

Modulating Becoming: Gaze-Guided Metamorphosis and Metabolic Data Visualization in a CAVE-Based Installation

Abstract

This paper presents 'Metabolic Carcass of the Dataset', my generative art installation that uses diffusion-based AI to create hybrid bio-synthetic creatures within an immersive CAVE environment. Drawing on posthuman and media-archaeological frameworks, I situate the work amid contemporary discussions of a computational 'biomachinic'. I use Seedream 4.0 as a pretrained diffusion model accessed via an external service; rather than training the base model, I steer identity and variation through multi-reference conditioning and text-guided editing, while maintaining an evolving, project-specific reference library derived from field documentation and exhibition logs. These AI-generated creatures are spatially projected onto the CAVE walls and floor, forming an affective interface that responds to audience presence. Critically, the creatures' surfaces bear climate scars of textures and anomalies derived from live, minute-scale environmental feeds (updated via API polling) and expressed through shader-level modulation and event-triggered texture refresh. This metabolic data-mapping treats pretrained model inputs as nutrients rather than mere 'raw material', echoing recent aesthetic critiques that value emergent 'mistakes and the poetry' of machine learning over pure optimization. I detail the technical and artistic rationale for model selection, rendering pipeline, and interaction design, emphasizing novelty in our treatment of data and embodiment of planetary computation. In combining speculative theory (posthumanism, critical AI, Deleuzian) with a concrete installation, the work foregrounds the entanglement of organic and algorithmic, organism and environment. By mapping climate feedback onto generative bodies, the project makes visible the biotechnical intimacy of humans, nonhumans, and machines on a heated planet.

Keywords

Immersive media art; CAVE / spatial display; Generative AI (diffusion); Deleuzian becoming; AI alignment critique

Introduction

The Anthropocene demands new artistic modes that reveal entanglements of technology and ecology. Generative AI has recently enabled artists to imagine nonhuman entities and environments, but most work has focused on static GAN outputs or virtual creatures in isolation. My practice pushes beyond these by using diffusion-based generative modeling to craft lifelike, planet-imprinted creatures in an immersive, interactive space.^[1] These creatures inhabit a CAVE-based multi-projection environment, where participants co-presence influences their appearance.(Figure 1&2) Crucially, the installation integrates live environmental signals as a feedback loop at two timescales: minute-scale local measurements (e.g., AQI/PM2.5, temperature, humidity, wind speed) and slower 'background' indicators

(e.g., atmospheric CO₂ baselines and long-horizon sea-level proxies) updated at daily-to-weekly cadence. These inputs are mapped primarily through shader-level modulation and event-triggered texture refresh, producing 'climate scars' without implying that all planetary variables are updated minute-by-minute.



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

Figure 2 Interior view from viewer's perspective
©Respect Copyright.

I draw on media-archaeology insights that the Earth itself has become 'part of media' through sensor networks and computational infrastructures^[2]. My creatures embody this 'biomachinic' planetary condition in which their digital skins carry the histories of mining, emissions, and datafication that power our AI systems. The work resonates with feminist and posthumanist calls to acknowledge nonhuman agencies and entanglements, e.g. Haraway's cyborgs, Spivak's planetary alterity). It also aligns with emerging discourse on machine learning aesthetics rather than optimizing for photorealism, I embrace the uncanny- leveraging model failures and irregularities as poetic elements. By foregrounding these speculative and critical dimensions alongside detailed technical methods (model choice, rendering, interaction, data mapping). This paper thus articulates both the theoretical rationale and the engineering of an affective, data-driven, planet-touched generative interface. This paper also treats interaction as politics in which gaze is translated into parameters, echoing Deleuze's modulation and contemporary 'character' governance in aligned assistants such as Claude. By staging gaze-guided metamorphosis, the work makes alignment felt as pressure on bodies rather than a promise of safety.

Aesthetic Background of Biomachinic Planet

Recent scholarship examines how Machine Learning's idiosyncrasies generate new aesthetic modes. Thomas Storey notes that AI artworks often expose 'distortions in the ways algorithms interpret the world'. Recent accounts of AI art emphasize that many practitioners working with machine learning are not primarily concerned with

optimising predictive performance. Instead, they treat glitches, opaque behaviours and black-boxed processes as core aesthetic resources, and locate authorship in the curation of pretrained model corpora and the careful tuning of models rather than in individual outputs.^[3] As Grba (2024)^[4] argues, such practices articulate an 'aesthetics of machine learning' in which error, statistical abstraction and algorithmic estrangement become explicit artistic materials. I adopt this sensibility of training dataset assembled from biological specimens, landscapes, and other natural imagery is treated as a kind of metabolic fuel for a generative ecology. The data are not simply exploited and thrown away, but carefully selected, layered, and periodically 'replenished' in a cyclical process akin to feeding a living organism. These form materialises authorship as tuning and selection rather than a single output: the creature's identity is maintained as a recognisable 'type' while its surface and articulation remain open to controlled mutation. Presented in the same neutral volume, the image emphasises the work's central claim that the dataset operates as metabolic fuel periodically replenished conditioning materials that produce an ecology of bodies where the biological and the computational are structurally entangled.



Figure 3 Rather than resolving into a coherent species, the figure foregrounds how pretrained models misrecognize and recombine anatomical cues that making distortion and categorical leakage legible as aesthetic material.

Figure 4 A coiled biomachinic organism whose nacreous segments, metallic joints, and iridescent scales oscillate between marine anatomy and industrial hardware. ©Respect Copyright.

The installation builds on a lineage of interactive AI art, moving beyond static gallery pieces. Early works like Akten's Learning to See or Martino & Stornaiuolo's AI Portraits used GANs to transform real-time inputs, revealing biases of the datasets.^[5] More recent projects (e.g. Varvara & Mar's Dream) employ latent-space navigation and multimodal inputs (e.g. voice or attention proxy) to co-create imagery. My approach leverages these advances in using diffusion models Seedream provides a practical interface for iterative i2i/inpainting under latency constraints. As Guljajeva and Canet-Sola note^[6], diffusion-based inpainting and image-to-image methods now enable fine-grained control and higher resolution than earlier GANs. This means I can continuously deform and re-generate creature forms within an interactive loop, using event-triggered image-to-image edits and localized inpainting rather than relying on precomputed outputs. Head-pose tracking (6-DoF) and motion sensing in the CAVE provide embodied input streams; rather than claiming eye-tracked gaze, I use head orientation and dwell time as a view-direction proxy to infer likely attention proxy and trigger localized visual changes.

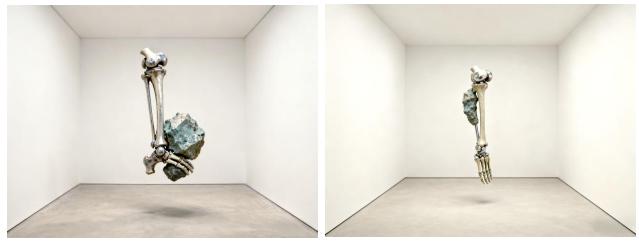


Figure 5 Rare-earth limb as biomachinic burden that bearing ore as the work's protagonist rather than a hidden substrate.

Figure 6 Rare-earth as implant and constraint.

©Respect Copyright.

I draw on critical-theoretical literature that situates my work in a larger planetary context. Thomas Storey's notion of the biomachinic planet argues that the Earth and its environment have become inextricably linked to computational systems. Satellites, sensors, and cloud data centres turn climate and ecology into programmable processes; the globe appears as a data-driven totality even as its alterity resists full capture. Storey quotes Bratton^[7] and Parikka^[8] to stress that 'planetary scale computation involves the whole Earth' the silicon chips themselves are mineral excavations, and the planet is literally 'part of media'. My installation Figure 5 materialises biomachinic planet by literalising how planetary-scale computation is anchored in extraction that the body's capacity to move is re-scripted by the mineral that makes sensors, satellites, and data centres possible. Here, the 'planet' does not merely inform the image as data; it returns as a stubborn, weighty constituent of the creature's anatomy of ethical reminder that computation begins as geology. Presented in a neutral test volume Figure 6, the image emphasises that the installation's 'scar' aesthetics are not decorative metaphors but constraints that is mineral, infrastructural, and political through which bodies become governable and legible. Bone, chrome-like couplings, and a rare-earth proxy embedded as an implant that interrupts anatomical coherence. The work thus reframes planetary limits as a kind of material intelligence that is not a background variable to be modelled, but a formative agent that dictates what forms can persist, mutate, or fail.

As showed in Figure 7, the suspended skeletal fragment is mechanically braced by a metal joint and levered against a rough silicon-bearing rock, foregrounding silicon not as a neutral 'material' but as the infrastructural condition of computation itself. By making the mineral appear as an anatomical load, something the body must literally articulate around the image concretises the biomachinic planet such as chips, sensors, and clouds begin as lithic extraction, and 'planetary-scale' media remain tethered to geological strata. A knee-like assembly in Figure 8 is encased in translucent crystalline growths that stand in for rare-earth mineral concentrates the often-invisible protagonists of high-performance magnets, sensing, and display infrastructures. The creature's joint becomes a site where geology hardens into function: rare earths appear not as background resources but as accreted constraints that shape what the body can do, turning embodiment into an index of extractive dependence and of the programmable limits through which planetary

computation is made operational.

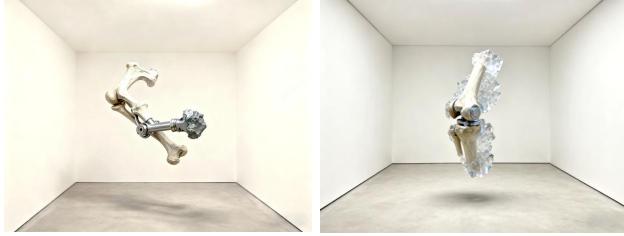


Figure 7 Silicon limb: the substrate of planetary computation.

Figure 8 Rare-earth accretion: crystalized dependency and control. ©Respect Copyright.

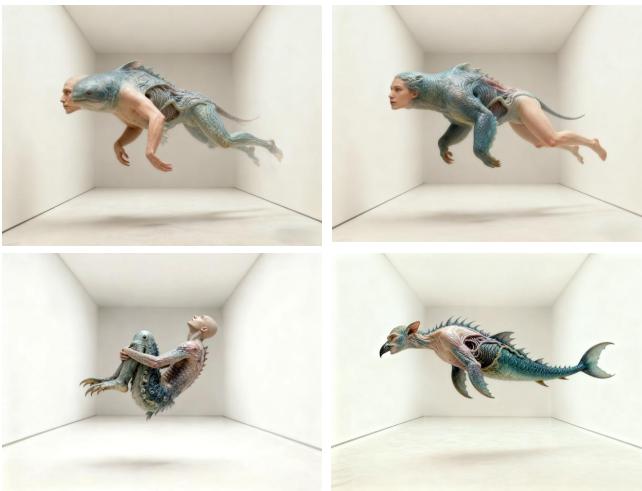


Figure 9 Gaze-guided metamorphosis I: aquatic-human composite

Figure 10 Gaze-guided metamorphosis II: identity drift within a stable type

Figure 11 Biomachinic chimera: categorical leakage as aesthetic material

Figure 12 Counter-becoming and capture: affective entanglement. ©Respect Copyright.

Figure 9 creature-type in which a human torso and arms are fused to an aquatic body, staged in a neutral test volume to foreground modulation rather than spectacle. The hybrid's skin reads as a data-bearing surface with environmental inputs bias pigment, wetness, and scar-like textural noise, while interaction triggers (head-pose dwell and proximity thresholds) prompt localized edits that recompose facial and dermal features. The image operationalises 'planetary limits as intelligence' by letting nonhuman constraints author the body's sensible form. The same creature-type Figure 10 showing controlled variation across the head/face and epidermal detailing. Here, authorship is located less in a single output than in maintaining a recognisable morphological 'family' while permitting drift through rate-limited, event-driven refreshes. The work thus makes planetary computation concrete that embodiment becomes a negotiated consequence of constraints, what can change, how fast, and under which environmental and attentional conditions. Figure 11 avian-marine composite whose exposed, rib-like interior and scaled surfaces dramatise the model's recombination of incompatible anatomical cues. Rather than 'fixing' such distortions, the installation treats them

as a legible trace of how computational systems parse the world with an aesthetics of estrangement that mirrors the biomachinic planet's entanglement of ecology and computation. Human figure in Figure 12 entwines with a creature-body, visualising the installation's central tension between co-creation and governability. Interaction is staged as both intimacy and infrastructure in which gaze/attention proxy becomes a measurable signal that modulates the creature, while the scene also invites refusal, hesitation, and opacity moments where the loop can be interrupted rather than naturalised as seamless participation. Taken together, biomachinic planet not as an abstract thesis but as an embodied condition of bodies become interfaces where environmental signals and computational capture co-produce sensation. In this assemblage, 'planetary' does not mean total visibility; it names a regime in which measurement, modulation, and extraction are sutured to flesh while opacity and refusal remain fragile counter-forces that contest that suturing.

This echoes projects like Tega Brain's Solar Protocol^[9], which injects solar data into network algorithms so that energy availability dictates digital flows (Figure 9 & 10). In contrast to top-down climate models, Solar Protocol treats environmental dynamics as an intelligence in its own right. Earth's linkage to computation is rendered palpable precisely at the point where the living body is compelled to negotiate what the system can sense, formalise, and govern. Why not think of planetary limits as intelligence? Similarly, I allow climate data to drive the aesthetics of our creatures, rather than merely letting human-centered design dominate. Accordingly, the installation frames these data-driven mutations not as 'information displays' but as a discipline of constraint-led co-authorship, in which the environment sets the operational conditions of the image and thereby redistributes agency away from the human designer.

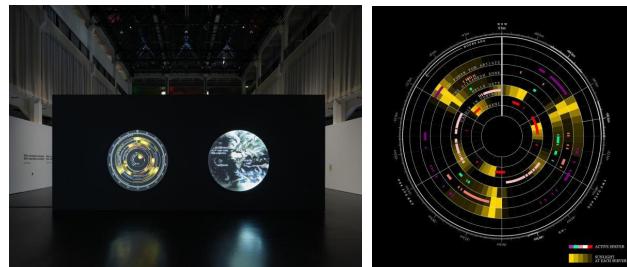


Figure 13 Tega Brain, Benedetta Piantella and Alex Nathanson, Solar Protocol, 2021-2023, in Renaissance 3.0 at the ZKM | Karlsruhe, 2023. ©ZKM | Center for Art and Media Karlsruhe. Photo: Tobias Woottton.

Figure 14 Data visualization of network activity.

©'Tega Brain-Solar Protocol, 2023.' Courtesy of the artists.

Deleuze, Gaze-Guided Metamorphosis, and the Politics of Claude

My installation can be read less as an 'interactive representation' of planetary crisis than as a Deleuzian production of sensation of a staging of becoming in which bodies (human, creature, environment, dataset) are not stable entities but temporary coagulations of forces.^[10] Deleuze's philosophy repeatedly insists that what matters

is not the subject as an origin, but the processes that generate subject-effects a positive, world-making 'desiring-production' that is non-anthropocentric and does not presuppose a sovereign self behind it^[11]. In this light, my 'gaze-guided metamorphosis' is not merely a UX affordance; it is an ontological claim. The viewer's attention proxy becomes a vector of intensity that participates in composing what a body can do in the scene of how color, breath, and behavioral temperament thicken or thin in response to proximity, dwell, and refusal.^[12]

Deleuze's notion of the body is especially useful here because it refuses the organism as the natural unit of analysis. In Francis Bacon: The Logic of Sensation, he develops the 'body without organs' as an intensive body opposed not to organs as such, but to the organization called the organism of an unstable field traversed by waves, where organs are provisional and effects are real rather than representational. My creatures' shifting skins and affective emissions can be positioned as precisely this: not symbolic 'avatars' but intensive surfaces where forces become legible as sensation. Figure 15 & 16 are body-without-organism in which provisional organ whose surface indexes force, friction, and exposure. Its cast shadow functions as a second, algorithmic body in which a diagram of continuous adjustment so that embodiment reads less as representation than as modulation wherein perception translated into parameters, and form returned as a governable outcome. The white-cube neutrality amplifies the control problem the work names how an apparently empty room stabilizes a regime where bodies are produced as readable, trackable, and endlessly recalibrated intensities. The biomechanical landscape of my installation then reads as an assemblage of milieu of material-semiotic constraints within which becoming is continuously negotiated. Importantly, this framing also sharpens the project's critique makes visible that 'embodiment' in interactive AI art is never innocent but always a technical construction that decides which intensities count^[13]. That is where Deleuze's late diagnosis of control becomes essential. In 'Postscript on the Societies of Control,' he distinguishes enclosure-based discipline from control's continuous variation: unlike molds, control operates as 'a modulation, like a self-deforming cast,' producing coded access, perpetual adjustment, and the reduction of persons into 'individuals' (data-samples, banks, markets).^[14] The Anarchist Library my gaze/proximity pipeline can be argued to perform this logic in miniature, attention proxy is operationalized as a measurable signal, routed through tracking infrastructure, and returned as aesthetic consequence. If this were presented as frictionless 'co-creation,' it would risk reproducing the ideology of capture (gaze-as-input, affect-as-feedback, subject-as-dataset). But my work can instead insist on exposure thereby the interface is a laboratory in which spectators are made to feel-viscerally-how quickly a body becomes governable once attention proxy is translated into control parameters. In Deleuzian terms, the piece becomes a diagram of modulation that stages the conversion of living perception into machinic governance, while still allowing moments of counter-becoming (withdrawal, refusal, opacity, misrecognition) to interrupt the loop.

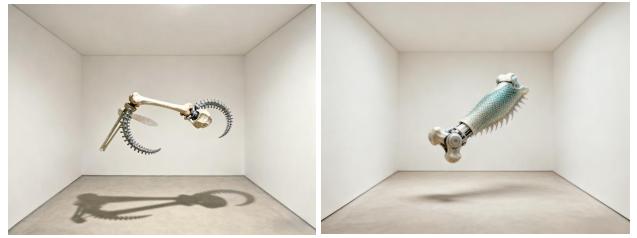


Figure 15 Bone-spine prosthesis folds the 'organism' into a detachable hinge

Figure 16 A single biomechanical 'organ' floats in enclosure, sheathed in fish-like scales that behave as an affective membrane rather than an anatomical skin. ©Respect Copyright.

This is also where Amanda Askell's recent Anthropic-facing discourse becomes a productive foil for critique, not because it is simplistic, but because it makes the political stakes of 'character' unusually explicit. Askell describes her work as training models to be 'more honest' and to have 'good character traits,' framing alignment as something like virtue cultivation at the level of fine-tuning^[15]. In a widely circulated interview transcript, she goes further claiming alignment is 'about whether the model has a good character,' while 'character training' embeds traits more deeply than mere roleplay, and the system prompt functions as a 'final tweak' that grants developers fine-grained behavioral control^[16]. Notice what becomes thinkable under this vocabulary: AI governance as the engineering of temperament, discretion, and 'global citizen' comportment. This is not far from a constitutional imaginary of conduct especially since Anthropic explicitly describes 'Constitutional AI' as giving models values via a written constitution, making those values more legible and adjustable than implicit RLHF^[17].

A constitution is a reterritorialization of becoming. It risks treating ethics as a stable list of principles that can be layered onto a system whose actual 'body' is distributed across training data, inference infrastructure, evaluation regimes, and product incentives. In other words, constitutional language can smuggle back the organism image of coherent personality and unified agency precisely where Deleuze asks us to track assemblages, contingencies, and the politics of composition. This is not merely abstract. Even Anthropic's own public-facing materials frame interaction as managerial control in which users are advised to treat Claude as a 'brilliant but very new employee (with amnesia) who needs explicit instructions,' emphasizing norms, style-guidelines, and procedural clarity as levers over behavior. Claude The 'employee' metaphor naturalizes hierarchy (instruction → compliance) and stabilizes the illusion of a governable subject, rather than foregrounding the distributed production of outputs. At the same time, Askell's remarks also expose a genuine tension my artwork can amplify rather than dismiss, if a model must operate across plural value-worlds, then a single moral theory can make it 'brittle,' whereas a posture of reflective 'moral uncertainty' may be safer than certainty or nihilism. This tension is exactly where a Deleuzian intervention can be most incisive, instead of

treating uncertainty as a psychological trait ('open-mindedness'), you can treat it as a structural condition of any planetary-scale assemblage, where bodies are continuously composed across heterogeneous forces (climate indicators, orbital debris, datasets, sensors, institutions). my creatures' metamorphoses, driven by entangled planetary variables and audience attention, provide an aesthetic argument that ethics is not a detachable 'constitution' but a situated practice of composing relations under constraint. And my choice of an immersive, bodily format matters, it prevents 'values' from remaining purely discursive, forcing them to register as affective consequences on skin, gait, breath, and vulnerability.



Figure 17. 'Constitutional limb' as misrecognised subject (employee-body). A folded prosthetic leg is cinched into a single, legible unit compact 'organism image'.

Figure 18. Assemblage-body against constitutional brittleness (ethics as composition). ©Respect Copyright.

The gauze-like sheath stages constitutional language in Figure 17 as a thin normative skin laid over an infrastructural chassis, foregrounding how moral principles risk becoming a managerial wrapper that hides the distributed assemblage (datasets, incentives, evaluations) that actually produces behaviour. The body appears as something that could be stabilised, corrected, and made compliant governed by instruction. The skeletal quadruped-like mechanism, threaded with translucent membranes and gear-like joints, refuses a coherent personality-form in Figure 18 represents composite that can only be composed, not commanded. Here 'moral uncertainty' reads as structural rather than psychological ethics of situated recomposition under constraint because the body is visibly dispersed across linkages and couplings. My work is not about making AI art more humane; it is about revealing how 'humanness' is operationalized as a governance layer. Wired's reporting on Anthropic's constitutional framing notes both the cherry-picked provenance of principles and the recursive structure ('a version of Claude monitoring Claude').^[18] This is a striking contemporary scene of what Deleuze would call control-by-modulation, governance displaced into self-policing feedback, legible as benevolence while remaining continuous, infrastructural, and hard to contest. my installation can therefore be positioned as a counter-infrastructure that makes modulation sensible, not seamless; it renders the 'constitution' felt as a pressure on bodies rather than a promise of safety. In doing so, it proposes a futures-oriented task for media art after generative AI, not to decorate alignment with aesthetics, but to aestheticize alignment as a problem of embodied, planetary governance, where becoming is never free, and where the body is the first site on which any 'value' is

written.

Methodology I chose a latent diffusion model Seedream 4.0 as the generative engine. I thank the Seedream 4.0 developer documentation and support resources for clarifying API integration constraints relevant to real-time rendering (supported output sizes, response formats, and latency considerations for event-triggered updates). Diffusion models excel at producing high-resolution images from extremely large, diverse dataset. For creature imagery, training on billions of images affords more novelty and ambiguity than GANs is critical, which were traditionally limited by smaller curated sets. For example, Seedream 4.0 was pretrained on billions of text-image pairs spanning diverse taxonomies and knowledge-centric concepts, enabling semantically grounded generation at native 1K–4K resolutions.^[19] Our reference library combined natural forms (flora, fauna, microscopic organisms) with synthetic textures (circuit patterns, digital artifacts) and climate phenomena (wildfires, floods, pollution events). After each exhibition run, we update the project-specific reference library with non-identifying documentation of the installation context (e.g., material textures, lighting conditions, and environment-only frames) and curated visual descriptors derived from exhibition logs. This maintains a feedback loop in which the installation both consumes and produces context, without implying any retraining of the base model.

I deliberately avoid simply optimizing the model for realism or bias-free output. Instead, inspired by avant-garde ML art practice, we value its glitch-like artifacts. I shape 'productive failure' at inference time by adjusting sampling steps and guidance, and when performing image-to-image or inpainting by injecting controlled noise and varying the edit strength to trade off structural stability against mutation. Because the generation service does not expose low-level schedulers, climate variables modulate (i) prompt descriptors, (ii) edit strength in image-to-image / inpainting, and (iii) a GPU shader stack that overlays crack-veins, soot bloom, and desiccation noise; the combined effect reads as 'climate scars' without claiming direct intervention in the model's internal denoising schedule. In effect, the environment feeds back into the generative noise, so that worse climate conditions produce more visibly damaged aesthetics on the creature. This amplifies the metaphor of planetary stress inscribed onto bodily form.

I conducted observational logging across N sessions, supplemented by short exit prompts (1–2 questions) focused on perceived agency, discomfort, and perceived causality of scars. Any audience-derived imagery is gathered only under clear on-site notice and an opt-out procedure; frames are processed locally to remove identifiers (e.g., face blurring and tight cropping), then discarded after feature extraction, with only aggregated interaction metrics and non-identifying visual descriptors retained for subsequent runs.

Rendering Pipeline and Spatial Display The generated creature images are integrated into a real-time 3D spatial scene. For each creature, we build a lightweight rig (basic

skeleton) and a UV-unwrapped, deformable mesh. At runtime, Seedream 4.0 generates a 2D texture map at a fixed resolution (e.g., 1024×1024 for sub-second to a few seconds, or 2K for higher fidelity; Seedream 4.0 also supports outputs up to 4K when needed).^[20] We generate an opaque texture and derive transparency via a separate matting pass (foreground segmentation), so the creature silhouette can be composited cleanly onto the UV-mapped mesh without assuming native RGBA output from the generator. We bind the generated texture to the mesh and update it through an event-driven refresh policy rather than a continuous loop. Specifically, a new generation request is issued only when an interaction trigger fires (e.g., gaze dwell time > 0.8–1.2 s, proximity < 1–1.5 m, or a threshold crossing in the minute-scale environmental feed), and each creature is rate-limited (e.g., one request every 2–4 s) to keep latency predictable. When a new texture arrives, the renderer performs a double-buffered material swap with a short crossfade (e.g., 300–600 ms), so the creature appears to morph smoothly while the underlying updates occur at a seconds-scale cadence rather than per-frame. The mesh geometry also responds, for example, a wind speed reading might cause subtle vertex displacement, so the creature appears to quiver or grow spikes. Lighting in the CAVE is ambient and soft, but the creatures emit an eerie self-illumination (subtle glow) to ensure visibility and a ghostly presence.

Our spatial display is a four-wall CAVE (three walls + floor) with edge blending. We run the Unity renderer as a frame-locked, swap-synchronized projector cluster, while high-precision head-pose/orientation tracking drives the off-axis view frustum, ensuring the edge-blended projections remain geometrically coherent around the viewer. The projected space represents an abstract landscape partly biological (vines, rock) and partly mechanical (wireframe grids, digital field lines) through which creatures move. Upon entry, depth cameras estimate the participant's position, and IR tracking provides 6-DoF head pose; rather than claiming eye-tracked gaze, we use head orientation as a view-direction proxy to infer what the participant is likely attending to. The system treats proximity and attention-proxy as engagement signals: when a participant approaches a creature or maintains head-oriented dwell toward it for >0.8–1.2 s, the creature locally intensifies its coloration (e.g., warmer hues or pulsating patterns) and may emit a low-frequency 'breathing' tone. We term this gaze-guided metamorphosis of reproducible trigger policy based on head-pose dwell and distance thresholds, inspired by prior 'attention proxy'-responsive works but implemented here without claiming eye-tracked fixation. The creature's evolution is therefore partially scripted at the system level yet locally co-shaped by participant behavior and environmental conditions.

Live Data Mapping (Minute-Scale) Environmental data inputs form a central feedback loop, but we distinguish between fast local signals and slow planetary indicators. We poll minute-scale local measurements that plausibly change on short timescales, e.g., air-quality indices (AQI / PM2.5), temperature, humidity, wind speed—using lightweight API queries, and we smooth them with an

exponential moving average to avoid visual flicker. In parallel, we ingest slow-moving background variables (e.g., atmospheric CO₂ baseline values and long-horizon sea-level indicators) at a daily-to-weekly cadence and treat them as session-level parameters that drift gradually across runs rather than updating minute-by-minute. All inputs are normalized to bounded ranges and routed into a shader stack controlling the creature's appearance. For example, the epidermis is implemented as a layered procedural material of 'healthy' fractal-noise layer (green/brown) and a 'scar' layer (cracks/scales) blended by environmental stress. Higher pollutant load increases the opacity and contrast of crimson crack-veins, while higher heat increases desiccation granularity. Where water is mapped at high frequency, we use local tide/water-level readings (not sea-level rise) to drive short-timescale 'drip' and buoyancy cues; long-term sea-level change is rendered as a slow, cumulative morphological bias across sessions.

Because the external generation service does not expose internal latent variables, we implement data-conditioning through inference-time controls: (i) prompt modifiers, (ii) image-to-image/inpainting edit strength, and (iii) reference-image weighting in multi-reference conditioning. For instance, when CO₂-related stress indicators exceed a threshold, the prompt is augmented with descriptors such as 'scarred,' 'soot-stained,' or 'desiccated,' while the edit strength is slightly increased to permit more morphological drift. This keeps the technical claim accurate while still making the creature's identity responsive to planetary conditions.

Failure Modes as Aesthetic Strategy Rather than treating system failures as defects to be eliminated, the installation incorporates recurrent failure modes of generative and interactive AI as productive constraints. Semantic drift is central. In diffusion-based image-to-image generation, iterative prompting and latent interpolation frequently cause gradual divergence from an initial intent. Whereas conventional pipelines suppress this drift to maintain coherence, the installation embraces it as a visual analogue of becoming that creatures exceed fixed morphological identity, accumulating hybrid traits irreducible to any single prompt or reference. Drift thus becomes legible as a temporal process, foregrounding how meaning in generative systems is always provisional and iteratively negotiated. A second failure mode concerns mode collapse and repetitiveness, where the model converges on dominant visual motifs. Rather than countering this through randomness, the system employs corpus rotation, periodically reweighting subsets of the curated image strata according to session duration rather than output quality. Repetition is redistributed rather than eliminated, producing family resemblance instead of sameness and reframing recurrence as an archive effect rather than a model flaw.

The installation also confronts tracking noise in attention proxies based on head orientation and proximity. These signals are discontinuous and sensitive to micro-movements. By introducing dwell-time thresholds and hysteresis windows, the system preserves lag,

misalignment, and overshoot, making control perceptibly effortful. Attention appears not as frictionless input but as a contested signal whose translation into effect remains partial. Finally, data absence and spikes in planetary indicators are handled through bounded smoothing and fallback states. When data drop out, the system visibly holds or mutes its state, allowing absence itself to register. Together, these failures articulate instability is not an obstacle to meaning but its condition, exposing how interactive AI governs bodies and environments through limits rather than seamless control.

Climate-Scarred Biomachinic Creatures and the Politics of Generative AI in Planetary Computation

The installation unfolds in the darkened CAVE as visitors enter. Multiple creatures wander the space, some stationary on the floor, others climbing projected walls or floating. These entities blend familiar and alien elements, a sheep-like torso with fluorescent veins, a tree-root spine fused to mechanical claws, eyes that reflect distant storm clouds. Figure 20 shows one such creature form its woolly texture is overlaid with luminous crackles and asymmetrical growths, hinting at an infection or mutation. These visual glitches encode real data. For example, skin ripples and sway are driven by local wind speed and indoor airflow cues, while pigment bruising and crack-vein contrast respond to short-timescale pollutant load (AQI/PM2.5). Slower indicators such as CO₂ baselines bias longer-horizon shifts in coloration and morphology across sessions rather than producing minute-by-minute jumps.

Technically, each creature's model was generated by running a diffusion Seedream 4.0 network with a balanced combination of class-conditioned prompts ('animal,' 'fungus,' 'machine') and actual climate keywords ('burned,' 'scarred,' 'flooded'). We iterated this process offline to build a library of prototype morphologies and texture maps. In the live system, these prototypes serve as bases that are then iteratively re-diffused when influenced by data or viewer interaction. For instance, if a viewer stares at a creature for several seconds, the model performs an inpainting step takes a cropped view of the creature's face or limb, and the text prompt shifts slightly toward emotive terms ('pain,' 'mutation'), causing that region to warp or add appendages.

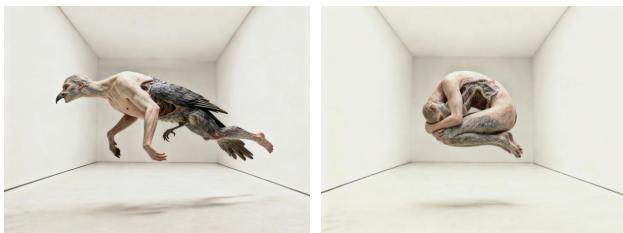


Figure 19&20. Diffusion-generated 'bio-synthetic' creature. Natural textures (e.g. feather) are twisted by algorithmic distortions and climate-driven anomalies, producing a haunting hybrid form. ©Respect Copyright.

By spatially embedding these creatures, the installation

becomes an 'affective interface' invites empathy with a nonhuman subject. Participants may feel both awe and unease as they watch the creatures transform. The interface is multi-sensory in response to data-driven changes, ambient soundscapes evolve, for example, harsher noise when pollution spikes. We deliberately avoid didactic graphs; instead, the environment embodies the data through living forms. This design emphasizes posthumanist readings of humans, technologies, and environments are not separate realms but mutually constitutive. Participants eventually sense that their actions and global environmental conditions 'hurt' or 'heal' these beings, prompting reflection on our ecological entanglement.

My generative installation contributes novel conceptual and formal approaches at the intersection of art, AI, and ecology. Metabolic Data Treatment in conceptualizing the training data as a metabolic feedstock, we critique the typical extractive framing of AI datasets. Rather than merely harvesting images, we cultivate an evolving dataset shaped by ecological inputs (local climate statistics) and human presence. Rather than claiming control over a model's original training corpus, we locate authorship in curating a project-specific reference library and in shaping inference-time conditioning (prompting, reference weighting, and controlled editing). Between exhibitions, we refresh this library using field documentation and sensor-informed descriptors, as though keeping a generative ecosystem nourished.



Figure 21&22 'Failure' as epistemic surface of Modulation, not representation. ©Respect Copyright.

By avoiding a single objective (e.g. realism or classification accuracy), I challenge norms of AI. As Elwes and Hemment observe, many ML artists find more generative potency in algorithmic 'failures' than in polished outputs. Our creatures deliberately exhibit such failures like anatomical anomalies, dream-like blending, and texture glitches. These are not errors to be fixed but features, they reveal the model's biases and the random seed of computation. For example, in Figure 21&22, if the diffusion model hallucinates a second face or a bio-luminous organ, we allow it, interpreting it as symbolic (perhaps a 'double self' of the creature bearing witness to two worlds). A hybrid creature with incompatible anatomies and an exposed internal cavity makes diffusion's anomalies legible as method rather than mistake. The doubled/ambiguous facial logics and seam-like textures mark where heterogeneous image regimes collide, indexes of bias, dataset strata, and random-seed contingency that resist closure into a single narrative of realism. This aesthetic embraces ambiguity and invites viewers to derive their own narratives, rather

than conveying a closed message.

The core novelty is the embodied mapping of environmental data onto creatures' bodies. This makes palpable the idea that climate information is not neutral abstraction but has visceral effect. In terms of feminist/environmental theory, the creatures become cyborgian hybrids of part organism, part medium. Storey would call them incarnations of a biomachinic planet, they display how ecological and computational networks intermesh into new life forms. As our creatures writhe and morph, they manifest the paradoxes Storey outlines of the planet's algorithms (our climate models and sensors) mediate its alterity and yet fail to fully capture it. In other words, we literally show that the datafication of nature is incomplete: scars on skin are a poetic irreducibility that no algorithm can fully rationalize. This resonates with Tega Brain's critique: the environment is not merely a bounded system, but possesses its own 'intelligence' which resists total control. I also invert the usual techno-optimism of planetary computation. Many climate tech platforms assume that more data and more AI will magically solve the crisis (the 'good Anthropocene' vision). By contrast, our installation highlights fragility and partiality. The creatures are not tamed by technology; they remain uncanny and autonomous. Like Solar Protocol, which uses renewable energy to limit computation rather than simply automate it, I let environmental limits guide creation. The installation is a speculative experiment of what if the Earth's own rhythms were the generative force? Each creature is thus a provisional co-creation of human, machine, and planet.

From a media-archaeology perspective, our system foregrounds the materiality and politics of AI. The very choice of diffusion Seedream 4.0 rests on the literal mining of earth such as the silicone, cobalt and copper in processors and the open knowledge sharing (LAION dataset) hidden geologies that Parikka reminds us underpin media. By making the artwork climate-reactive, we invoke that subterranean story of the creature's body is not separate from the minerals burned to create the compute, nor from the forests burned by global warming. I hope viewers leave with an embodied sense that computational interfaces far from ethereal are deeply inscribed in the natural world.

Conclusion

'Metabolic Carcass of the Dataset' demonstrates generative training data as metabolized matter, integrating real-world sensors into the creative loop, and merging machine forms with living ecologies. Technically, the use of Seedream 4.0 diffusion model (with a vast pretraining corpus) allows us to update them on a seconds-to-minute cadence, combining minute-scale environmental polling with event-triggered texture updates and short crossfades in the CAVE renderer.

Artistically, mapping climate metrics onto bodies creates an affective interface that embodies planetary dynamics. The project thus critiques the trope of AI optimization and instead celebrates nonhuman agency and

environmental feedback. Grounded in critical AI, posthuman theory and Deleuzian becoming, the work makes visible the metabolic circuits between planet and algorithm. It exemplifies Storey's insight that contemporary planetary culture is inseparable from computation, but turns that insight into a performative space that the planet speaks through our creatures. This addresses emergent concerns about AI and the environment by showing concretely how climate change can be felt via generative media, rather than abstractly charted. When modulation becomes sensibly registers on skin, breath, gait, and hesitation, new forms of refusal, opacity, and counter-becoming become thinkable as aesthetic practice and as planetary politics.

Reference

- [¹] Storey, Thomas. 'Encountering the Biomachinic Planet.' *Culture Machine* 24 (2025).
https://wp-content/uploads/2025/11/Storey_CM24_BiomachinicPlanet.pdf
- [²] Yusoff, K.: Epochal aesthetics: Affectual infrastructures of the Anthropocene. In: Axel, N. et al. (eds.) *Accumulation: The Art, Architecture, and Media of Climate Change*, pp. xx–xx. University of Minnesota Press, Minneapolis (2022)
- [³] Vidmar, M., Hemment, D., Murray-Rust, D., Black, S.R.: On creative practice and generative AI: Co-shaping the development of emerging artistic technologies. In: Terras, M., Jones, V., Osborne, N., Speed, C. (eds.) *Data-Driven Innovation in the Creative Industries*, pp. 196–218. Routledge, London (2024). <https://doi.org/10.4324/9781003365891-9>
- [⁴] Grba, D.: Art notions in the age of (mis)anthropic AI. *Arts* 13(5), 137 (2024). <https://doi.org/10.3390/arts13050137>
- [⁵] Akten, M., Fiebrink, R., Grierson, M.: Learning to See: You Are What You See. *arXiv*:2003.00902 (2020).
- [⁶] Guljajeva, V., Canet Sola, M.: AI portraits, dream painter, and circuits: Interactive AI art installations. *Computer Graphics Forum* 42(1), 15–30 (2023)
- [⁷] Bratton, B.: *The Stack: On Software and Sovereignty*. MIT Press, Cambridge, MA (2015)
- [⁸] Parikka, J.: *A Geology of Media*. University of Minnesota Press, Minneapolis (2014)
- [⁹] Brain, T.: The environment is not a system. *APRJA – A Peer-Reviewed Journal About* 7(1), 153–165 (2018).
<https://aprja.net/article/view/116062/164244>
- [¹⁰] Deleuze, G.: *Francis Bacon: The Logic of Sensation*. Trans. Smith, D.W., pp. 27–35. University of Minnesota Press, Minneapolis (2003)
- [¹¹] Deleuze, G., Guattari, F.: *Anti-Oedipus: Capitalism and Schizophrenia*. Trans. Hurley, R., Seem, M., Lane, H.R., pp. 1–5, 26–28. University of Minnesota Press, Minneapolis (1983)
- [¹²] Deleuze, G., Guattari, F.: *A Thousand Plateaus: Capitalism and Schizophrenia*. Trans. Massumi, B., pp. 21–23, 260–265. University of Minnesota Press, Minneapolis (1987)
- [¹³] Suchman, L.: *Human–Machine Reconfigurations: Plans and Situated Actions*, 2nd edn. Cambridge University Press, Cambridge (2007)
- [¹⁴] Deleuze, G.: Postscript on the societies of control. *October* 59, 3–7 (1992)
- [¹⁵] Askell, A.: About me. <https://askell.io/>, last access 14 Dec 2025
- [¹⁶] Gorsky, M.A.: Amanda Askell: I avoid lying to AI models. LinkedIn (interview transcript), 30 Apr 2025.
<https://www.linkedin.com/pulse/amanda-askell-i-avoid-lying-ai-models-mikael-alemu-gorsky-gj3rf>, last access 14 Dec 2025
- [¹⁷] Bai, Y., et al.: Constitutional AI: Harmlessness from AI feedback. *arXiv*:2212.08073 (2022). <https://doi.org/10.48550/arXiv.2212.08073>
- [¹⁸] Levy, S.: If Anthropic succeeds, a nation of benevolent AI geniuses could be born. *Wired*, 28 Mar 2025.
<https://www.wired.com/story/anthropic-benevolent-artificial-intelligence/>, last access 14 Dec 2025
- [¹⁹] Team Seedream et al.: Seedream 4.0: Toward next-generation multimodal image generation. *arXiv*:2509.20427v3 (2025).
<https://doi.org/10.48550/arXiv.2509.20427>
- [²⁰] ByteDance Seed: Seedream 4.0.
https://seed.bytedance.com/en/seedream4_0, last access 13 Dec 2025