge. *AI & Society* 36 (2021), 1–12.
- 758 John O’Keefe and Jonathan Dostrovsky. 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research* 34, 1 (1971), 171–175. Nobel Prize winning work on Place Cells.
- 759 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research* (2024).
- 760 Cynthia Owsley, Robert Sekuler, and Dennis Siemsen. 1983. Contrast sensitivity throughout adulthood. *Vision Research* 23, 7 (1983), 689–699.
- 761 Albert R Powers, Christoph Mathys, and Philip R Corlett. 2017. Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. *Science* 357, 6351 (2017), 596–600.
- 762 René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*. 12179–12188.

- 781 David P Reichert, Peggy Seriès, and Amos J Storkey. 2013. Charles Bonnet Syndrome: Evidence for a Generative Model in the Cortex? *PLOS Computational*
782 *Biology* 9, 7 (2013), e1003134.
- 783 Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang,
784 Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. 2024. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks.
785 arXiv:2401.14159 [cs.CV]
- 786 Jun'ichiro Seyama and Ruth S. Nagayama. 2007. The Uncanny Valley: Effect of Realism on the Impression of Artificial Human Faces. *Presence: Teleoperators*
787 *and Virtual Environments* 16, 4 (2007), 337–351.
- 788 Michael N Shadlen and William T Newsome. 1998. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and
789 Information Coding. *Journal of Neuroscience* 18, 10 (1998), 3870–3896.
- 790 Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël
791 Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts,
792 Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. 2025. DINOv3.
793 arXiv:2508.10104 [cs.CV] <https://arxiv.org/abs/2508.10104>
- 794 Susan Sontag. 2003. *Regarding the Pain of Others*. Farrar, Straus and Giroux.
- 795 Katherine R. Storrs, Tim C. Kietzmann, Alexander Walther, Johannes Mehrer, and Nikolaus Kriegeskorte. 2021. Diverse deep neural networks all predict
796 human inferior temporal cortex well, after training and fitting. *Journal of Cognitive Neuroscience* 33, 10 (2021), 2044–2064.
- 797 Francisco J Varela. 1996. Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies* 3, 4 (1996), 330–349.
- 798 Tongzhou Wang and Phillip Isola. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In
799 *International Conference on Machine Learning (ICML)*. PMLR, 9929–9939.
- 800 Elizabeth K Warrington. 1975. The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology* 27, 4 (1975), 635–657.
- 801 Daniel LK Yamins and James J DiCarlo. 2016. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience* 19, 3 (2016),
802 356–365.
- 803 Hannah Zeilig. 2013. Dementia as a cultural metaphor. *The Gerontologist* 54, 2 (2013), 258–267.
- 804 Joanna Zylińska. 2020. *AI Art: Machine Visions and Warped Dreams*. Open Humanities Press.

805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832