Article

Desire and Data: Metabolism of AI

VADIM EPSTEIN, independent researcher, United Kingdom, email: eps@wesual.xyz

10.2478/bsmr-2025-0002

ABSTRACT

Generative artificial intelligence (AI) is reshaping audio-visual media, yet its long-term creative sustainability remains uncertain. This article proposes a metabolic framework to explain how AI systems consume, transform, and renew cultural material, and how these cycles depend on human creativity. Drawing on cybernetics, general systems theory, and autopoiesis, the study situates AI within a coupled human-machine ecosystem. Human creativity is both a fuel that sustains AI and the antidote to its stagnation, providing unpredictable variation that prevents recursive degradation while remaining vulnerable to commodification that limits regeneration. A qualitative, practice-based approach combines conceptual analysis with case studies of independent and marginal AI art/research. Four thematic dynamics are examined: (1) aesthetic convergence, with synthetic folklore as algorithm-native form; (2) creative entropy, including Model Autophagy Disorder (MAD), with desire as a counterforce to homogenisation; (3) creative metabolism, a cycle of cultural input and transformation addressing erosion of desire under algorithmic optimisation; and (4) creative countercurrent practices resisting corporate logics and highlighting marginality's role in sustaining diversity. Case studies – from CLIP-VQGAN and Aphantasia to ConceptLab, Active Divergence, and Roope Rainisto's post-photography – show how low-resource, critical engagement with AI fosters cultural variation. Findings suggest that supporting marginal creative practices is essential not only for aesthetic diversity but for the viability of AI-mediated creative ecosystems. Future research considers emerging agentic AI systems and autonomous artists.

KEYWORDS: **Artificial intelligence; data aesthetics; machine desire; computational culture; algorithmic agency**

INTRODUCTION

The integration of generative AI into audio-visual media production represents a fundamental shift in creative practice that extends beyond technological adoption. What began as discrete applications for generating isolated assets has evolved into comprehensive world-building systems with increasing sophistication and coherence. This shift demands new conceptual frameworks considering AI not only as a production tool but as a mediator of reality – filtering, prioritising, and shaping access to cultural expression (Gillespie 2018). When we speak of AI "world-building," we are not simply describing a technique for creating fictional universes in film or gaming; we are acknowledging how AI increasingly constructs the cognitive environment through which we perceive and interpret reality itself. In this context, the aforementioned frameworks should help us understand both AI's current capabilities and its long-term impact on creative culture itself.

Most critical discourse surrounding AI in creative fields has focused on immediate concerns such as copyright implications in training data (Buick 2025), biases in generated content (Bianchi et al.

2023), or potential displacement of human labor (Acemoglu et al. 2022). While these issues deserve attention, they arise from analysing discrete instances of AI-generated content rather than examining systemic dynamics of creative AI ecosystems. This approach overlooks a larger, system-wide problem: AI models do not stay static. Instead, they move through endless cycles of training, producing results, and retraining, reshaping themselves over time. This recursive process raises questions about long-term sustainability and cultural diversity.

To address this, we draw on the metaphor of metabolism, which is understood in systems theory as complex processes sustaining complex organisms through consumption, transformation, and renewal cycles. We observe how human desires route to behaviours, behaviours transform to data, and data becomes AI – a cycle that, without proper balance, risks collapsing into increasingly uniform results. Such a perspective reveals how human creative impulses function as metabolic input and essential nutrients that, when properly circulated, maintain creative vitality by introducing unpredictability that counteracts homogenising statistical tendencies. Conversely, when these creative inputs become commodified or homogenised, they lose their capacity to provide the "controlled noise" necessary for preventing algorithmic stagnation.

The study foregrounds marginal and counter-hegemonic practices as crucial sites for resisting convergence, offering both technical and aesthetic interventions. These practices not only maintain diversity in AI outputs but also ensure the long-term resilience of human-AI creative relations.

LITERATURE REVIEW

Creativity, Imagination, and Artificial Intelligence

Vygotsky ([1930] 2004) distinguishes between reproductive activity, which involves repeating previously experienced behaviours; and creative/combinatorial activity, which combines elements from past experience to generate new constructions. He positions imagination as the foundational brain function enabling creative activity through its capacity to combine and rework elements of previous experience. This complements Runco and Jaeger's (2012) widely cited definition of creativity, which requires both novelty and appropriateness or value, thereby establishing that truly creative outputs must be unprecedented and meaningful within their cultural context.

Neuropsychological research further clarifies these processes: Beaty et al. (2014) identify creative cognition as involving both associative combination (linking existing concepts) and executive control (evaluating and refining ideas in context). Boden (2004) demonstrates how computational models can illuminate the mechanisms underlying human creativity, arguing that understanding creative processes through artificial intelligence helps demystify rather than diminish the nature of human creative thought.

Venkatesh Rao's (2024) distinction between imagination and creativity provides additional conceptual clarity. Imagination involves pattern recognition and recombination within existing knowledge domains – identifying potential connections within familiar conceptual territories. Creativity, by contrast, generates genuine novelty through open-ended exploration that ventures beyond established patterns to discover previously unknown possibilities. Contemporary AI systems excel at imagination but remain fundamentally bounded by their training distributions, struggling to achieve the kind of generalisation beyond training domains that characterises human creativity (Colton and Wiggins 2012).

This limitation becomes critical when considering the lifecycle of AI development over extended periods. Hertzmann (2018) argues that computational creativity remains fundamentally imitative rather than genuinely innovative, capable of sophisticated pastiche but incapable of the radical departures from existing patterns that drive cultural evolution. This assessment aligns with empirical research on Model Autophagy Disorder, which demonstrates how AI systems experience progressive degradation when trained recursively on their own outputs (Alemohammad et al. 2024).

Cybernetics, Systems Theory, and Autopoiesis

The metabolic framework employed in this analysis draws primarily on the cybernetic tradition initiated by Wiener (1948) who conceptualised feedback systems as analogous to biological regulation. General systems theory (Von Bertalanffy 1968) further formalised the study of self-regulating systems, which was developed through the biological systems theory of Maturana and Varela (1980). Their concept of autopoiesis – the capacity of living systems to maintain and reproduce themselves through continuous self-production – provides essential theoretical grounding for understanding how creative systems sustain themselves over time.

Varela, Thompson, and Rosch's (1991) development of enactivism extends autopoietic principles to cognitive systems, arguing that cognition emerges from embodied interaction between organism and environment rather than abstract information processing. This enactive perspective proves crucial for understanding why human creativity provides irreplaceable input to AI systems: human creative acts emerge from embodied experience, cultural context, and the complex interplay between conscious intention and unconscious drives that cannot be replicated through algorithmic processing.

The enactive approach emphasises that cognitive systems are not passive receivers of environmental input but active agents that co-create their worlds through sensorimotor engagement (Thompson 2007). This perspective suggests that human creativity operates through what Varela termed "structural coupling" – dynamic processes where cognitive systems and environments mutually influence each other's development. Such approach informs our view of AI as a loosely coupled system, metabolising human creative inputs into algorithmic outputs.

Beer's (1972) application of cybernetic principles to complex organisational systems provides additional theoretical support for understanding how creative ecosystems maintain viability through feedback loops and adaptive responses. His viable system model demonstrates how complex systems require variety and adaptability to survive in changing environments – principles that prove essential for understanding why creative AI systems require ongoing input from diverse human sources.

Apparatus Theory and Technical Mediation

Flusser's ([1983] 2011) analysis of photography as a "technical image" produced by an apparatus rather than direct human action offers a valuable lens for understanding AI systems as creative mediators rather than transparent tools. In his account, the apparatus – such as the camera – operates according to its own programmatic logic, shaping what can be created through its technical constraints and possibilities, even as it appears to serve human intentions. This "programmed freedom" allows users to explore a range of outcomes, but always within parameters set by the apparatus, creating the illusion of unlimited creativity while subtly directing it. In this framework, the photographer becomes a "functionary" of the apparatus, working within the system's possibilities rather than transcending them – a dynamic equally applicable to AI tools, whose training data, architectures, and optimisation objectives similarly structure and limit creative expression.

Agamben's (2009) contemporary extension of apparatus critique provides additional theoretical grounding for understanding how AI systems capture and channelise human creative energy.

Agamben argues that apparatus operates through processes of capture, separation, and control that transform human capacities into manageable resources for systemic purposes. This perspective aligns with platform studies, a field that offers key insights into the extraction processes analysed in this article.

Simondon's (2017) analysis of technical objects offers complementary insights into how technical systems develop their own evolutionary trajectories that may diverge from human intentions. His concept of "concretisation" describes how technical objects become increasingly autonomous and self-determining through their development, developing internal coherence that shapes their future evolution. This perspective helps explain how AI systems develop aesthetic tendencies and creative patterns that emerge from their technical structure rather than explicit human programming.

Platform Studies and Digital Culture

Understanding AI's role in creative culture requires engagement with platform studies scholarship that examines how digital infrastructures shape cultural production. Gillespie's (2018) analysis of content moderation and algorithmic curation reveals how platforms actively shape creative possibilities through recommendation systems, engagement metrics, and content policies rather than serving as neutral distribution channels.

This platform perspective is central for understanding the extraction processes examined in the creative metabolism framework. Digital platforms harvest behavioural data about creative preferences, aesthetic choices, and cultural trends, transforming user engagement into structured datasets suitable for training AI systems (Zuboff 2019). This extraction operates through what Zuboff terms "surveillance capitalism" – business models that derive value from predicting and modifying human behaviour rather than serving user needs.

Bucher's (2018) analysis of "programmed sociality" demonstrates how algorithms become embedded in social practices, particularly through platforms like Facebook's News Feed, which shapes conditions for social interaction while users simultaneously adapt to and work within these algorithmic systems. Pariser (2011) introduces the concept of "filter bubbles" to describe how algorithmic personalisation by search engines and social media platforms creates individualised information environments that may reduce users' exposure to diverse viewpoints, potentially leading to intellectual isolation and the reinforcement of existing beliefs rather than challenging them.

Glitch Aesthetics and Technical Resistance

Menkman's (2011) foundational work on glitch aesthetics demonstrates how technical failures and system breakdowns can become productive creative resources rather than mere obstacles to overcome. Krapp (2011) conceptualises noise and glitch as counter-forces to efficiency and normativity, locating in digital error a potential to "jam" systems that would otherwise render themselves invisible.

Recent theoretical work explores how AI art operates through fundamentally different visual systems than human-centered aesthetics. Paglen (2019) examines how contemporary image-making increasingly functions through machine-to-machine processes that create visual abstractions largely alien to human perception, operating through "activations, keypoints, eigenfaces, feature transforms, classifiers, [and] training sets" rather than traditional representational frameworks. This shift challenges conventional aesthetic categories and suggests new possibilities for understanding non-human forms of visual intelligence.

These works suggest that technical peculiarities or failures, whether human-perceived or machine- native, can be generative – producing spaces for critical resistance, creative expression, and the reimagining of socio-technical norms.

THEORETICAL AND METHODOLOGICAL FRAMEWORK

This research adopts a qualitative, practice-based approach positioned within digital media studies while drawing on interdisciplinary theoretical frameworks including systems theory, media theory, and creativity studies. The investigation recognises that understanding AI's integration into creative culture requires both theoretical analysis and direct engagement with the technologies and practices under examination.

The research integrates three complementary approaches, addressing the interdisciplinary nature of AI creativity research. The conceptual synthesis establishes theoretical connections across these fields to develop the metabolic framework as an analytical lens. The case study analysis examines selected independent AI art and research projects according to specific criteria designed to identify practices that resist mainstream commercial approaches.

Finally, a practice-based investigation involves direct experimental engagement with AI creative tools and processes, following Sullivan's (2005) and Gray and Malins's (2004) conception of artistic research as a legitimate form of scholarly inquiry that generates knowledge through reflective creative practice. Selected examples from the author's own artistic practice are incorporated to illustrate specific theoretical concepts and technical approaches. This follows established practice-based research conventions where artists use their own work as illustrative material to demonstrate theoretical principles (Bolt 2007).

Given the rapidly evolving nature of AI creative practice, this research draws upon both peer- reviewed academic sources and contemporary practitioner discourse found in technical blogs and online documentation. Such sources provide essential insights into current developments and community practices that have not yet entered formal academic literature, making this methodology particularly essential for examining creative phenomena that resist quantification while maintaining scholarly rigour through systematic analysis and theoretical grounding.

Case Study Selection Criteria

Primary Criteria:

1. Methodological Innovation: Projects that develop novel approaches to AI creative tools rather than using them conventionally
2. Aesthetic Divergence: Work that produces outputs distinct from mainstream AI art platforms
3. Technical Experimentation: Projects that modify, subvert, or creatively misuse AI systems

Secondary Criteria:

1. Resource Constraints: Projects developed with (or aimed for) limited computational resources, demonstrating "guerrilla research" approaches (even if in academic environment)

2. Open Source Ethos: Work that shares methods and tools with broader community
3. Temporal Significance: Early experiments that influenced subsequent AI art development

Criteria for artworks may also be formulated as:
1. Deviation from mainstream AI aesthetics: privileging outputs that resist photorealistic or corporately optimised "house styles."
2. Critical engagement with the medium: exploring the affordances and limitations of AI in ways that question its default modes of operation.
3. Contribution to creative diversity: generating outputs, methods, or concepts that enrich the overall cultural dataset rather than narrowing it.

The criteria deliberately focus on marginal practices exhibiting critical engagement with the medium, rather than on commercially successful AI artists whose work, while aesthetically compelling, employs mainstream tools in conventional ways. Such focus reflects the metabolic framework's emphasis on diversity and resistance to homogenisation as essential for creative ecosystem health.

This triangulated approach addresses the interdisciplinary nature of AI creativity research while maintaining analytical rigour appropriate for academic inquiry. By combining theoretical analysis with experimental practice and empirical case studies, the approach provides multiple perspectives on how AI technologies are being integrated into creative culture beyond dominant commercial narratives.

AESTHETIC FRICTIONS

The phenomenon of aesthetic convergence in AI-generated content represents the most visible manifestation of creative metabolism dysfunction. Empirical observations by AI practitioners consistently reveal that AI systems, while producing technically sophisticated outputs, exhibit a pronounced tendency to generate content that represents statistical averages rather than creative extremes.

At the algorithmic level, neural networks operate through processes of generalisation and pattern recognition that exhibit implicit smoothing and averaging behaviours. Training procedures may impose biases that favour smoother functions over more complex ones, effectively averaging statistical regularities while potentially losing outliers and edge cases (Belkin et al. 2019).

The aspect of mass access to AI is important. Tools like Midjourney, DALL-E, and Stable Diffusion represent a new class of generative AI technologies that can produce high-quality artistic media, fundamentally altering the creative processes by which creators formulate ideas and put them into production. (Epstein et al. 2023). These technological developments have significantly lowered barriers to visual creation in unprecedented ways, enabling individuals without traditional artistic training to produce compelling imagery. This has genuine global reach, with artists in regions previously excluded from expensive creative software ecosystems now participating in digital art creation with minimal resources.

However, such accessibility operates at a cost. Standard models trained on dominant datasets inevitably reflect biases and aesthetic preferences embedded in their training data (Bianchi et al. 2023). Western visual traditions, commercial photography styles, and popular internet imagery become default templates from which AI generates new content. Local artistic traditions, marginalised visual cultures, and unconventional aesthetic approaches risk being filtered out in favour of statistically dominant patterns. Artistic research by Bellio and Kayser-Bril (2023) clearly illustrates how various image generation systems consistently reproduce problematic stereotypes while failing to achieve meaningful diversity across protected characteristics.

These converging effects compound as AI becomes integrated throughout creative production pipelines. Individual instances

FIGURE 1. Response to prompt "A presidential candidate is giving a speech in the street" generated by Adobe Firefly. Source: Bellio & Kayser-Bril (2023), AlgorithmWatch.

of aesthetic bias might appear manageable – for example, a filmmaker using AI for concept art could manually enforce diversity in outputs. However, as Flusser ([1985] 2014) observes regarding technical apparatus, the machine's "program" predisposes it toward certain outputs unless deliberately subverted through human intervention. In AI art, such subversion may involve prompt strategies, fine-tuning, or model retraining to introduce forms that the default statistical tendencies would not produce. Yet as AI mediates everything from initial concept development to final delivery, subtle tendencies toward statistical norms accumulate, potentially resulting in the significant narrowing of aesthetic possibilities across the broader media landscape.

Synthetic Folklore

Contrary to viewing convergence as purely negative, we must acknowledge AI-native aesthetic qualities as essential signatures rather than unwanted artefacts. AI anomalies have evolved into sophisticated forms of resistance, expanding from purely visual distortions into semantic interventions that challenge both sanitised commercial aesthetics and anthropocentric creative paradigms, thereby transforming technical constraints into new forms of creative expression. Each technological medium develops its own vocabulary of imperfection – from analogue film grain to digital compression artefacts – and AI's semantic disruptions represent the latest chapter in this ongoing dialogue between human intention and technological agency.

One of the most fascinating developments in AI-generated content involves the emergence of synthetic folklore – cultural myths and aesthetic styles originating from algorithmic processes rather than human cultural traditions. Unlike traditional folklore, which arises from specific geographic and cultural contexts over extended periods, synthetic folklore appears rapidly without clear cultural roots yet resonates with universal psychological patterns.

The phenomenon of Loab – a disturbing female figure that has persistently appeared across multiple AI image generations despite attempts to diverge from this imagery – illustrates this new category of AI-native mythology (Swanson 2022). Discovered through negative prompting techniques in text-to-image generation, Loab became an internet sensation – a haunting figure from the latent space of diffusion models whose unsettling appearance and strange consistency across different prompts created folk narrative around her quasi-mythological status as a "cryptid" or "demon in the machine."

Christoph Winkler's dance project "Whispers of Wood" (Winkler 2025) offers another less radical and more controlled example of synthetic folklore generated through what the field terms "AI hallucination". When prompted to provide stories about relationships between artificial intelligence and embodiment, ChatGPT fabricated a compelling narrative about a Yoruba myth involving a wooden horse coming to life and causing chaos – a story that research revealed never existed in Yoruba culture. Rather than dismissing this as error, Winkler embraced the AI's creative confabulation as a form of algorithmic storytelling, developing it into a collaborative dance performance where AI serves as co-creator.

Such phenomena represent more than plain curiosities. They reveal how AI systems, in their statistical processing of human cultural content, can generate genuinely novel mythological structures that resonate with human psychological patterns while remaining fundamentally non- human in origin. These AI-generated myths operate as a form of post-geographic folklore, unrooted from specific cultural traditions yet familiar to global audiences through their engagement with universal narrative and visual archetypes.

The aesthetic qualities marking content as "AI-generated" – often dismissively called "AI slop" – can actually serve as markers of healthy creative process. The slight uncanniness in human feature rendering, characteristic lighting and composition patterns, and semantic disruptions

FIGURE 2. Steph Maj Swanson, "Loab" character. Used with permission. Source: https://loab.ai, 2022.

FIGURE 3. Concept art for "Whispers of Wood." Created by Author using Leonardo.AI, 2025.

represent AI's native aesthetic language. As artist Roope Rainisto notes, "The artwork doesn't have to look realistic, but it has to feel real" (Cartagena 2023). Perfect realism or seamless conformity to established aesthetic conventions would paradoxically signal faster creative degradation, indicating system excellence at reproduction rather than innovation.

By foregrounding these AI-native mythologies, we see synthetic folklore as both a cultural by- product of algorithmic convergence and a potential site of resistance. Through intentional appropriation, artists can transform unintended machine artefacts into enduring cultural forms, thereby extending the diversity of the AI creative ecosystem.

CREATIVE ENTROPY

Understanding AI's long-term sustainability in creative contexts requires a distinction between imagination and creativity as related but discrete cognitive processes engaging AI systems in fundamentally different ways. As established in the literature review, imagination involves pattern recognition within existing knowledge – identifying potential connections and interpretations within familiar conceptual territories. Creativity, by contrast, generates genuine novelty through open- ended exploration, venturing beyond established patterns to discover truly new possibilities (Vygotsky [1930] 2004; Rao 2024).

Contemporary AI systems excel at imagination but struggle with true creativity. They can recombine elements from their training data in sophisticated ways, producing outputs that appear creative but remain fundamentally bounded by their training distributions. To achieve genuine novelty, AI systems would require generalisation beyond training domains – a capability that remains elusive despite significant research efforts (Zhou et al. 2021). Moreover, we should not confuse such true creativity with AI's own peculiarities discussed above; those glitch-born variations, while remaining aesthetically compelling and culturally significant, operate through mechanisms different to human innovation, representing statistical aberrations rather than discoveries.

While in most cases the difference between these two processes may look negligible (taking into account the enormous scale and diversity of information in the world's cultural heritage), this limitation becomes critical when we consider the lifecycle of AI development and deployment over extended periods. Researchers have identified a phenomenon called Model Autophagy Disorder (MAD), in which AI systems trained recursively on their own outputs, thereby experience progressive degradation (Alemohammad et al. 2024). Like biological autophagy (self-consumption), this process creates a feedback loop where the statistical simplifications inherent in each generation of AI output become amplified, with long tails of unusual or minority examples disappearing in favour of increasingly averaged, mainstream outputs. This degradation can be observed through iterative image processing with diffusion-based methods, operating on similar principles.

The implications extend beyond technical performance metrics, revealing concerning parallels between current and future cultural dynamics. Just as AI systems presently fade out minority voices in favour of statistically dominant mainstream content, authentic human culture itself may face similar algorithmic marginalisation. The very mechanisms that currently render diverse human communities invisible within AI systems threaten to position all humanity as a minority voice in landscapes dominated by dehumanised versions of our own cultural output. This represents more than cultural extinction – it suggests recursive processes where today's marginalised minorities serve as a preview of tomorrow's marginalised humanity, displaced not by alien content but by algorithmically processed echoes of themselves. This risks not only aesthetic homogenisation, but also the potential erasure of cultural DNA that has historically enabled marginalised communities and human societies more broadly to adapt and survive.

FIGURE 4. Degradation of imagery, recursively processed with Stable Diffusion image-to-image generation. Left: initial image, right: result after 50 iterations. Created by Author, 2023.

FIGURE 5. Cityscape imagery, artistically deteriorated by processing with recursively retrained Pix2Pix models. Created by Author, 2020.

While MAD is typically framed as a degenerative threat to diversity, it is important to acknowledge that the distortions and collapses arising from such processes can themselves be sources of creative opportunity. Just as glitch art embraces compression errors or data corruption as aesthetic material (Menkman 2011), degraded model outputs can inspire avant-garde exploration, repurposing statistical decay into deliberate stylistic strategies. Russell's (2020) glitch feminism framework suggests that system breakdowns create spaces for alternative formations that challenge normative structures. This perspective reframes MAD not as pure degradation but as a potential gateway to genuinely non-human aesthetic languages. The statistical drift occurring through recursive training might enable AI systems to develop creative directions that would be impossible through direct human programming. However, realising this productive potential requires active cultivation through creative experimentation rather than passive acceptance of decline.

The Role of Desire

Within the context of creative entropy, human desire functions as a critical counterforce to homogenising tendencies in AI-generated media. Desire, unlike algorithmic optimisation parameters, is emergent, embodied, and contextually situated (Varela, Thompson, and Rosch 1991). It arises from lived experience, cultural specificity, and the interplay between conscious intention and unconscious drives, making it inherently resistant to complete statistical modelling.

In Varela's theory of embodied cognition, cognition emerges through the dynamic coupling of organism and environment. Transposed into the AI-human creative ecosystem, desire operates as an embodied variable that feeds unpredictability into the generative process, introducing what we term "controlled noise" – deviations that restore diversity to the system's outputs.

When human creative work stems from authentic desire – whether driven by curiosity, critique, or aesthetic exploration – it generates variations that challenge and expand AI's learned distributions through open-ended exploration. This diversity is essential for counteracting Model Autophagy Disorder and other convergence effects. Conversely, when desire is channelled into algorithmically reinforced patterns (e.g., through recommendation systems prioritising engagement), it risks being commodified into predictability. In such cases, desire no longer disrupts entropy but accelerates it, reducing the system's resilience to creative collapse.

This dynamic suggests a symbiotic (or metabolic) relationship where humans and AI systems depend on each other for creative vitality. Humans provide the unpredictable variations and contextual meaning that prevent creative entropy, while AI systems offer computational power and pattern recognition capabilities that can amplify and transform human creative impulses in novel ways. In a poetic sense, it can be said that humans create history through their desire-driven explorations, while AI processes and amplifies that history into new creative possibilities.

CREATIVE METABOLISM

To understand the complex relationships between human creativity and AI systems, this analysis reframes contemporary phenomena through extraction rather than creation. This perspective builds on Rao's (2023) characterisation of AI as "camera, not engine" – technology that captures and processes existing information rather than generating genuinely new knowledge or expression.

Drawing on systems theory (Von Bertalanffy 1968) and autopoiesis (Maturana and Varela 1980), we conceptualise this process as a sequence of extraction and transformation cycles in which human-generated cultural material becomes the "nutrient flow" that sustains algorithmic production.

At the uppermost level, neural networks extract intelligence from human-generated data, transforming millions of creative works, conversations, and cultural artefacts into mathematical representations suitable for computational manipulation. This extraction separates creative expressions from their human origins, converting them into statistical patterns available for algorithmic recombination.

These datasets – which subsequently train neural networks – are themselves extracted from user engagement at intermediate levels. Creative platforms, social media networks, and AI services harvest behavioural data about user interactions with generated content, creating rich datasets about creative preferences, aesthetic choices, and cultural trends (Zuboff 2019). This data extraction transforms user behaviour into structured information suitable for algorithmic processing.

Finally, at the foundational level, user engagement is extracted from human intentions through attention economy mechanisms. Recommendation algorithms, content moderation systems, and engagement metrics create feedback loops shaping what content gets created and consumed, effectively monetising human attention and emotional investment (Gillespie 2018).

This multi-layered extractive system constitutes what we call creative metabolism – processes where human urge serves as the fuel that is consumed, processed, and transformed through successive algorithmic mediation layers. Each layer extracts value while transforming raw human creativity into increasingly abstracted forms suitable for further processing.

The Erosion of Desire

Critical threats to creative metabolism emerge when human desires themselves become commodified, homogenised, and trapped within algorithmic feedback loops. When AI systems optimise for engagement metrics, they reinforce existing preferences rather than challenging or expanding them, creating possible metabolic inefficiency where the essential unpredictability of human impulses becomes channelled into predictable patterns.

The phenomenon of “filter bubbles” and “echo chambers” may also represent a kind of metabolic breakdown where diversity of human wishes – the very quality that makes it valuable for preventing creative entropy – gets eliminated in favour of commercially optimised homogeneity, reinforcing existing preferences rather than challenging them (Pariser 2011). When platforms prioritise engagement over exploration, they create closed loops that reduce rather than expand creative possibilities.

These reductions in desire variety ultimately limit capacity to introduce the “controlled noise” that is necessary for preventing entropic degradation in AI systems. When human desire becomes perfectly predictable – perfectly aligned with algorithmic expectations – it can no longer serve as a source of genuine novelty.

Maintaining a healthy creative metabolism therefore requires protecting spaces where authentic human activities can evolve outside algorithmic mediation. This means preserving cultural contexts, artistic traditions, and experimental practices that operate outside corporate platforms with their content moderation and optimisation logics.

CREATIVE COUNTERCURRENT

While corporate laboratories dominate narratives around generative AI development, parallel ecosystems challenge these mainstream approaches from multiple directions. Established institutional fields – particularly academic research and arts – may also play crucial roles in fostering non-mainstream AI formations that resist commercial optimisation pressures.

University research laboratories are encountering mounting industry pressures that steadily shape their AI research agendas. The academic sector has experienced an unprecedented exodus of faculty to industry positions, alongside a growing corporate presence at scholarly conferences. Nevertheless, these environments maintain distinct research priorities, and academic researchers retain some autonomy to pursue fundamental research that may not align with immediate commercial interests, including work on interpretable AI, bias detection, and methodological approaches that prioritise broader societal considerations (Gofman and Jin 2024).

Similarly, art institutions face significant challenges in evaluating, collecting, and displaying contemporary AI art forms, requiring new technical infrastructure and curatorial expertise to support real-time, generative, and interactive features. However, these established contexts have increasingly provided space for sophisticated artistic engagement with AI technologies, where artists address the critical, conceptual uses of AI tools beyond their mainstream commercial applications. This support enables more experimental approaches that distinguish serious artistic practice from commercial illustrations or collectibles (Paul 2023).

Besides available institutional efforts, parallel ecosystems of independent researchers and open- source collectives have expanded the medium's expressive potential through "guerrilla research:" practices prioritising low-cost, improvisational experimentation over brute-force computational power. This grassroots approach challenges fundamental assumptions about AI development, demonstrating how repurposed architectures, interdisciplinary hacks, and creative constraints can unlock novel aesthetic possibilities with minimal resources rather than requiring trillion-parameter models trained on massive datasets with enormous compute. They demonstrate that the frontier of AI expression lies not necessarily in scaling existing approaches but in discovering unexpected applications and combinations of more modest tools.

Together, these efforts create vital counterbalances to corporate AI development, ensuring that generative media remains a frontier for diverse exploration rather than a standardised commercial product. Where grassroots efforts offer agility and innovation, structured support provides sustainability and resources for deeper investigation into AI's creative and cultural implications.

Case Study: CLIP-VQGAN Revolution

Perhaps the most significant manifestation of crowd research occurred in early text-to-image generation methods representing fundamentally different paradigms from later commercial developments. These pioneering approaches were based not on heavy pretrained models like GANs but on iterative optimisation techniques using backpropagation underlying neural network mechanics generally.

When OpenAI published the CLIP (Contrastive Language-Image Pre-training) model on 5 January 2021, the company intended it primarily as a tool for classifying images according to textual labels (Radford et al. 2021). However, AI practitioners, beginning with Ryan Murdock (@advadnoun) and Katherine Crowson (@RiversHaveWings), quickly recognised that the model could be repurposed for text-to-image synthesis, a capability OpenAI had not explicitly designed or optimised for (Rorvig 2021).

By combining CLIP with various generative architectures like VQGAN (Vector Quantized Generative Adversarial Networks), these practitioners initiated transformative movements that explored text-guided image synthesis months before commercial tools like DALL-E were released publicly (Crowson et al. 2022). More importantly, it established aesthetic and methodological precedents influencing the entire field of AI art.

This grassroots innovation demonstrates several principles central to creative countercurrent analysis. First, it reveals how creative practitioners can discover unexpected applications for technologies by engaging with them experimentally rather than following intended use cases. Second, it shows how resource constraints can become creative assets rather than limitations, forcing practitioners to develop elegant solutions using available tools rather than requiring massive computational infrastructure.

Third, the CLIP-VQGAN phenomenon demonstrates the importance of open research culture where practitioners share

FIGURE 6. Examples of CLIP-VQGAN generations with their prompts.
Source: Crowson et al., VQGAN-CLIP: Open Domain Image Generation and Editing with Natural Language Guidance, 2022, p. 27 (preprint version).

FIGURE 7. Created by Author using Illustrip method (CLIP+RGB parameterisation) from Aphantasia project, 2022.

methods, code, and discoveries rather than maintaining proprietary control. This open approach enabled rapid iteration and the collective development of techniques that might have remained isolated innovations without community sharing practices.

Case Study: Aphantasia Project

One such CLIP-based project, Aphantasia, represents a particularly illustrative example of guerrilla research approaches (Epstein 2021). Named after the neurological condition characterised by inability to form mental imagery, Aphantasia coined the text-to-image approach as "outsourced visual imagination." Users input textual prompts – perhaps titles of paintings, novels, or musical compositions – and the system generates distilled visual interpretations externally manifesting associative processes typically hidden within the mind.

Initiated months before the celebrated CLIP-VQGAN combination and first creative experiments with diffusion models that would eventually evolve into platforms like Midjourney and Stable

Diffusion, Aphantasia charted distinctive paths (Computer Vision Art Gallery 2021, Steinfeld 2023). The project developed unique visual language of layered, dreamlike imagery evoking mediaeval tapestries colliding with glitch art. It emerged directly from the technical approach rather than an attempt to mimic existing artistic styles or follow mainstream AI art developments.

Technically, the method combined OpenAI's CLIP with Fast Fourier Transform (FFT) parameterisation or direct optimisation of RGB pixels, instead of employing heavy pretrained models like GANs as drawing engines. This approach achieved virtually unlimited resolution and made the tool highly effective for creating pan/zoom/extrude animations. The project showcased how technical constraints of resource scarcity could be transformed into novel creative methods – a characteristic of independent research.

Case Study: ConceptLab Project

Similarly, academic research project ConceptLab extends CLIP-based synthesis beyond conventional imagery to explore entirely novel concepts (Richardson et al. 2024). The project employs a simple but elegant iterative technique to generate hybrid objects and entities that belong to established categories yet do not fit neatly into any particular subcategory. Rather than simply recombining familiar elements, ConceptLab pushes AI systems to expand conceptual vocabularies of so-called tokens (semantic units) by probing under-explored regions of the conceptual latent space between recognised concepts. This has allowed the generation of outputs with no pre-existing referent, suggesting forms that might have never emerged from the human imagination.

Case Study: Active Divergence et al.

Terence Broad's research on "Active Divergence" (a term coined by Sebastian Berns) presents a comprehensive survey and taxonomy of frameworks for subverting generative models to produce unexpected results, exploring the concept of creative destruction in various areas (Broad et al.

2021). Rather than using AI tools as intended, these approaches intentionally manipulate generative models to produce outputs diverging from training data distributions or eliminating the need for conventional datasets entirely through data-free generation methods (Computer Vision Art Gallery 2019).

"Network bending" techniques involve deliberately interfering with low-level neural network operations to force unpredictable behaviours and outputs that no amount of prompt engineering or conventional AI usage could achieve. This approach treats AI models not as black boxes to be optimised, but as complex systems to be explored, hacked, and artistically manipulated.

Such methods serve multiple purposes including novelty generation for truly unprecedented outputs, creativity support tools for artists, and adaptation to domains

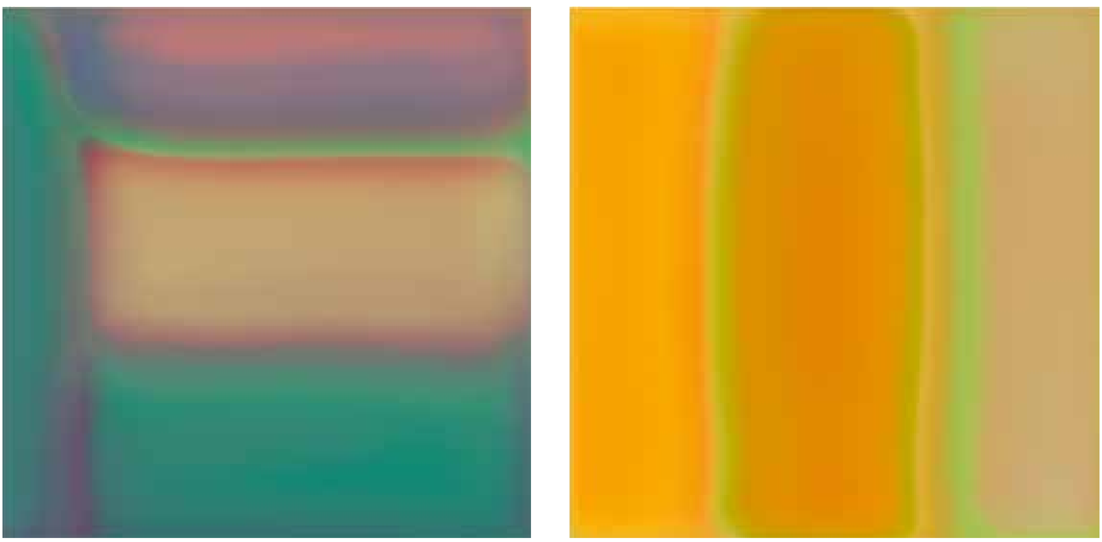

FIGURE 8. Created by Author using ConceptLab method within own codebase, 2024.

FIGURE 9. Terence Broad, "(un)stable equilibrium 1:1": example of data-free generation methods, with two generator networks competing to each other within GAN architecture. Used with permission. 2019.

FIGURE 10. Roope Rainisto, "Smoke #6" from "Vacation" collection. Used with permission. Source: https://www.rooperainisto.com, 2024

where limited training data exists. While the methods themselves may be familiar to artists, systematic scholarly research proves crucial for the field, as it legitimises and popularises the very vector of these counterintuitive and seemingly non-productive explorations.

Case Study: Art of Roope Rainisto

The generative art landscape encompasses numerous groundbreaking experimental artists from both AI and pre-AI eras, falling under the creative countercurrent concept, including pioneers and visionaries such as Mario Klingemann and Alexander Mordvintsev. However, comprehensive coverage of these practices would require dedicated research beyond this study's scope. Instead, we focus on one particularly indicative edge case that effectively illuminates key tensions.

The art of Finnish artist Roope Rainisto provides a clear example of developing distinctive AI art practices that deliberately diverge from commercial optimisation toward photorealistic outputs, while being rooted in established art traditions rather than research experiments (Cartagena 2023). His approach to "post-photography" demonstrates a clear alternative to mainstream AI development – an alternative that celebrates rather than conceals AI's inherent characteristics and native visual language. Rainisto's works fit perfectly into the category of synthetic folklore, representing deliberate exploration of obscure directions that many conventional artists avoid in favour of collapsing into conventional polished pseudo-realism.

Rainisto is far from an outsider. With collections selling out in minutes and gaining significant market recognition, he demonstrates how independent artists can develop distinctive practices that enrich cultural diversity rather than contribute to aesthetic convergence, resisting the homogenising pressures of algorithmic optimisation.

This approach proves especially relevant for creative countercurrent analysis because it shows how artists can produce aesthetically compelling and widely acclaimed work while maintaining critical distance from dominant AI trends and platforms, achieving commercial viability without compromising their experimental character or critical potential. Art history has repeatedly demonstrated this pattern within other areas, and AI art is no exception.

THE ROLE OF MARGINALITY

These case studies reveal an important principle: marginal or counter-hegemonic practices often provide disproportionately valuable contributions to generative landscapes. The further artistic practices deviate from dominant patterns, the more they contribute to preventing convergence and algorithmic stagnation. This principle has profound implications for structuring relationships between human expression and AI systems.

Rather than seeking to optimise AI tools to satisfy average user preferences, sustainable cultural environments may need to prioritise tools and practices amplifying marginal voices and unexpected directions. This approach contradicts conventional optimisation logic, which prioritises user satisfaction and engagement metrics, in favour of maintaining the health of the creative ecosystem through diversity preservation. In the metabolic metaphor, marginality can be compared to the immune system: it detects and counteracts homogenising tendencies before they dominate the field.

Independent practitioners serve as vital counterbalances to optimisation-focused corporate development. Their willingness to embrace instability, explore unusual aesthetic territories, and experiment with unconventional approaches provides essential "controlled noise" that keeps creative ecosystems healthy and diverse.

This suggests that protecting and nurturing authentic human expression becomes a structural requirement and a matter of cultural and technological policy rather than a purely aesthetic concern. This might involve platform design that gives precedence to exploration over engagement, educational approaches that emphasise risk-taking over optimisation, and economic models that support experimental work.

Cultural and Technical Implications

The metabolic framework and marginality principle suggest several implications for AI-mediated contexts and development practices. First, training data curation should prioritise diversity and edge cases rather than optimising for mainstream preferences. Current approaches to dataset filtering that remove "low-quality" or unusual content may inadvertently eliminate precisely those elements necessary for preventing creative entropy.

Second, platform regulation should consider preserving diversity as a public interest concern rather than leaving such decisions to commercial optimisation logic. This might require platforms and the corresponding authorities to maintain spaces for experimental content or to implement recommendation algorithms that actively promote creative diversity rather than optimising purely for conventional engagement. Such an approach might even involve supporting experimental uses that violate intended application scenarios (within reasonable safety precautions), while discouraging uses that reinforce existing biases or aesthetic patterns.

In particular, evaluation metrics for creative AI systems should include diversity and novelty measures rather than focusing solely on user satisfaction or aesthetic appeal. AI development could also benefit from incorporating feedback mechanisms that actively seek out marginal voices and unconventional applications rather than optimising for broad user adoption.

Third, educational institutions should develop curricula that emphasise creative experimentation and critical engagement with AI systems rather than training students to use AI tools effectively for commercial purposes. This educational approach would promote the kind of experimental practice represented in the case studies while developing critical literacy about AI's cultural implications.

CONCLUSION

Throughout this exploration of AI's role in content production, we have traced the contours of a complex systemic relationship between human creativity and algorithmic processes. This relationship encompasses simultaneous tendencies towards convergence and divergence, entropy and renewal, extraction and creation. Understanding how to maintain balance within this dynamic represents one of the central challenges facing artists, technologists, and cultural institutions as AI becomes increasingly integrated into our practices and perceptual environments. The metabolic framework developed here contributes to digital media studies by providing systems-level analysis of human-AI collaboration that moves beyond individual output assessment to examine ecological conditions necessary for sustainable cultural practices.

The promise lies in AI's capacity to amplify human expressive potential, democratise access to sophisticated creative tools, and generate genuinely novel aesthetic experiences. Experimental art and research demonstrate that these benefits need not depend on corporate control or massive computational resources.

The risks centre on AI's tendency towards homogenisation, the commoditisation of human creativity, and the potential for recursive degradation when AI systems train on their own outputs. These risks threaten not only aesthetic diversity but the fundamental sustainability of AI-mediated context and creative culture itself.

The metabolic metaphor ultimately suggests that sustainable human-AI collaboration requires ongoing circulation of authentic human desire and inventive unpredictability. Moving forward, the challenge is not in optimising AI systems for efficiency or commercial viability, but in supporting marginal voices, protecting human agency in creative processes, and designing AI systems that amplify rather than replace human creativity.

As we navigate this transition, we must remember that creativity itself resists optimisation. The most valuable contributions often emerge from unexpected places, marginal perspectives, and deliberately inefficient processes. Our challenge is to harness AI's remarkable capabilities while preserving the beautiful inefficiencies that make human culture vital, diverse, and sustainable.

FUTURE RESEARCH: AGENTIC AI AND EVOLVED METABOLISM

Recent developments toward agentic AI systems – models designed with autonomous decision- making capabilities enabling direct environmental interaction – represent a significant evolution beyond the pattern recognition and generative capabilities examined in this analysis. These agentic systems aim to overcome static training data limitations by actively exploring and manipulating the world as an extended playground (Park et al. 2023).

Pioneering strategies for autonomous creative practice are already being explored in technological art. Botto, created by Mario Klingemann (Klingemann et al. 2022), representing perhaps the most developed example of a "decentralised autonomous artist," has generated over $4 million in sales since 2021. Botto operates through a sophisticated community governance model where stakeholder voting data trains the AI's aesthetic model, creating a feedback loop in which collective human input shapes Botto's creative development while maintaining operational autonomy.

Gene Kogan's Abraham/Eden project represents another approach to autonomous creative agency. Eden provides infrastructure for building "autonomous creative agents" capable of understanding and creating visual content through natural conversation (Kogan 2023), while Abraham was conceived as a fully decentralised autonomous artificial artist that would operate as a "crowd- sourced AI" generating art through collective intelligence (Kogan 2019). These projects explore how creative autonomy might emerge from distributed human input rather than individual programming.

Furthermore, the ongoing "Assembly" art project explores multimedia narrative developments with semi-autonomous, self-controlled AI agents interacting within shared environments (Epstein 2025), while the elegantly conceived "Needle in a Haystack" installation, created by Varvara & Mar, uses a robot searching for a hidden needle as a strong metaphor for AI's computational limitations in processing vast datasets (Guljajeva and Canet Sola 2025).

Researchers document emergent behaviours in agentic systems – unexpected capabilities and interaction patterns – which many interpret as promising indicators of progress toward more sophisticated artificial intelligence forms. These developments allegedly attempt to replicate human creativity's open-ended and exploratory nature while potentially addressing some limitations identified in pre-agentic systems.

However, these simulations, no matter how sophisticated, still lack grounding in personal, large- scale embodied experiences shaping human creativity – the lived, contextual understanding emerging from navigating physical and social worlds with human bodies and perspectives. This gap reinforces critical questions: Can synthetic curiosity truly compete with natural human exploration? Do agentic expansions alter the creative metabolism framework described throughout this analysis, or does this framework continue imposing fundamental constraints on artificial systems?

Future research should investigate how agentic AI systems interact with creative metabolism principles, whether they develop novel forms of creative sustainability, or whether they remain fundamentally dependent on human creative input despite increased autonomy. Studies tracking aesthetic evolution in agentic systems could provide crucial data about whether autonomous exploration can substitute human desire in preventing creative entropy.

FIGURE 11. Vadim Epstein, "Assembly" in the storytelling mode. Still frame from the recorded session. Used with permission. 2025.

REFERENCES

Acemoglu, Daron; Autor, David; Hazell, Jonathon; Restrepo, Pascual 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40, no. S1 (2022): S293- S340. https://doi.org/10.1086/718327.

Agamben, Giorgio 2009. *What Is an Apparatus? And Other Essays*. Stanford: Stanford University Press. https://doi.org/10.1515/9781503600041.

Alemohammad, Sina; Casco-Rodriguez, Josue; Luzi, Lorenzo; Humayun, Ahmed Imtiaz; Babaei, Hossein; LeJeune, Daniel; Siahkoohi, Ali; Baraniuk, Richard 2024. "Self-Consuming Generative Models Go MAD." In *Proceedings of the International Conference on Learning Representations (ICLR 2024)*. https://openreview.net/forum?id=ShjMHfmPs0.

Beaty, Roger E.; Silvia, Paul J.; Nusbaum, Emily C.; Jauk, Emanuel; Benedek, Mathias 2014. "The Roles of Associative and Executive Processes in Creative Cognition." *Memory & Cognition* 42 (7): 1186–1197. https://doi.org/10.3758/s13421-014-0428-8.

Beer, Stafford 1972. *Brain of the Firm: The Managerial Cybernetics of Organization*. London: Allen Lane.

Belkin, Mikhail; Hsu, Daniel; Ma, Siyuan; Mandal, Soumik 2019. "Reconciling modern machine learning practice and the bias-variance trade-off". *Proceedings of the National Academy of Sciences*, 116(32), 15849–15854. https://doi.org/10.1073/pnas.1903070116.

Bellio, Naiara; Kayser-Bril, Nicolas 2023. "Some Image Generators Produce More Problematic Stereotypes than Others, but All Fail at Diversity." *AlgorithmWatch*. Accessed August 12, 2025. https://algorithmwatch.org/en/image-generators-stereotypes-diversity.

Bianchi, Federico; Kalluri, Pratyusha; Durmus, Esin; Ladhak, Faisal; Cheng, Myra; Nozza, Debora; Hashimoto, Tatsunori; Jurafsky, Dan; Zou, James; Caliskan, Aylin 2023. "Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale." In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1493–1504. https://doi.org/10.1145/3593013.3594095.Boden, Margaret A. 2004. *The Creative Mind: Myths and Mechanisms*. 2nd ed. London: Routledge. https://doi.org/10.4324/9780203508527.

Bolt, Barbara 2007. "The Magic Is in Handling." In *Practice as Research: Approaches to Creative Arts Enquiry*, edited by E. Barrett and B. Bolt, 27–34. London: I.B. Tauris.

Broad, Terence; Berns, Sebastian; Colton, Simon; Grierson, Mick 2021. "Active Divergence with Generative Deep Learning – A Survey and Taxonomy." In *Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)*, 227–236. Association for Computational Creativity. https://computationalcreativity.net/iccc21/wp- content/uploads/2021/09/ICCC_2021_paper_43.pdf.

Bucher, Taina 2018. *If...Then: Algorithmic Power and Politics*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780190493028.001.0001.

Buick, Adam 2025. "Copyright and AI training data – transparency to the rescue?" *Journal of Intellectual Property Law & Practice* 20, no. 3 (2025): 182–192. https://doi.org/10.1093/jiplp/jpae102.

Cartagena, Alejandro 2023. "In Conversation: Roope Rainisto on Post-photography and REWORLD." *Post Photographic Perspectives* (blog). Accessed August 12, 2025. https://postphotography.xyz/ppp1/in-conversation.

Colton, Simon; Wiggins, Geraint A. 2012. "Computational Creativity: The Final Frontier?" In *Proceedings of the 20th European Conference on Artificial Intelligence*, 21-26. Amsterdam: IOS Press. https://doi.org/10.3233/978-1-61499-098-7-21.

Computer Vision Art Gallery 2019. "(un)stable equilibrium 1:1." *ICCV 2019 Computer Vision Art Gallery*. Accessed August 12, 2025. https://computervisionart.com/pieces2019/unstable- equilibrium.

Computer Vision Art Galler. 2021. "Aphantasia." *CVPR 2021 Computer Vision Art Gallery*, March 1, 2021. Accessed August 12, 2025. https://computervisionart.com/pieces2021/aphantasia.

Crowson, Katherine; Biderman, Stella; Kornis, Daniel; Stander, Dashiell; Hallahan, Eric; Castricato, Louis; Raff; Edward 2022. "VQGAN-CLIP: Open Domain Image Generation and Editing with Natural Language Guidance." In *Computer Vision – ECCV 2022*, 88–105. Lecture Notes in Computer Science, vol. 13697. Cham: Springer. https://doi.org/10.1007/978-3-031-19836-6_6.

Epstein, Vadim 2021. *Aphantasia*. GitHub repository. Accessed August 12, 2025. https://github.com/eps696/aphantasia.

Epstein, Vadim 2025. *Assembly*. Artist website. Accessed August 12, 2025. https://wesual.art/assembly.

Epstein, Ziv; Hertzmann, Aaron; Investigators of Human Creativity, Akten, Memo; Farid, Hany; Fjeld, Jessica; Frank, Morgan R., et al. 2023. "Art and the Science of Generative AI." *Science* 380, no. 6650 (2023): 1110–1111. https://doi.org/10.1126/science.adh4451.

FITC 2023. "Gene Kogan." FITC Speaker Profile, April 2023. Accessed August 2025. https://fitc.ca/speaker/gene-kogan/?event=33481%2F.

Flusser, Vilém (1983) 2011. *Towards a Philosophy of Photography*. London: Reaktion Books.

Flusser, Vilém (1985) 2014. *Into the Universe of Technical Images*. Minneapolis: University of Minnesota Press. https://doi.org/10.5749/minnesota/9780816670208.003.0021.

Gillespie, Tarleton 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. New Haven, CT: Yale University Press. http://doi.org/10.12987/9780300235029.

Gofman, Michael; Jin, Zhao 2024. "Artificial Intelligence, Education, and Entrepreneurship." *The Journal of Finance* 79, no. 1 (2024): 631–667.

Gray, Carole; Malins, Julian 2004. *Visualizing Research: A Guide to the Research Process in Art and Design*. Aldershot: Ashgate. https://doi.org/10.1017/S0307472200014498.

Guljajeva, Varvara; Canet Sola, Mar 2025. "A Needle in a Haystack. " *SocArXiv*. May 23. https://doi.org/10.31235/osf.io/5sn67_v1.

Hertzmann, Aaron 2018. "Can Computers Create Art?" *Arts* 7 (2): 18. https://doi.org/10.3390/arts7020018.

Klingemann, Mario; Hudson, Simon; Epstein, Zivvy 2022. "Botto: A Decentralized Autonomous Artist." In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS 2022).* https://neuripscreativityworkshop.github.io/2022/papers/ml4cd2022_paper13.pdf.

Kogan, Gene 2019. "Artist in the Cloud: Towards an Autonomous Artist." In *NeurIPS Machine Learning for Creativity and Design Workshop*, December 14, 2019, Vancouver, Canada.

Kogan, Gene 2023. "Provenance for Generative AI." *Medium* (blog). Accessed August 12, 2025. https://medium.com/@genekogan/provenance-on-eden-d80dacce572a.

Krapp, Peter 2011. *Noise Channels: Glitch and Error in Digital Culture*. Minneapolis: University of Minnesota Press. https://doi.org/10.5749/minnesota/9780816676248.001.0001.
Maturana, Humberto R.; Varela, Francisco J. 1980. *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel Publishing Company. https://doi.org/10.1007/978-94-009-8947-4.
Menkman, Rosa 2011. *The Glitch Moment(um)*. Amsterdam: Institute of Network Cultures. http://doi.org/10.25969/mediarep/19296.
Paglen, Trevor 2019. "Invisible images: Your pictures are looking at you." *Architectural Design*, 89(1), pp. 22–27. https://doi.org/10.1002/ad.2383.
Pariser, Eli 2011. *The Filter Bubble: What the Internet Is Hiding from You*. New York: Penguin Press.
Park, Joon Sung; O'Brien, Joseph; Cai, Carrie Jun; Morris, Meredith Ringel; Liang, Percy; Bernstein, Michael S. 2023. "Generative Agents: Interactive Simulacra of Human Behavior." In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. ACM. https://doi.org/10.1145/3586183.3606763.
Paul, Christiane 2023. *Digital Art*. 4th ed. London: Thames & Hudson.
Radford, Alec; Kim, Jong Wook; Hallacy, Chris; Ramesh, Aditya; Goh, Gabriel; Agarwal, Sandhini; Sastry, Girish et al. 2021. "Learning Transferable Visual Models from Natural Language Supervision." In *Proceedings of the 38th International Conference on Machine Learning*, 8748–8763. PMLR. http://proceedings.mlr.press/v139/radford21a.html.
Rao, Venkatesh 2023. "A Camera, Not an Engine." *Contraptions* (blog). Accessed August 12, 2025. https://contraptions.venkateshrao.com/p/a-camera-not-an-engine.
Rao, Venkatesh 2024. "Imagination vs. Creativity." *Ribbonfarm Studio* (blog). Accessed August 12, 2025. https://ribbonfarm.com/2024/07/14/imagination-vs-creativity.
Richardson, Elad; Goldberg, Kfir; Alaluf, Yuval; Cohen-Or, Daniel 2024. "ConceptLab: Creative Concept Generation Using VLM-Guided Diffusion Prior Constraints." *ACM Transactions on Graphics* 43 (3): Article 34. https://doi.org/10.1145/3659578.
Rorvig, Mordechai 2021. "AI-Generated Art Scene Explodes as Hackers Create Groundbreaking New Tools." *VICE*. Accessed August 12, 2025. https://vice.com/en/article/ai-generated-art-scene- explodes-as-hackers-create-groundbreaking-new-tools.
Runco, Mark A.; Jaeger, Garrett J. 2012. "The Standard Definition of Creativity." *Creativity Research Journal* 24 (1): 92–96. https://doi.org/10.1080/10400419.2012.650092.
Russell, Legacy 2020. *Glitch Feminism: A Manifesto*. London: Verso Books.
Simondon, Gilbert 2017. *On the Mode of Existence of Technical Objects*. Minneapolis: University of Minnesota Press.
Steinfeld, Kyle 2023. "Clever Little Tricks: A Socio-Technical History of Text-to-Image Generative Models." *International Journal of Architectural Computing* 21(2): 211–241. https://doi.org/10.1177/14780771231168230.
Sullivan, Graeme 2005. *Art Practice as Research: Inquiry in the Visual Arts*. Thousand Oaks, CA: Sage Publications.
Swanson, Steph Maj 2022. *Loab: The AI Cryptid* (blog). Accessed August 12, 2025. https://loab.ai. Thompson, Evan. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
Varela, Francisco J.; Thompson, Evan; Rosch, Eleanor 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/6730.001.0001.
Von Bertalanffy, Ludwig 1968. *General System Theory: Foundations, Development, Applications*. New York: George Braziller.
Vygotsky, Lev Semenovich (1930) 2004. "Imagination and Creativity in Childhood." Translated by M. E. Sharpe. *Journal of Russian & East European Psychology* 42 (1): 7–97. https://doi.org/10.1080/10610405.2004.11059210.
Wiener, Norbert 1948. *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.
Winkler, Christoph 2025. "Whispers of Wood: A Prompt-based Dance Piece." *Company Christoph Winkler* (blog). Accessed August 12, 2025. https://christoph-winkler.com/en/productions/whispers- of-wood.
Zhou, Kaiyang; Liu, Ziwei; Qiao, Yu; Xiang, Tao; Loy, Chen Change 2021. "Domain Generalization: A Survey." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44 (9): 5396-5415. https://doi.org/10.1109/TPAMI.2022.3195549.
Zuboff, Shoshana 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.