

Monstrous Hallucinations as Planetary Interface: Designing Scar-Based Visualisation for AI-Driven Immersive Installations

Abstract

This paper develops aesthetics of training data through a practice-based study of Metabolic Carcass of the Dataset, a CAVE-based immersive installation where AI-generated animal-human hybrids are coupled to planetary indicators of climate and orbital damage. Rather than treating 'hallucination' as a technical flaw, the project stages monstrous outputs as planetary interfaces: composite bodies whose scars and deformations are driven by ice-sheet loss, species extinction, atmospheric CO₂ and space debris. Within the CAVE, viewers share a continuous spatial field with these data-scarred creatures, navigating policy scenarios and triggering a reverse-annotation system that briefly exposes plausible training images behind specific organs. The paper proposes scar-based visualisation as an alternative pattern for environmental communication and formalises the installation's logic as coupled mappings between archives, models, indicators and behaviours. It situates this practice within debates on data colonialism and planetary media, reconfiguring creative AI as an arena for negotiating data governance and more-than-human spectatorship.

Keywords: AI art; training data aesthetics; planetary media; immersive installation; climate communication;

My practice-based analysis of CAVE-based immersive media installation 'Metabolic Carcass of the Dataset' in which hybrid animal-human creatures roam a stark white cubic room, read as 'monstrous hallucinations,' these figures are approached not as fantastical beings but as condensed interfaces between generative models, planetary data infrastructures, and the material politics of extraction. Building on media archaeology and planetary media theory and on critical accounts of AI training data and hallucination, this paper contributes to digital creativity by operate as planetary media devices that materialize how images, sensors, and models participate in Earth-scale governance.

I propose three design moves for visual communication research. First, each monster is coupled to planetary indicators of ice-sheet loss, species extinction, greenhouse gas emissions, orbital debris so that climate data deform the creature's skin, gait, and morphology in real time, generating what I call scar-based visualization. Second, an interaction layer allows viewers to explore different policy futures (mitigation, extraction, militarization...), which reconfigure the creatures as liberated, commodified, weaponized, or reintegrated into multispecies ecologies. Third, a reverse-annotation system lets viewers point to any organ on the projected bodies and briefly reveal fragments of the training images from which such organs might plausibly have been synthesized, mapping the data politics buried in the hallucinations.

The CAVE-based environment¹ thus becomes a field laboratory for planetary media design. Sharing space, sound and experience with these scale-shifting, data-scarred bodies forces a re-thinking of species boundaries, planetary identity, and spectatorship itself. Much current AI art and immersive installation work remains formally dazzling yet conceptually evasive, treating generative systems as neutral spectacle rather than contested infrastructures stitched from extractive training datasets, asymmetric compute, and planetary-scale energy expenditure. In contrast, this paper proposes a transferable design grammar the scar-based visualization, reverse annotation, and planetary coupling that other practitioners can adapt to materialize data violence and environmental cost as palpable, negotiable features of immersive experience.

Introduction: From Generative Monsters to Planetary Media

My CAVE-based installation 'Metabolic Carcass of the Dataset'(Figure 1&2) is an immersive white-cube environment where AI-generated hybrid creatures circulate around the viewer. Projected on three walls and the floor, the monsters function as embodied hallucinations of training data and planetary damage, condensing archives of anatomical imagery, industrial husbandry, and environmental crisis into a single, enclosing scene. By forcing spectators to co-inhabit the cube with these unstable bodies, the work reconfigures the gallery into an interface for interrogating species boundaries, algorithmic vision, and planetary vulnerability. In the projected scene, spliced animal-human hybrids pace, leap, and convulse inside a minimalist white box. Some have furless canine torsos threaded with exposed bone; others sport taloned bird legs attached to muscular humanoid chests; a towering two-headed creature stares down the viewer. The creatures are neither photorealistic nor openly cartoonish. I created them using a specific coded visual regime of generative AI that slightly blurred edges, implausible joints, thin seams where training images appear to have collided. The space is at once clinical and claustrophobic. There is no 'outside' diegetic world, only the cube and its inhabitants, who are endlessly observed and yet never fully knowable. Viewers stand or sit inside the projection volume; their bodies cast shadows across the floor, occasionally intersecting with the monsters' paths. The work already reads as an allegory of machinic vision and enclosure, but it can be pushed further.

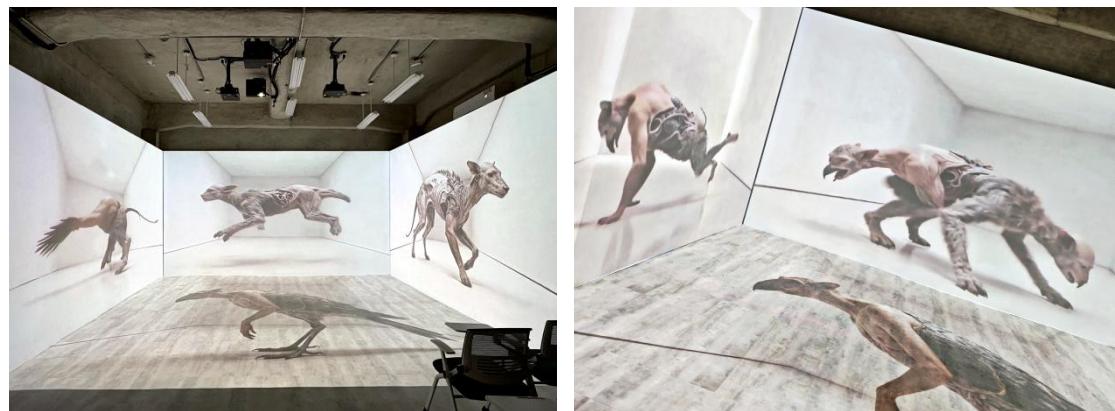


Figure 1. CAVE-based immersive installation 'Metabolic Carcass of the Dataset'.

Figure 1. Alt Text: the CAVE-based installation Metabolic Carcass of the Dataset: three rear-projected white walls and a projection floor form a seamless cubic envelope in which large, AI-generated hybrid animal-human creatures move around the viewer standing inside the space.

Figure 2. Interior view from a participant's perspective, where multiple hybrid creatures and the viewer's tracked position define a co-habited field of risk; planetary indicators continually reshape bodies and atmospheres rather than appearing as external overlays.

Figure 2. Alt Text: Interior view of the CAVE showing a closer framing of one wall and the floor, with several deformed hybrid creatures pacing and lunging in a stark white virtual room that visually doubles the physical projection space, creating a perceptual confusion about where the monsters are located.

This project treats climate and orbital indicators as forces that reorganise a shared, volumetric habitat. Planetary variables do not appear as overlays on a distant globe, but as lesions, limps, and atmospheric frictions that co-condition how bodies move and breathe in the CAVE. This shift from informational representation to co-habited risk reframes AI art and data visualisation as an epistemic-ethical practice: not simply making patterns visible, but choreographing who must live with which scars, and under what planetary assumptions. Following Jennifer Gabrys's proposal that 'becoming planetary' involves recognizing how media infrastructures compose planetary problems as calculable objects (Gabrys, 2024, 2019)² I propose turning the installation into an experiment in planetary media design. What happens if each monster is bound to streams of planetary data including cryospheric loss, biodiversity decline, CO₂ concentration, space-debris density so that the creatures' bodies become dynamic visualizations of Earth's injuries? What if the same monsters are also made to disclose the hidden training images that underwrite their hallucinated anatomies? The resulting environment would treat generative AI not as an opaque author of wonders but as a planetary apparatus whose outputs can be philosophically and sensorially dissected. By combining climate data, training-set forensics, and multi-species avatars inside the CAVE environment, the work aims at an aesthetics of training data, which is an approach that makes the normally invisible substrate of generative media in which its archives, extractions, and omissions available to experiential critique.

Monstrous Hallucinations and the Politics of Training Data

The term 'hallucination' has quickly become shorthand for the ways generative models produce convincing but fabricated outputs. Recent research traces how the metaphor frames error as quasi-pathological property of models, often deflecting attention from structural issues in data sourcing and system design (Förster, 2025)³. Bender et al. (2021)⁴ insist that large language models are essentially 'stochastic parrots,' remixing

patterns in pre-existing corpora without understanding, and warn that their scale amplifies the biases and exclusions encoded in their training sets. For visual generative systems, the situation is analogous. Kate Crawford's *Atlas of AI*⁵ documents how computer-vision datasets are entangled with extractive mining, exploitative labeling labor, and imperial forms of surveillance (Crawford, 2021). Couldry and Mejias (2019)⁶ theorize such pervasive data capture as 'data colonialism,' a reorganization of social life around continuous extraction (Couldry & Mejias, 2019). Manovich and Arielli's *Artificial Aesthetics* (2024)⁷ argue that generative imagery reconfigures visual culture by embedding aesthetic decisions in training-set composition and model architecture, demanding new literacies of data disclosure and algorithmic authorship.

Against this backdrop, the monsters in the installation can be read as embodiments of training data politics. Their 'impossible' anatomies with too many joints, asymmetrical musculature, scars that seem to be both organic and metallic index the heterogeneity of archival fragments from which such bodies are assembled that veterinary photographs, medical illustrations, bodybuilding stock images, animal-rights campaigns, industrial meat-packing footage, videogame textures. The fact that the creatures are trapped in an empty, surveilled white cube evokes both the laboratory and the dataset, which is a space where bodies are standardized, recorded, and circulated without regard for their singular histories. To treat these monsters as hallucinations is therefore to foreground not their unreality but their over-determination. Each limb points back to multiple possible training images, each gait traces paths through an archive of animal experimentation, war photography, or fitness advertising. The hallucination is not a random glitch but a symptom of the underlying social and ecological relations that made the dataset possible. Aesthetically, the monsters are 'too much' thereby they condense dystopian industrial relations, planetary violence, and the seductive smoothness of machine-generated imagery into a single, uncomfortable figure.

Aesthetics of training data takes this excess seriously. Instead of correcting hallucinations to restore accurate representation, it proposes to stage them as critical surfaces where the data relations of AI become sensuously legible. As an artistic medium, the CAVE immersive environment offers rare capacity for co-presence of bodies, images, and data share a continuous spatial field, allowing algorithmic processes to be felt as ambience, pressure, and proximity rather than merely seen on a screen. At the same time, this aesthetics risks becoming a new formalism fixated on ever more sophisticated ways of rendering hallucinations visible while leaving the material regimes of extraction and labor that sustain them largely untouched. By remaining primarily within curatorial and discursive circuits, it may inadvertently aestheticize harm turning data colonialism, ecological devastation, and cognitive exhaustion into objects of refined contemplation rather than fronts of organized struggle. Its task, therefore, is structurally limited unless it forges durable alliances with political economy, decolonial theory, and climate justice movements that can

translate sensuous legibility into constraints on how training archives are built, governed, and eventually dismantled.

Planetary Media and Scar-Based Visualization

Planetary media theory provides conceptual frame for this experiment. Gabrys (2018, 2019, 2022)⁸ describes planetary media as infrastructural assemblages of sensor networks, satellites, models, dashboards through which the Earth is rendered as knowable, governable entity. Lee-Morrison et al. (2023)⁹ similarly examine 'machinic visions of the planetary,' emphasizing how satellite imaging, climate modeling, and AI reshape what counts as planetary evidence and how such evidence circulates (Lee-Morrison et al., 2023). Kim (2025)¹⁰ in a comparative study of European and Chinese digital-twin programs, analyzes digital twins as planetary media that couple high-resolution models of territories to optimization regimes in science and policy.

Much of this work focuses on operational images such as maps, dashboards, simulation outputs that are designed to be read by experts and machines more than by lay publics (Parikka, 2015¹¹; Marks, 2020). The 'Metabolic Carcass of the Dataset' monsters propose a different mode of planetary mediation. Instead of depicting the planet as a globe, grid, or heat map, they act as multi-species avatars whose bodies register planetary data as lesions, tumors, limps, and deformations. Each creature is algorithmically linked to a small bundle of indicators. Technically, the creatures are generated with a latent diffusion model rather than a GAN because diffusion architectures offer greater control over conditional sampling, better stability under heavy fine-tuning, and higher tolerance for extreme morphological variation. I trained a custom checkpoint on a stratified corpus of anatomical illustration, veterinary documentation, industrial husbandry, and speculative creature design, using low learning rates (1e-5-5e-5), 200-400 epochs, and classifier-free guidance (scale 6-9). Inference parameters (denoising steps, guidance scale, seed variability) were iteratively tuned to maximise anatomical plausibility while preserving visible artefacts of the training corpus.



Figure 3. Cryosphere monster



Figure 4. Extinction monster

Figure 3. Alt text: Rendered still of the 'cryosphere monster,' a pale, quadrupedal hybrid with translucent crystalline growths along its back; the crystals appear cracked and slumping, while the creature's hind legs strain to balance, visually encoding accelerated ice-sheet melt and sea-level rise.

Figure 4. Alt text: Rendered still of the 'extinction monster,' a furred, partly humanoid creature whose coat thins into bare patches, eyes turn cloudy, and ghostly phantom limbs flicker at its sides, representing species loss drawn from biodiversity and Living Planet Index data.

Cryosphere monster(Figure 3): tied to time-series of ice-sheet mass balance and sea-level rise. As melt accelerates, translucent crystals on its back crack and slump, its hind legs struggle to stabilize on the floor.

Extinction monster(Figure 4): mapped to the Living Planet Index and species-richness projections; its fur thins, eyes cloud over, and phantom limbs flicker as unobserved species disappear from databases.



Figure 5. Carbon monster



Figure 6. Orbital debris monster

Figure 5. Alt text: Rendered still of the 'carbon monster,' a muscular hybrid torso veined with dark, pulsating lines that branch across the skin; within the veins, faint silhouettes of pipelines and flare stacks intermittently appear, linking bodily deformation to rising atmospheric CO₂ and fossil-fuel infrastructure.

Figure 6. Alt text: Rendered still of the 'orbital debris monster,' a bulky creature whose flank is pierced by sharp metallic fragments that either orbit around or protrude from its body, visualizing the accumulation of trackable and untrackable space debris as wounds in a more-than-human anatomy.

Carbon monster(Figure 5): coupled to atmospheric CO₂ parts per million and fossil-fuel infrastructure; dark veins pulse across its torso as emissions climb, occasionally revealing the silhouettes of pipelines and flare stacks.

Orbital debris monster(Figure 6): bound to estimates of trackable and untrackable space junk; as orbital congestion increases, shrapnel-like fragments orbit, embed in, and perforate the avatar's body, rendering space-junk statistics as persistent, injurious proximity.

Rather than smoothing these data into legible graphs, the system translates them into scars. Changes in the indicators modulate texture maps, skeletal rigs, and animation parameters, producing a slow choreography of injury. I call this logic scar-based visualization. Scar-based visualization rejects the fantasy of disembodied overview associated with planetary dashboards. It does not show the planet as a calm object of knowledge but as a wounded, twitching composite whose injuries are unevenly distributed. It resonates with what Rob Nixon (2011)¹² calls 'slow violence' incremental, attritional harms that are difficult to represent and with Gabrys's insistence that planetary sensing should account for more-than-human vulnerabilities rather than purely technocratic metrics. By binding abstract data streams to grotesque yet affectively charged bodies, the 'Metabolic Carcass of the Dataset' acts as a planetary interface that lets viewers feel, not just know, how climate and orbital infrastructures are transforming the conditions of life for human and nonhuman beings. The monsters' scars are not metaphorical; they are algorithmic consequences of actual measurements, albeit stylized and scaled for experiential impact.

The scar-based visualization pipeline operates as follows. Planetary time-series are ingested, normalized, and mapped onto animation parameters and procedural textures for each creature, which bind climate data to visible bodily degradation rather than colour maps. Some mappings are direct, e.g., higher CO₂ increases the amplitude of thoracic expansion and constriction, while others are temporally offset to emphasize long-term lag, e.g., extinction metrics influence the slow fading of certain bodily features. Above 420 ppm, vessel-texture opacity is reduced by ~30% while a 0.5 Hz pulsation of normal intensity is activated, producing a faint, anxious throbbing under the skin. Extinction and pollution indices drive lesion density and colour shift. Creatively, scars are treated as a live interface where abstract climate data returns as damaged, metabolically overdriven flesh. Viewers can select different policy scenarios via a discreet tablet interface or physical sliders such as aggressive mitigation, business-as-usual extraction, militarized orbital expansion, strong multispecies governance. Each design interface modifies the projected evolution of the indicators across the duration of the visit, driving alternative trajectories of bodily transformation. Under extractionist futures, the monsters grow more weaponized or commodified; under care-oriented futures, scars may heal into new forms of symbiosis rather than simply vanish.

Sound Design Using Logical Operations of Monstrous Hallucinations

I present this formalization as a design notation that other practitioners can adapt when scripting multimodal hallucinations in immersive environments. My practice articulate the agential logic in explicit philosophical-mathematical form and then treat

this formalization as a compositional score, using it to script the sound design as a direct sonic enactment of those operations. To clarify how hallucinated bodies and planetary scars operate within the project, I schematize the installation as a set of coupled mappings between archives, models, planetary indicators, and embodied avatars.

Let D be the heterogeneous training archive, decomposed into subsets

$$D = D_{med} \cup D_{vet} \cup D_{ind} \cup D_{mil} \cup D_{ent} \cup \dots$$

(medical, veterinary, industrial husbandry, military, entertainment, etc.).

M be the generative model with parameters θ , such that a hallucinated image h is given by

$$h = H(z; D, \theta)$$

where z is a latent seed and H is the hallucination operator.

Each monster-body is segmented into regions

$$B = \{b1, b2, \dots, bn\}$$

(limbs, torsos, heads, organs, prostheses).

We can then define a data provenance relation

$$R \subseteq D \times B,$$

where $(d, bi) \in R$ reads: 'element d in the archive contributes to the synthesis of body region bi .' In practice, we only have access to a plausible relation R^\wedge , reconstructed through reverse-annotation in the installation, but we treat R as a theoretical ground.

The monstrous hallucination of a given creature c can then be written as a composite:

$$c = \sum_{i=1}^n \phi_i (H(z_i; D, \theta)),$$

Where ϕ_i is a spatial mask that extracts and positions the relevant fragment of hallucinated imagery for body region bi . Over-determination appears where multiple archival strata converge on the same region:

$$\forall b_i \exists d_j, d_k \in D: (d_j, b_i) \in R \wedge (d_k, b_i) \in R \wedge d_j \neq d_k.$$

Philosophically, the 'hallucination' is thus not a simple error but an over-saturated mapping from incompatible archival lineages into a single, unstable corporeal form.

For the planetary coupling, let P be a vector of planetary indicators:

$$P(t) = (p_{\text{cry}}(t), p_{\text{ext}}(t), p_{\text{CO}_2}(t), p_{\text{orb}}(t), \dots)$$

where each component is a time-dependent series (cryosphere loss, extinction indices, atmospheric CO₂, orbital debris density, etc.).

E_c be the set of parameters that define the expressive state of creature c : skeletal angles, gait frequencies, texture fields, translucency, and lesion maps.

$$F: P(t) \times E_c^{(0)} \rightarrow E_c(t),$$

Scar-based visualization can then be formalized as a modulation function

$$F: P(t) \times E_c^{(0)} \rightarrow E_c(t),$$

such that

$$E_c(t) = F(P(t), E_c^{(0)}),$$

Where $E_c^{(0)}$ is a baseline configuration and F encodes the choreography of injury (e.g., higher p_{CO_2} intensifies vascular darkening; higher p_{ext} increases phantom-limb flicker). A scar, in this sense, is a derived field

$$S_c(t) = \sigma(P(t)),$$

with σ mapping planetary indicators into localized deformations at the level of skin, bone, and motion.

The planetary interface property of the installation can be captured by the following condition:

$$\forall t \forall v \in V: (v \text{co-inhabits the CAVE with } c(t)) \rightarrow \text{Perceive}(v, S_c(t)),$$

where V is the set of viewers and $\text{Perceive}(v, S_c(t))$ denotes the phenomenological exposure of viewers v to the embodied planetary scars at time t . The installation is planetary not because it shows a globe, but because the mapping $P \mapsto S_c$ organizes how planetary processes become sensible in shared space.

To extend this logic operation into sound, I treat audio not as secondary ornament, but as an additional channel in the same compositional calculus. Let

$X_v(t)$ denote the state of viewers v (position, orientation, physiological arousal).

Ω be the set of sound sources and filters distributed in the ‘Metabolic Carcass of the Dataset’ (speakers, transducers, spatialization buses).

$A_c(t)$ be the auditory state of creature c : timbral envelopes, rhythmic patterns, spatial trajectories.

I define a sound-mapping function

$$G: P(t) \times E_c(t) \times X_v(t) \rightarrow A_c(t),$$

so that sonic behaviour is co-determined by planetary indicators, corporeal deformation, and viewers presence. Concretely:

Planetary drones

$$f_{\text{base}}(t) = f_0 + \alpha_{\text{CO}_2} p_{\text{CO}_2}(t) + \alpha_{\text{orb}} p_{\text{orb}}(t),$$

where $f_{\text{base}}(t)$ controls a subsonic or low-frequency drone diffused through floor transducers. As CO_2 and orbital debris increase, the spectrum thickens and destabilizes, making the planetary metabolic load haptically audible.

Extinction clicks and gaps

Let λ_{ext} be a rate parameter derived from extinction indicators. A stochastic process

$$N_{\text{ext}}(t) \sim \text{Poisson}(\lambda_{\text{ext}}(t))$$

governs high-frequency clicks and micro-ruptures in the sound field. As more species vanish from the index, the pattern paradoxically grows more sparse and irregular, encoding absence as temporal silence.

Scar sonification

Each localized scar field in $S_c(t)$ is mapped to a micro-resonance:

$$\forall s_k \in S_c(t): \quad a_k(t) = \beta_k \cdot \|s_k(t)\|,$$

Where $a_k(t)$ modulates the amplitude or filter depth of a narrow band in the spectrum. The deeper the scar, the more pronounced the corresponding resonance or distortion. Viewers moving closer to a wounded region experience intensified, localized sonic abrasion.

Data-provenance whispers

For each relation $(d, b_i) \in \widehat{R}_i$, I assign a spectral 'voice' ν_d . The composite vocal texture around body region b_i is:

$$V_{b_i}(t) = \sum_{(d, b_i) \in \widehat{R}} w_d \nu_d(t),$$

where w_d weights each archive subset (e.g., medical, industrial, activist). These voices do not resolve into intelligible language; instead, they form shifting phonetic smears, half-buried traces of the archives metabolized into the hallucination.

In posthuman terms, this sound design refuses anthropocentric voice-music hierarchies. The creatures do not 'speak for' the planet in human language; rather, the planet's metrics, archives, and injuries are folded into a distributed sonic metabolism that envelops both monsters and viewers. The 'Metabolic Carcass of the Dataset' becomes an audible logic machine in which each transformation $P(t) \rightarrow S_c(t), R \rightarrow V_{b_i}(t), X_v(t) \rightarrow A_c(t)$ renders otherwise opaque relations perceptible as pressure, vibration, and interference. Listening within this environment is an encounter with the algebra of training data and planetary media, written directly into the auditory field of a shared, more-than-human body.

The White Box That Compresses Data-Driven Deformation

The scene of white box doubling of the physical CAVE by its virtual twin produces a powerful effect (Figure 7). The viewer inhabits a room that appears identical to the room they see onscreen, creating a subtle confusion about where the monsters are 'really' located. In contemporary art theory, the white cube remains an unresolved problem because it operates simultaneously as technology of purification and machine of disavowal. As an spatial device, it claims neutrality erasing context, labor, and infrastructure while in practice stabilizing specific regimes of visibility, authorship, and value (O'Doherty, 1999)¹³. The most urgent question is how this putative emptiness now intersects with algorithmic and planetary scales. White cubes are not only ideological containers for modernist autonomy, but terminal nodes in extractive media ecologies that extend from rare-earth mining to global real-estate speculation (Bennett, 2015; Parikka, 2015)¹⁴. What forms of exhibition-making could expose, rather than occult, these supply chains and climatic entanglements? And can the white cube be reprogrammed as critical interface of 'operating theater' for environmental and data politics without simply aestheticizing crisis?

Interpreted media-archaeologically, the white box evokes the modernist white cube of the gallery, the psychology experiment chamber, the slaughterhouse inspection corridor, and the clean room of semiconductor fabrication. Parikka(2015)¹⁵'s call to trace the deep media histories embedded in such spaces of rock, dust, energy infrastructures helps situate the 'Metabolic Carcass of the Dataset' not as an abstract digital environment but as an extension of planetary resource chains. By filling this white cube with planetaryly coupled monsters, the installation compresses dispersed eco-political processes into a single room. Ice-sheet dynamics, forest dieback, agribusiness, fossil extraction, and orbital capitalism are all made to inhabit one claustrophobic enclosure. The 'Metabolic Carcass of the Dataset' turns into a small-scale planetary prison, a place where the consequences of global systems are confined for inspection, but also where the viewer realizes that they, too, are inside the apparatus.

From a design-research perspective, this system is less about predictive accuracy than about epistemic experimentation. It allows us to investigate how people perceive and emotionally respond to planetary models when those models are incarnated in strange bodies rather than in graphs. Do viewers treat the monsters as victims(Figure 8) whose suffering should be minimized, as tools to be optimized, or as uncanny kin with whom they share a planetary fate?



Figure 7. The CAVE as 'white box' apparatus.

Figure 7. Alt text: Schematized diagram of the 'scar-based visualization' pipeline, depicting arrows between training archives, generative model, planetary indicator time-series, and segmented monster bodies; planetary data streams are mapped to parameters such as texture, translucency, and gait to produce evolving scars and deformations.

Figure 8. A single body in which heterogeneous archives, species, and imaging regimes are fused into a co-present, volumetric creature.

Figure 8. Alt text: Rendered still of a single hybrid creature whose body fuses multiple anatomical fragments, species traits, and imaging regimes into one volumetric form: patches of fur, bare skin, translucent tissue, and scan-like overlays coexist on the same torso, suggesting that heterogeneous

archives and more-than-human histories have been composited into one co-present, data-scarred body.

For the moment, I bracket the question of whether the CAVE should itself be understood as a residual, leftover form of new media art. Yet the installation as an immersive environment is itself a compromised apparatus of tightly controlled, rectilinear chamber that reproduces the modernist fantasy of total visual enclosure even as it seeks to critique it. Its volumetric field is optically dominant and anthropocentric, privileging the upright human viewer and marginalizing other sensory and more-than-human modes of experience of thermal, olfactory, microbial that exceed projection-based vision. The reliance on high-resolution projection, heavy computation, and air-conditioned infrastructure also risks reproducing the very extractive energy regimes and planetary metabolisms that the work seeks to problematize. Epistemically, the CAVE environments can too easily stabilize complex data and contested archives into coherent spectacles, foreclosing the ambiguities, silences, and counter-knowledges that should remain politically negotiable. Unresolved, therefore, are questions of how such environments might be decolonized, multispeciesed, and energetically re-engineered and how a future CAVE environments could function not as a sealed theater of synthetic control, but as a porous, low-carbon, and collectively governed interface for genuinely planetary forms of mediation and encounter.

Reverse Annotation, Species Boundaries and Planetary Identity

The second experimental layer concerns training-set forensics. Whereas climate indicators feed into the creatures' evolution, another subsystem invites viewers to probe the archival origins of specific body parts. Handheld pointers or gaze-tracking allow a viewers to select, say, the junction where a bird leg meets a humanoid pelvis. For a few seconds, the wall behind the creature dissolves into a dense collage of images such as X-rays of fractured hips, racing-greyhound photographs, Renaissance anatomical drawings, industrial poultry photographs, biomechanical concept art. The collage does not claim to show the exact dataset but an algorithmically constructed plausible ensemble drawn from open repositories tagged with relevant keywords.

Each such revelation constitutes a micro-investigation of data politics. Viewers may notice that most reference images come from Western medical archives, military sources, or industrial farming in the Global North corroborating critiques of data colonialism and uneven visibility (Couldry & Mejias, 2019; Crawford, 2021). In other cases, they may see scraped selfies or animal-rights campaign imagery, raising questions about consent, activism, and appropriation. The point is not to reconstruct any proprietary training set with forensic precision but to dramatize the epistemic gap between output and archive. A monstrous hallucination becomes a portal onto overlapping image economies in which scientific, commercial, activist whose entanglement is usually obscured by polished aesthetics of generative media (Gillespie, 2024¹⁶; Khokhar, 2025). Through repeated interactions, viewers begin to

grasp that every body part is contested site where multiple visual regimes have been sampled, flattened, and recombined. The creatures are no longer merely victims of planetary damage; they are cross-sections through a stratified archive of colonial science, factory farming, fitness culture, and speculative design.

If the installation is a planetary prison, viewers are not outside its jurisdiction. They share the cube with the monsters; their shadows fall across the projected floor; the creatures occasionally appear to lunge toward them. This co-presence matters. It shifts the work from a spectacle about distant planetary harms to situation in which the human spectator is implicated in a shared condition. The installation thus becomes a laboratory for reimagining species boundaries and planetary identity. The monsters are neither straightforwardly nonhuman nor fully anthropomorphic. They are in between among cyborgs in Donna Haraway's sense, but also planetary avatars that condense multi-species suffering and resilience (Haraway, 2016)¹⁷. Their very illegibility disrupts the neat division between 'human subject' and 'environmental object.'

From a phenomenological perspective, viewers experience a form of forced conviviality with hybrid others. They cannot simply glance at an image and move on; they must inhabit the same air, listen to the same echoing footsteps, and endure the same enclosure. This proximity invites what Susan Schuppli (2020)¹⁸ calls 'material witnessing,' an attunement to how evidence is embodied in material forms. Here, the evidence is not photograph or legal document but a monstrous gait whose limp signifies decades of data accumulation. This situation also recodes planetary identity. In mainstream discourses, the planet often appears as a distant blue-marble object or a crisis dashboard quantified into stripes and graphs. In the 'Metabolic Carcass of the Dataset', by contrast, the planetary is proximal and embodied, yet never fully human. Viewers might feel themselves addressed not as sovereign individuals but as partial fragments of a larger damaged organism distributed body whose limbs are ice sheets, coral reefs, and orbital shells of debris.

The monstrous hallucinations thereby offer a speculative answer to Gabrys's question of how to 'sense a planet in crisis' when perception is mediated by infrastructures and models (Gabrys, 2019)¹⁹. Instead of aiming for total overview, they cultivate what might be called partial monstering that is experience of being simultaneously inside and outside the planetary body, complicit in its injuries yet capable of re-imagining its futures.

Toward an Aesthetics of Training Data

Aesthetics of training data critically treats machine-learning datasets not as invisible substrates of computation but as culturally and sensorially charged artefacts. It foregrounds how decisions about what is collected, labeled, excluded, and cleaned are already aesthetic operations that shape the look, feel, and intelligibility of algorithmic outputs. Rather than focusing on the beauty or novelty of generated images, this

perspective relocates aesthetic inquiry to the level of the archive containing the composition of corpora, their scalar organization, their histories of colonial extraction, labor, and environmental cost, and the ways these histories surface as glitches, biases, or hallucinated forms. Aesthetics of training data thus investigates how datasets encode norms of race, species, gender, class, and planetarity, and how artistic and design practices can expose, reconfigure, or refuse these encodings. It's the study of the perceptual, affective, and ethical forces condensed in training sets, and of the artistic strategies that make those forces publicly thinkable. However, it also exposes blind spot in much AI art that fixates on the strangeness of outputs while leaving the longue durée of infrastructural violence untheorized; by contrast, Calculating Empires makes clear that any aesthetics of training data worth the name must take empire, extraction, and logistics as its primary materials. My project takes this as provocation to push the genealogy one step further into the flesh, if Crawford and Joler diagram the empire that trains the model, Metabolic Carcass of the Dataset asks what it means to inhabit, sensorially and ethically, the mutilated bodies that this empire leaves in its wake.

Ethical critique here must move beyond individual bias and consent to interrogate how training corpora crystallize corporate power, geopolitical asymmetries, and planetary extraction. Karen Hao's (2025)²⁰ account of OpenAI as an 'empire' of infrastructural control shows that any aesthetic engagement with training data must also be a critique of enclosure, secrecy, and the privatization of planetary-scale archives. Olli Hellmann's (2025)²¹ study of Sora's colonial visual tropes around Aotearoa New Zealand demonstrates that training sets silently reinscribe imperial historiographies, making decolonial audit and reparative data curation central, not optional, to aesthetic judgment. Bartlett's (2025)²² Decentring Ethics further suggests that AI art can operate as method of staging situations in which those most affected by extractive data regimes co-produce criteria for refusal, redistribution, and counter-archiving.

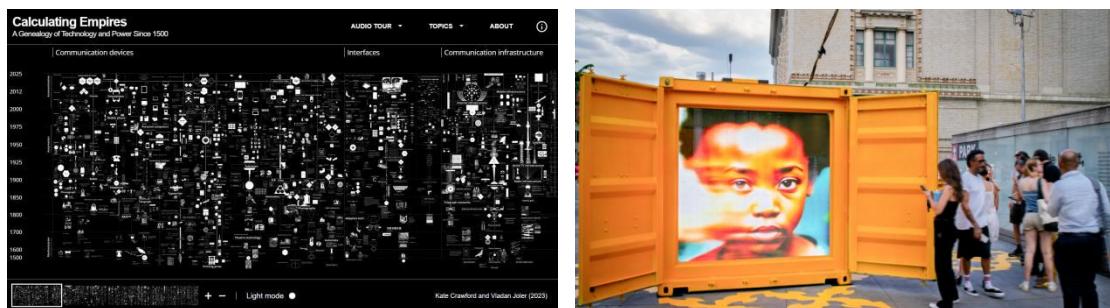


Figure 9. Calculating Empires (Kate Crawford and Vladan Joler, 2023-2025), a large-scale research diagram that stages five centuries of technology-power entanglements as dense cartography. Source: Kate Crawford and Vladan Joler.

Figure 9. Alt text: Installation view of Calculating Empires (Kate Crawford and Vladan Joler, 2023-2025), a long, wall-sized research diagram composed of dense black-and-white graphical

strata that map five centuries of entanglements between communication, computation, classification, and control, turning technological and imperial infrastructures into a continuous, walkable cartography of power.

Figure 10. If We Don't, Who Will? (Stephanie Dinkins, commissioned by More Art, The Plaza at 300 Ashland, Downtown Brooklyn Cultural District, 2025), public AI laboratory that translates community-contributed stories into live generative imagery. Source: Downtown Brooklyn / More Art.

Figure 10. Alt text: Installation view of If We Don't, Who Will? at The Plaza at 300 Ashland in the Downtown Brooklyn Cultural District, featuring a multi-screen urban media façade with animated, text-and-image-based visualizations foregrounding climate justice and community organizing in the midst of passing pedestrians and city lights.

A sharper precedent is Kate Crawford and Vladan Joler's *Calculating Empires: A Genealogy of Power and Technology, 1500-2025* (2023-2025)²³. Staged as a 24-metre diptych, the work maps five centuries of 'communication, computation, classification, and control,' diagramming how colonial archives, biometric registries, military sensing, and corporate data infrastructures coalesce into preconditions for contemporary AI. (Figure 9) Rather than showing individual training images, *Calculating Empires* operates at the scale of infrastructural genealogy visualises the long circuits of extraction, enclosure, and logistical coordination that silently underwrite today's training datasets. In this sense it prototypes aesthetics of training data that is less about spectacular outputs than about historical composition of the corpus itself treating datasets as the latest iteration of imperial calculation, not as neutral 'big data.' By forcing viewers to walk along a dense cartography of wires, ships, cables, categories, and camps, the piece frames every AI model as a node in much older empire of classification. It thereby opens crucial horizon for my project, if *Calculating Empires* charts the macroscopic genealogy of training regimes, 'Metabolic Carcass of the Dataset' collapses those regimes into a single hallucinated body, where the same histories of selection and control appear as wounds, grafts, and scars.

Stephanie Dinkins' If We Don't, Who Will?²⁴ transforms shipping container into public 'data classroom,' (Figure 10) where Black and brown communities feed stories, images, and vernaculars into a live AI system, turning training data into site of reparative practice and community governance. The aesthetics of training data must therefore remain suspicious of works that merely 'visualize' datasets while leaving extractive pipelines and institutional capture untouched. The danger is new genre of critical spectacle in which bias, opacity, and colonial archives are theatrically acknowledged yet materially unchanged, precisely what Crawford (2025)²⁵ diagnoses as the 'metabolic logic' of AI slop continuous ingestion and excretion without structural repair. The 'AI slop' Crawford(2024)²⁶ indicates is not a minor glitch but the popular surface of computational capitalism, where human culture is continuously rendered into a standing reserve of data.

AI slop is not contingent by-product of generative systems but predictable aesthetic of a media regime organized around continuous ingestion and recomposition of data. Building on Flusser's 'technical images'²⁷ and Farocki's 'operational images,'²⁸ , culture is treated as a standing reserve in which images, texts, and sounds are scraped, decomposed, and refashioned into synthetic outputs that circulate as if they were spontaneous creativity rather than the sediment of vast extractive operations. What appears on the surface as playful remix or 'magic' automation is, at the infrastructural level a tightly coupled metabolism linking data centers, energy grids, logistics chains, and water systems, fabricated by ingesting and decomposing billions of prior images through infrastructures with immense energy and water demands.

Yet this configuration is intrinsically unstable. The metabolic rift between human and machinic systems manifests on at least three fronts. First, model collapse emerges as models are trained increasingly on their own synthetic outputs, degrading diversity and amplifying artefacts in feedback loop of self-referential noise. Second, ecological collapse is accelerated by the resource intensity of training and deployment. Escalating demands for electricity, land, rare minerals, and cooling water are layered on top of existing extractive economies and climate stressors. Third, cognitive collapse looms as synthetic media saturate everyday perception, undermining epistemic trust, overwhelming attention and normalizing environments where verification becomes prohibitively costly. The present AI economy does not simply risk malfunction; it tends toward self-undermining conditions in which the systems that generate slop erode the very ecological and cognitive bases on which they depend.

The crucial point is that these are not external 'side effects' but consequences of the dominant metabolic logic itself. A regime premised on infinite scalability, frictionless appropriation of cultural material, and ever-expanding demand for computational power cannot sustain itself indefinitely on a finite planet. The pertinent question is therefore not whether current AI slop economies will exhaust their conditions of possibility, but how, and on whose backs, that exhaustion will be borne. If left unchallenged, the most likely trajectories involve uneven forms of managed collapse. Zones in which synthetic media remain abundant for affluent users while extractive burdens and environmental degradation are offloaded onto marginalized human and more-than-human communities. At the same time, this instability opens conceptual and practical space for what can be called post-synthetic cultural organization. Post-synthetic does not imply the disappearance of generative techniques, but a shift in metabolic logics under which they operate. Instead of treating training data as a free, inexhaustible input, post-synthetic forms would recognize limits, obligations, and reciprocal entanglements. They would be grounded in alternative energy systems that prioritize decarbonization and ecological repair; governance arrangements that redistribute ownership and control over data and models to communities rather than corporations; and cultural protocols that specify not only what may be ingested by algorithms, but under what conditions, for whose benefit, and with what mechanisms of refusal, redress, and restitution. What is needed are practices that reroute

infrastructural control redirecting value, authorship, and even model parameters toward those historically treated as training material rather than epistemic partners. Practices that accept the irreversibility of computational mediation while refusing its current hyper-metabolic form is what we need. This might include dataset commons governed by those historically rendered as mere training material; low-energy, small-scale models designed for situated use rather than planetary capture; or artistic and design experiments that slow down, annotate, and counter-map slop, transforming it into a site of collective inquiry rather than passive consumption.

Ethics adequate to aesthetics of training data must be structural, decolonial, and planetary in scope, willing to question not only how models look, but who owns the images they learn from, whose worlds they overwrite, and what forms of shared governance could re-route these archives toward more just and multispecies futures. These questions are not rhetorical but jurisdictional. 'Ownership' of images in massive, scraped corpora is less a matter of individual rights than of who is authorized to define the scope of the visible and the sayable. When models overwrite worlds such as Indigenous ecologies, minoritized lifeways, nonhuman habitats, they do so by normalizing one regime of description at planetary scale, relegating others to noise or anomaly. Shared governance would therefore require more than advisory boards or opt-out form which would mean community and multispecies stakeholders co-determining what may be collected, how long it may persist, under what obligations of return, and with what capacities for refusal. In such a regime, training data become a contested commons rather than a frictionless resource.

We may also question what kinds of desires and attachments are we willing to cultivate? We can acquiesce to AI slop as the default horizon of mediated experience, normalizing a culture of endless synthetic reproduction aligned with extraction and exhaustion. Or we can orient aesthetic, institutional, and infrastructural work toward other futures where cultural forms that treat images, texts, and datasets as vehicles for flourishing rather than raw material to be burned. Media ecologies that privilege regeneration over throughput and planetary interfaces that acknowledge their own metabolic costs and obligations. Aesthetics of training data reframes hallucination from a technical pathologization to a critical resource. Hallucinations are not merely bugs to be minimized but sites where the systemic conditions of data collection, labeling, and curation become visible. By turning hallucinated monsters into large-scale, explorable entities, the 'Metabolic Carcass of the Dataset' makes legible the normally invisible architectures of datasets, from their colonial genealogies to their sampling biases.

Moreover, aesthetics of training data foregrounds the non-innocence of archives. The reverse-annotation sequences in the installation expose training images as contested artifacts, saturated with histories of exploitation, activism, and desire. This resonates with media-archaeological calls to examine the strata of past media objects that undergird today's platforms (Parikka, 2012, 2015; Gabrys, 2011²⁹). Yet the monsters

push beyond textual or diagrammatic analysis by allowing viewers to feel how different archival regimes literally pull the body in distinct directions.

The approach also entangles data politics with planetary governance. By coupling creatures to climate and orbital indicators, the installation reveals how decisions about energy, agriculture, and space infrastructure reverberate through the bodies of more-than-human beings. The scar-based visualization pattern treats planetary models not as abstract projections but as ongoing injuries that accumulate in shared flesh. In doing so, it complements emerging work on planetary digital twins that warns against technocratic fantasies of optimization and control (Kim, 2025)³⁰ by insisting on the affective and ethical consequences of such modeling. This aesthetics reimagines human-AI relations as co-embodiment rather than interface alone. Viewers do not simply click buttons or prompt a model; they share a room with its hallucinated outputs, negotiate futures with them, and experience their own reflections in distorted form. This resonates with critical discussions of AI as extractive ecological arrangement rather than a discrete tool (Crawford, 2021; Marks, 2020³¹). The 'Metabolic Carcass of the Dataset' shows how such arrangements might be reworked into spaces of reflexive encounter, where violences of data colonialism and environmental extraction are not abstractly described but aesthetically confronted.

Finally, aesthetics of training data suggests new methodologies for art and design research in which planetary media theory and practice form a feedback loop.. Instead of documenting an existing system, the installation is itself a speculative instrument that generates empirical material including viewers trajectories, choices among policy scenarios, verbal reflections, and affective responses. These can be analyzed alongside the theoretical frameworks invoked here, creating a loop between planetary media theory and practice. The 'Metabolic Carcass of the Dataset' becomes both philosophical argument and experimental setup, embodying what media-environment scholars call 'critical technical practice' at planetary scale (Lee-Morrison et al., 2023). Designing with and against infrastructures simultaneously to repurposing sensing, modeling, and generative systems while keeping their material and geopolitical conditions permanently in view, rather than treating them as neutral tools. Computation is obligated to disclose its scars, share governance, and accept limits rather than quietly optimize toward exhaustion. It demands that technical choices of data schemas, model architectures, interface grammars be treated as sites of contestation where climate justice, decolonial repair, and multispecies care are negotiated in code and hardware as much as in discourse.

System Architecture and Creative Pipeline

The installation is built as a tightly coupled pipeline linking training datasets, generative models, planetary indicators, and a CAVE-based rendering environment. The aim is not simply to 'skin' a pre-existing avatar, but to let planetary signals and training-data provenance co-determine the morphology, surface condition, and comportment of each monstrous body.

Datasets and curation

Four partially overlapping image corpora underpin the system: (1) anatomical and veterinary atlases documenting musculature, vasculature, wounds and scars across multiple species; (2) industrial and medical imaging of fractures, corrosion, and impact damage; (3) remote-sensing and scientific visualisations of ice-sheet fractures, smoke plumes, dust storms, and orbital debris fields; (4) a smaller corpus of in-house sketches and 3D renders of hybrid creatures used to stabilise the target aesthetic. All material is licensed, de-identified where relevant, and standardised to common resolution and colour space. CLIP-based tagging pass produces semantic embeddings that later support reverse-annotation without exposing raw images in the installation.

I fine-tune a latent diffusion model rather than training from scratch, prioritising controllability and compatibility with downstream real-time rendering. A publicly available checkpoint trained on general imagery is adapted using low-rank adaptation (LoRA) modules and a curated subset of the corpora emphasising trauma, hybridity, and environmental textures. This preserves the base model's compositional flexibility while biasing it toward 'monstrous interface' regime. Training uses mixed-precision optimisation on commodity GPUs, with early stopping determined by qualitative inspection of samples and quantitative monitoring of overfitting via reconstruction loss on a held-out validation set. From hallucinated images to rigged bodies. The diffusion model is not tasked with producing fully formed 3D creatures, but with generating high-resolution 'epidermal atlases', in which multi-channel image sets containing albedo, normal, and displacement information for key body regions (torso, limbs, head). These atlases are projected onto base meshes sculpted and retopologised in Blender from a library of human, animal, and speculative anatomies. Retargetable skeletons are created using standard rigging tools, allowing gait and posture to be driven independently from surface condition. This division of labour keeps the 3D pipeline robust and makes the contribution of training-data aesthetics legible at the level of skin, scars, and local deformations.

Planetary indicators and mapping

Planetary variables such as atmospheric CO₂ concentration, global mean temperature anomaly, cryosphere mass balance, biodiversity indices, and orbital debris density are pre-processed into normalised time series within [0,1], using z-score normalisation and exponential smoothing to avoid jitter. During the installation, viewers select among scenario bundles that correspond to different plausible futures. Each variable is mapped to a set of visual and behavioural parameters, for example, CO₂ level modulates the opacity and emissive intensity of vascular textures and their pulsation frequency; ice-sheet loss drives the amplitude of mesh fragmentation and the instability of the creature's gait; debris density controls the density and tangibility of orbiting shrapnel fields that occasionally intersect the avatar. These mappings are implemented as shader uniforms and animation controllers updated at frame rate,

ensuring that planetary change is experienced as a continuous deformation rather than a set of discrete 'states.'

Reverse annotation and data legibility

To materialise the training dataset without directly exhibiting copyrighted or sensitive material, we compute CLIP embeddings for all curated images offline and store only their vector representations and textual descriptors. In the CAVE, when a viewer targets a body region with a tracked pointer, the system retrieves k-nearest neighbours in embedding space and generates a brief overlay of synthetic micro-montage or textual capsule ('veterinary archive: bovine fracture; satellite: Larsen C rift, 2016') that appears adjacent to the organ for a few seconds. This mechanism renders data provenance intermittently visible while maintaining ethical and legal constraints.

Audience experience and iterative development

During public showings, up to four viewers inhabit the CAVE simultaneously, with head and hand tracking captured at 90 Hz to drive parallax-corrected views and interaction rays. Fieldnotes and short post-experience interviews indicate that participants most viscerally registered planetary change when gait instability and skin degradation accumulated slowly over several minutes rather than as immediate effects, a finding that informed subsequent tuning of temporal smoothing and parameter ranges. Earlier desktop and single-wall prototypes exposed problems of information overload and 'dashboard fatigue,' leading to reduce textual overlays and emphasise bodily cues. Future iterations will port the core pipeline to a WebGPU/WebXR stack for browser-based access and to standalone VR headsets, preserving scar-based mappings while exploring lighter-weight interaction schemes suitable for domestic or classroom environments.

Conclusion

For digital artists, interaction designers, and creative coders, this project offers a transferable vocabulary for reconfiguring AI systems as critical interfaces rather than neutral tools. Scar-based visualisation turns climate and orbital indicators into lived, deforming conditions on an avatar's body, displacing dashboards with situated, affective co-habitation. Reverse annotation folds the opacity of training data back into the scene, making provenance and data violence intermittently perceptible. By forcing viewers to co-inhabit the cube with these creatures, it invites a re-imagining of species boundaries and planetary identity as matters of shared vulnerability rather than distance. Planetary coupling, finally, frames each technical choice of dataset, model, mapping as a normative decision about which worlds are rendered survivable, inviting future immersive works to treat hallucination as design space for negotiating accountability.

The aesthetics of training data developed here thus points toward a broader agenda for planetary media design. Instead of striving for ever more seamless, invisible infrastructures, designers and artists can cultivate spaces where the scars of planetary computation, its extractions, erasures, and hallucinations are made experientially salient. Such spaces will not deliver simple solutions; they will be messy, uncomfortable, and at times overwhelming. Yet precisely in this discomfort lies their political and philosophical potential to unsettle the common sense that AI is immaterial, that climate is elsewhere, and that the planet is a distant, smooth blue orb rather than a wounded body to which we irreducibly belong.

Conflict of interest

In accordance with Taylor & Francis policy and my ethical obligations as a researcher, I confirm that I have no financial or business interests, consultancy roles, or external funding relationships that could be affected by the research reported in the enclosed paper. I have no conflicts of interest to declare.

Funding

No funding was received for this research.

Reference:

¹ In the current prototype, the CAVE consists of three rear-projected walls (4.5 m × 2.6 m each) and a projection floor (4.5 m × 4.5 m), driven by five 10,000-lumen laser projectors at 1920×1200 px, 60 Hz, edge-blended to produce a continuous 5760×2400 px image envelope. Spatial tracking is provided by two overhead depth cameras (time-of-flight, 30 fps, 640×480 px) supply coarse body position and orientation for sound and image modulation; no biometric data are stored.

² 'The Forest Multiple: Composing and Digitalizing Wooded Worlds,' special issue edited by Jennifer Gabrys, Michelle Westerlaken, Kate Lewis Hood, Yuti Ariani Fatimah and Trishant Simlai. *Environment and Planning F: Philosophy, Theory, Models, Methods and Practice*.

³ Förster, S., Skop, Y. Between fact and fairy: tracing the hallucination metaphor in AI discourse. *AI & Soc* (2025). <https://doi.org/10.1007/s00146-025-02392-w>

⁴ Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. <https://doi.org/10.1145/3442188.3445922>

⁵ Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

⁶ Couldry, N., & Mejias, U. A. (2019). Data colonialism: Rethinking big data's relation to the contemporary subject. *Television & New Media*, 20(4), 336-349. <https://doi.org/10.1177/1527476418796632>

⁷ Manovich, L., & Arielli, E. (2024). Artificial aesthetics: Generative AI, art and visual media. <https://manovich.net/index.php/projects/artificial-aesthetics>

⁸ Citizens of Worlds: Open-Air Toolkits for Environmental Struggle. Minneapolis: University of Minnesota Press, 2022.

⁹ Lee-Morrison, Lila, Kathrin Maurer, Rikke Munck Petersen, and Dominique Routhier. 2023. 'Introduction: Machinic Visions of the Planetary.' *Media+Environment* 5 (1). <https://doi.org/10.1525/001c.88425>.

¹⁰ Kim, J. (2025). Digital twins as planetary media: Optimization, resilience, and the modeling of territory. *New Media & Society*. Advance online publication.

¹¹ Parikka, J. (2015). *A geology of media*. University of Minnesota Press. Retrieved from <https://culturetechnologypolitics.wordpress.com/wp-content/uploads/2015/11/jussi-parikka-a-geology-of-media-intro.pdf>

¹² Nixon, R. (2011). *Slow violence and the environmentalism of the poor*. Harvard University Press.

¹³ O'Doherty, B. (1999). *Inside the white cube: The ideology of the gallery space*. University of California Press.

¹⁴ Bennett, J. (2015). Systems and things: On vital materialism and object-oriented philosophy. In R. Grusin (Ed.), *The nonhuman turn* (pp. 223-239). University of Minnesota Press.

¹⁵ Parikka, J. (2012). *What is media archaeology?* Polity Press.

¹⁶ Gillespie, T. (2024). Generative AI and the politics of visibility. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241252131>

¹⁷ Haraway, D. J. (2016). *Staying with the trouble: Making kin in the Chthulucene*. Duke University Press.

¹⁸ Schuppli, S. (2020). *Material witnessing: Media, forensics, evidence*. MIT Press.

¹⁹ Gabrys, J. (2019). Sensing a planet in crisis. *Media+Environment*, 1(1). <https://doi.org/10.1525/001c.10036>

²⁰ Hao, K. (2025). Empire of AI: Dreams and nightmares of infrastructural control [Manuscript in preparation].

²¹ Hellmann, O. (2025). Colonial bias in AI training data: Prompting Sora to generate images of Aotearoa New Zealand's historical past. *Kōtuitui: New Zealand Journal of Social Sciences Online*.

²² Bartlett, V., Pfefferkorn, J., & Sunde, E. K. (Eds.). (2025). *Decentring ethics: AI art as method* [Edited volume, forthcoming].

²³ Crawford, K., & Joler, V. (2023). Calculating Empires: A genealogy of technology and power since 1500 [Interactive map]. <https://calculatingempires.net/>

²⁴ Downtown Brooklyn Partnership. (2025, June 26). Stephanie Dinkins' 'If we don't, who will?' unveiled at The Plaza at 300 Ashland. *Downtown Brooklyn*. <https://downtownbrooklyn.com/news/2025/if-we-dont-who-will/>

²⁵ Crawford, K. (2025). Eating the future: The metabolic logic of AI slop. *e-flux architecture*. <https://www.e-flux.com/architecture/intensification/6782975/eating-the-future-the-metabolic-logic-of-ai-slop>

²⁶ Antonelli, P. (2024). Kate Crawford: Metabolic images. *Aperture*. Retrieved from <https://archive.aperture.org/paper/2024/4/4/kate-crawford-metabolic-images>

²⁷ Flusser, V. (2011). *Into the universe of technical images* (N. A. Roth, Trans.). University of Minnesota Press. (Original work published 1985)

²⁸ Harun Farocki, *Eye/Machine I-III*, 2001-03. Video installation trilogy.

²⁹ Bartlett, V., Pfefferkorn, J., & Sunde, E. K. (Eds.). (2025). *Decentring ethics: AI art as method*. Open Humanities Press.

³⁰ Kim, J. (2025). Digital twins as planetary media in Europe and China: Optimization and resilience in scientific and socioeconomic domains. *New Media & Society*. Advance online publication. <https://doi.org/10.1177/14614448251336434>

³¹ Marks, L. U. (2020). *Radiant infrastructures: Media, environment, and cultures of uncertainty*. Duke University Press.