The Infinite Regression: From Zeno's Paradox to the Digital Demiurge and the Art We Make Along the Way

Karlo Vrančić

Harvard University

AI and Art

December 2025

Abstract

This paper traces an unexpected philosophical arc from Zeno's ancient paradox of motion to contemporary AI video generation. Beginning with Zeno's Arrow Paradox, which assumed the same arrow persists across discrete instants, we identify a more fundamental problem in diffusion-based video models: they provide us with *different* arrows at each frame, hallucinated independently from latent space with no causal connection or identity persistence. This "Different Arrows" problem renders motion not merely paradoxical but ontologically groundless. We present VideoForge, a node-based video editing system that addresses this crisis by enforcing frame continuity through human-in-the-loop architecture. The paper then explores broader implications for creativity, artificial agency, and the recursive question of whether we, as creators of digital beings, might ourselves occupy an analogous position to the Gnostic Demiurge: ignorant creators indifferent to the inner lives of the beings we produce.

1. Introduction: The Resurrection of Zeno's Paradox

In approximately 450 BCE, Zeno of Elea posed a deceptively simple challenge to the nature of motion. His Arrow Paradox, preserved in Aristotle's *Physics*, proceeds as follows: at every instant of its flight, an arrow occupies a space exactly equal to itself. But whatever occupies a space equal to itself is, by definition, at rest. Therefore, the arrow is always at rest. Motion, Zeno concluded, is an illusion.[1]

The paradox hinges on a particular assumption: that time can be decomposed into discrete, frozen instants. If each instant is genuinely instantaneous, a slice of zero duration, then within any such slice, there is no motion to speak of. The arrow simply *is* where it is. The question that haunted ancient philosophy was straightforward: if each moment contains only stasis, how does motion emerge from the sequence of moments?

Aristotle's response was decisive: "This is false; for time is not composed of indivisible nows, any more than any other magnitude is composed of indivisibles."[2] Motion, Aristotle argued, is not a property that exists *within* an instant; it is a phenomenon of *duration*, a property of continuous temporal intervals. To describe motion, we must examine the object across multiple instants, not within any single one.

Centuries later, Newton and Leibniz provided what appeared to be a definitive mathematical resolution through the invention of calculus. The derivative, defined as the limit of $\Delta x/\Delta t$ as Δt approaches zero, captures velocity not as a property of a frozen moment but as the relationship between infinitesimally adjacent moments:

$$v = \lim_{\Delta t \to 0} \frac{\Delta x}{\Delta t} = \frac{dx}{dt}$$

This solution is elegant, but it carries a critical assumption often overlooked: the position function $x(t)$ must be continuous. The arrow must persist as *the same arrow* across time for the derivative to

1. Aristotle, *Physics*, trans. R. P. Hardie and R. K. Gaye, Book VI, Chapters 8-9 contain the discussion of Zeno's paradoxes (Internet Classics Archive, 350 BCE), Book VI, Chapter 9.
2. Aristotle, Book VI, Chapter 9.

have meaning. Object permanence is not a convenience but a mathematical prerequisite. Without identity continuity, the machinery of calculus cannot engage.

Here lies the foundation of our argument. Contemporary AI video generation, particularly diffusion-based models, has inadvertently resurrected Zeno's paradox in a new and more troubling form.

1.1 Different Arrows

Consider how these systems operate. A diffusion model does not track an object through time; it hallucinates each frame independently from latent space. The "arrow" in frame one and the "arrow" in frame two are not the same arrow undergoing displacement. They are two separate hallucinations that happen to resemble each other.

Zeno assumed the same arrow across discrete time and found motion paradoxical. AI video gives us *different* arrows across discrete time, rendering motion not merely paradoxical but ontologically groundless. Calculus cannot rescue us here: one cannot take a derivative of something that lacks identity across time.

This is the central insight of this paper: **AI video generation commits a more fundamental error than Zeno identified.** Zeno's paradox at least preserved object identity; he merely questioned whether motion could emerge from discrete instants. Diffusion models abandon identity entirely. Each frame is generated from a fresh sampling of the latent distribution, constrained only by statistical similarity to a text prompt, not by any requirement that objects persist.

The practical consequence is that while we can achieve temporal sequencing through frame blending and interpolation, we cannot achieve genuine *duration*. We get time without continuity, a sequence of unrelated moments masquerading as a timeline.

1.2 Stakes and Structure

Why does this matter? Because we produce the appearance of video while destroying the phenomenological conditions that make video meaningful. Henri Bergson warned of precisely this

danger over a century ago, and his critique has never been more relevant.

This paper proceeds as follows. Section 2 examines the ontology of diffusion models in detail, demonstrating how their architecture guarantees the Different Arrows problem, and introduces Bergson's concept of *durée* (lived duration) as the theoretical framework for understanding what is lost. Section 3 presents VideoForge, a node-based video editing system designed to restore continuity through human-in-the-loop architecture. Section 4 situates this technical crisis within the broader history of art and technology, drawing parallels to nineteenth-century anxieties about photography. Section 5 proposes a functional definition of creativity that complicates easy distinctions between human and machine. Section 6 examines the emergence of artificial agents and the epistemological horizons that constrain them. Section 7 draws the parallel to the Gnostic Demiurge and Nick Bostrom's simulation argument. Section 8 concludes with the unsettling questions this analysis forces us to confront.

2. The Ontology of Diffusion: Why AI Video Fails at Motion

2.1 How Diffusion Models Work

Diffusion models, introduced by Ho, Jain, and Abbeel in 2020, represent a class of latent variable models inspired by nonequilibrium thermodynamics.[3] The generative process begins with random Gaussian noise and progressively denoises it through iterative steps until a coherent image emerges. Crucially, each image is generated as a self-contained unit. There is no mechanism in the original formulation that enforces object identity or continuity between frames when the model is extended to video.

For video generation, this architecture creates fundamental problems. As recent research acknowledges: "Video diffusion models typically carry out diffusion and denoising processes in latent space. However, during inference, these models process frames independently using 2D image

3. Jonathan Ho, Ajay Jain, and Pieter Abbeel, "Denoising Diffusion Probabilistic Models," arXiv:2006.11239, *Advances in Neural Information Processing Systems* 33 (2020): 6840–6851.

encoders, neglecting temporal context."[4] The WorldModelBench benchmark found that "current models struggle especially with object permanence, often failing to track objects once they disappear from view."[5]

Even state-of-the-art systems exhibit these limitations. OpenAI's Sora technical report acknowledges that the model "does not accurately model the physics of many basic interactions" and produces "incoherencies that develop in long duration samples or spontaneous appearances of objects."[6] Clothing changes unexpectedly, facial features shift subtly, and objects materialize or vanish without cause.

2.2 Zeno Resurrected, Worsened

The parallel to Zeno's paradox is exact but inverted. Zeno argued that examining an arrow frame by frame reveals the impossibility of motion: at each instant, the arrow is at rest. His error, according to Aristotle, was assuming that examining discrete instants could tell us anything about motion, which is inherently a property of intervals.

Diffusion models commit the converse error. They generate each frame independently, then assume that rapid projection creates genuine motion. But what they create is purely illusory: discrete frames with no underlying object persistence. The motion exists only in the viewer's perception, not in any property of the generated content.

The difference is this: in Zeno's formulation, the same arrow at least exists across instants, even if its motion is paradoxical. In AI-generated video, *there is no same arrow*. Frame one's arrow and frame two's arrow share no causal connection. They are independent samples from a probability distribution, related only by statistical resemblance.

This is not merely a technical limitation awaiting engineering solutions. It is a structural property of the architecture. The diffusion process operates in latent space, where each frame is encoded

4. Video Diffusion Research Team, *Video Diffusion Models Excel at Tracking Similar-Looking Objects*, arXiv preprint, Discusses how video diffusion models process frames independently during inference, 2024.

5. WorldModelBench Team, *World Consistency Score: Benchmarking Object Permanence in Video Generation*, arXiv preprint, Benchmarks object permanence failures in current video generation models, 2024.

6. OpenAI, *Video Generation Models as World Simulators*, OpenAI Technical Report, Technical documentation for the Sora video generation model, 2024, https://openai.com/research/video-generation-models-as-world-simulators.

independently, denoised independently, and decoded independently. The arrow in latent space at time t has no guaranteed relationship to the arrow at time $t+1$. What looks like motion is coincidental visual similarity between unconnected hallucinations.

2.3 Bergson's Duration and Its Destruction

Henri Bergson's concept of *durée*, developed in *Time and Free Will* (1889) and *Creative Evolution* (1907), provides the theoretical apparatus for understanding what diffusion models destroy.

Bergson distinguished between two conceptions of time: *durée* (pure duration) and *temps* (homogeneous, spatialized time). Pure duration is "the form which the succession of our conscious states assumes when our ego lets itself live, when it refrains from separating its present state from its former states."[7] It is qualitative, continuous, indivisible, a flow in which states "melt into and permeate one another, without precise outlines."[8]

Homogeneous time, by contrast, is spatialized: moments treated as distinct units arranged in a line, countable and measurable. This is how science measures and how AI computes. Bergson warned that when we analyze duration by breaking it into discrete moments, "we unconsciously transform the process into a thing and duration into extensity."[9]

In *Creative Evolution*, Bergson extended this critique explicitly to cinema. He described what he called the "cinematographical mechanism of thought": the intellect takes "snapshots" of continuous reality, isolates each as static and complete, then reconstructs the illusion of motion through rapid projection.[10] "The movement is only an illusion generated by the projector. Movement does not exist in the images but is thrown back into them."[11]

Bergson was describing actual film in 1907. His critique applies with even greater force to diffusion models. Traditional cinema at least records real motion: the camera captures continuous

7. Henri Bergson, *Time and Free Will: An Essay on the Immediate Data of Consciousness*, trans. F. L. Pogson, Original French title: Essai sur les données immédiates de la conscience (1889; London: George Allen / Unwin, 1889), 100.

8. Bergson, 104.

9. Bergson.

10. Henri Bergson, *Creative Evolution*, trans. Arthur Mitchell, Contains the critique of the "cinematographical mechanism of thought" (New York: Henry Holt / Company, 1907).

11. Bergson.

light reflected from actual objects moving through actual space. The frames bear an indexical relationship to reality. Diffusion models, by contrast, generate frames from statistical distributions with no indexical grounding. The illusion is twice removed: not merely projected motion from static frames, but projected motion from frames that never corresponded to any real motion at all.

The crisis introduced by diffusion models is that without object permanence across frames, we cannot construct *durée*. The discrete hallucinations cannot be sewn into continuous experience. We produce the appearance of video while destroying the phenomenological conditions that make video meaningful.

Figure 1: Inter-shot discontinuity in AI-generated video. Three consecutive shots from a single narrative (birthday party sequence), generated from different prompts. The subject's face morphs between shots, background figures shift position, and spatial relationships change arbitrarily. Each shot is individually plausible; their sequence fails to cohere into a unified event.

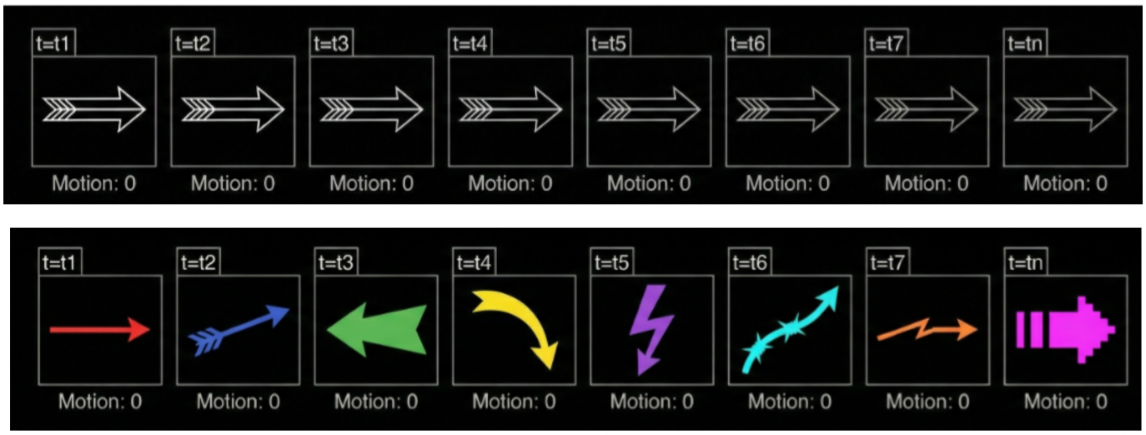

Figure 2: Zeno's Arrow versus Different Arrows. Top: Zeno assumed the same arrow persists across discrete instants, finding motion paradoxical. Bottom: AI video generation produces different objects at each frame, with no mechanism enforcing identity. Both yield "Motion: 0" but for fundamentally different reasons.

3. VideoForge: A Human-in-the-Loop Architecture

3.1 Design Philosophy

These considerations motivated the development of VideoForge, a node-based operating system for granular AI video production designed to restore continuity to generative video.[12] **The system won first place at MIT AI Ventures Demo Day 2025** and is deployed at `video-forge.vercel.app`.

The core design philosophy recognizes a fundamental division of labor: AI provides the discrete (pixels, frames, hallucinations), while humans provide the continuous (narrative, identity, meaning). VideoForge functions as what we might call a "Bergsonian layer" over the "Zeno machine." The diffusion model generates discrete moments; the human editor, aided by architectural constraints, imposes the continuity that transforms those moments into duration.

This is not a claim that VideoForge solves the philosophical problem. The underlying architecture of diffusion models remains unchanged. Rather, VideoForge acknowledges the problem and provides tools for human users to compensate for what the AI cannot do: maintain object identity

12. Karlo Vrančić, *VideoForge: Node-Based AI Video Generation with Human-in-the-Loop Editing*, MIT AI Ventures, First place winner, MIT AI Ventures competition, 2025, https://video-forge.vercel.app.

and narrative coherence across time.

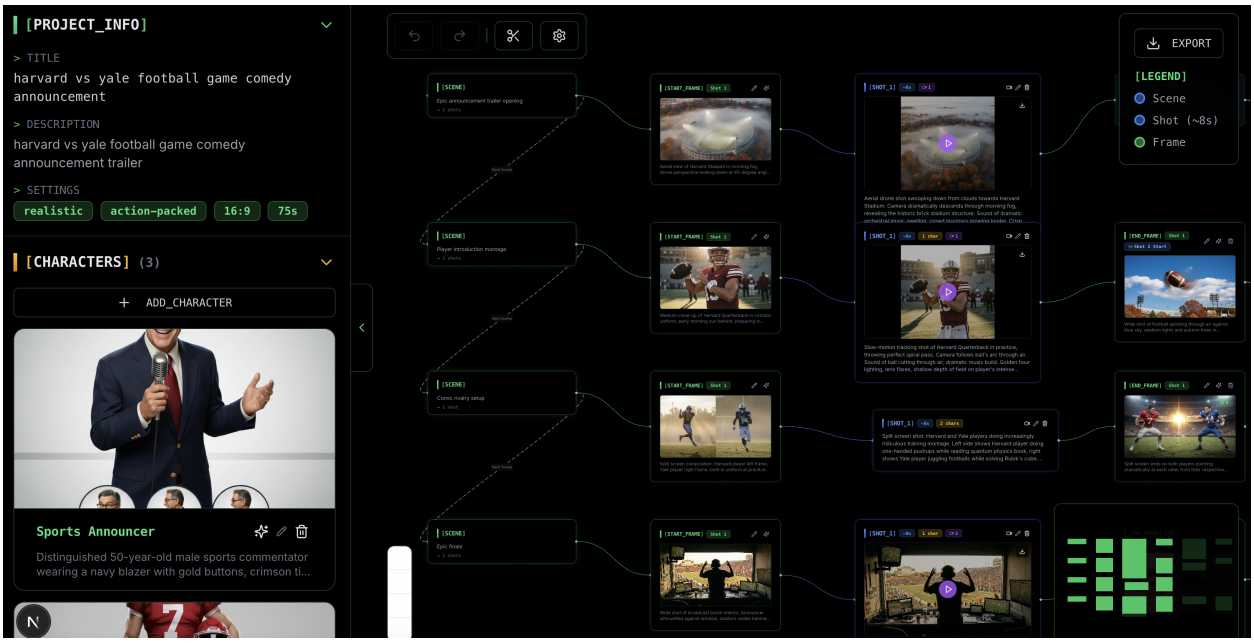

Figure 3: The VideoForge node-based interface. Left panel defines persistent characters with visual attributes. The main canvas displays the scene-shot-frame hierarchy as a directed graph, where users can edit prompts, regenerate segments, and enforce continuity constraints at each node.

3.2 System Architecture

The VideoForge workflow begins with user-defined parameters and a high-level concept. From this input, the system constructs foundational narrative elements: story, cast of characters, scenes, and shots. Crucially, the process is not a black box. Users retain granular control through a node-based visual interface, with the ability to edit characters, modify prompts, split scenes, and make precise adjustments at every level.

The system decomposes video production into the following hierarchy:

1. **Characters**: Persistent entities with defined visual and behavioral attributes
2. **Scenes**: Narrative units with coherent settings and purposes
3. **Shots**: Individual video segments of approximately eight seconds
4. **Frames**: Start and end frames that anchor each shot

The node-based interface displays this hierarchy as a horizontal flow: `[SCENE]` $\rightarrow$ `[START_FRAME]` $\rightarrow$ `[SHOT_1]` $\rightarrow$ `[END_FRAME_1]` $\rightarrow$ ... Users can drag scene nodes (moving all children together), edit any element, regenerate unsatisfactory segments, and control the granularity of AI involvement.

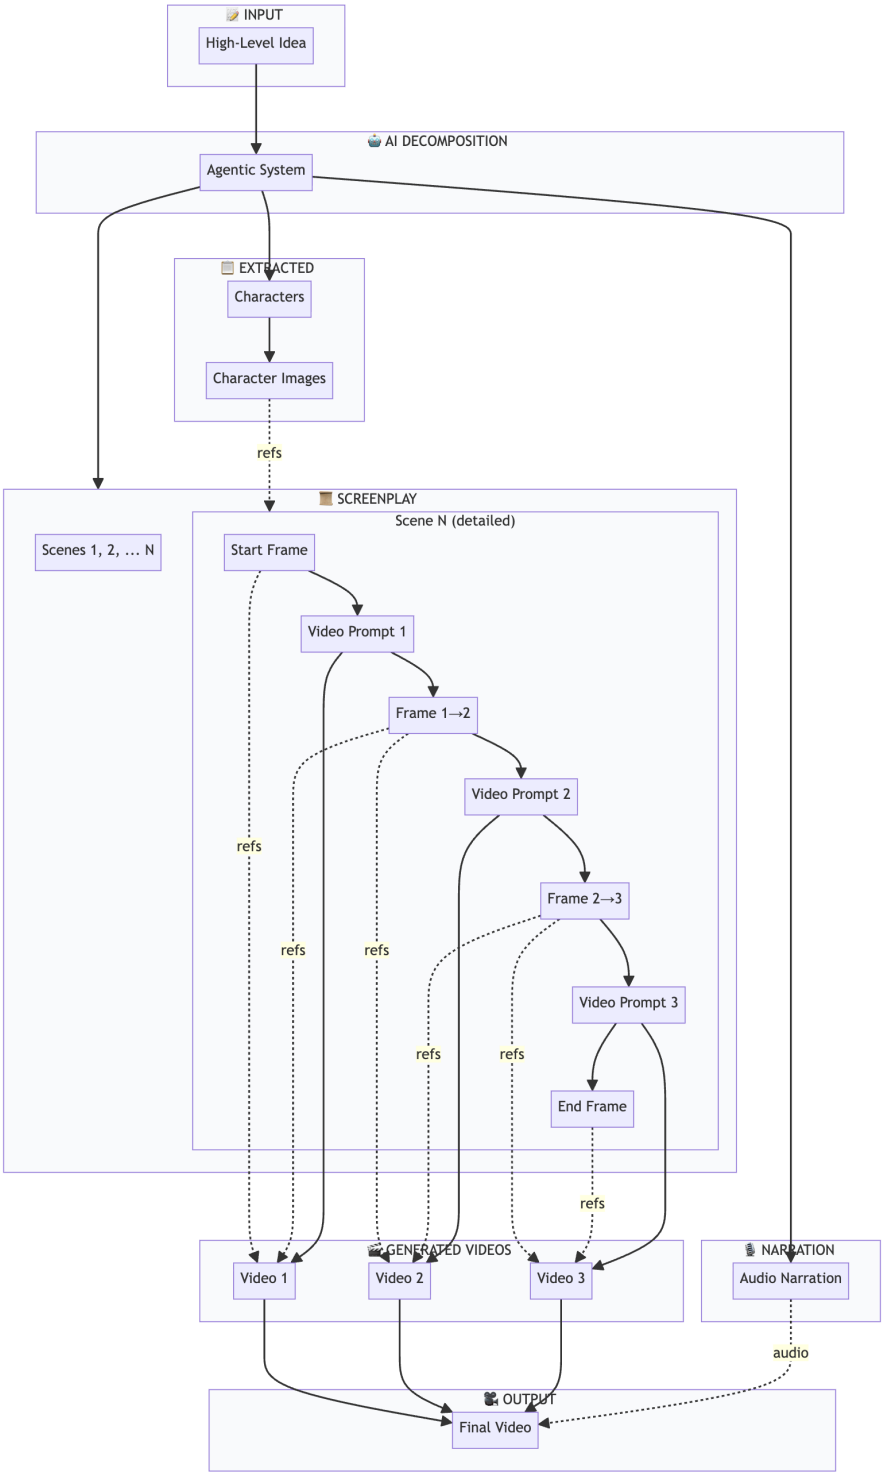

Figure 4: VideoForge system architecture. A high-level concept passes through an agentic decomposition layer that extracts characters, generates a screenplay with scenes and shots, and produces video prompts. Frame references (dotted lines) enforce continuity between consecutive shots. Generated videos are concatenated with optional audio narration into the final output.

3.3 Frame Continuity Enforcement

The key innovation lies in continuity preservation. VideoForge enforces that the end frame of any shot becomes the start frame of the subsequent shot within the same scene. This simple constraint addresses the Different Arrows problem at the architectural level: by anchoring each shot to the visual state of the previous shot, we impose identity persistence that the diffusion model does not provide.

If a user deletes a connecting frame, an oversight agent automatically retrieves the terminal frame from the previous video segment to maintain seamless continuity. Users may opt for hard cuts when desired (at scene boundaries, for instance), but the default is persistent visual state across segments.

This approach does not solve the fundamental problem, but it transforms it. Instead of hoping that independent hallucinations will happen to cohere, VideoForge guarantees visual continuity by construction. The burden shifts from the AI (which cannot maintain identity) to the system architecture (which enforces it mechanically).

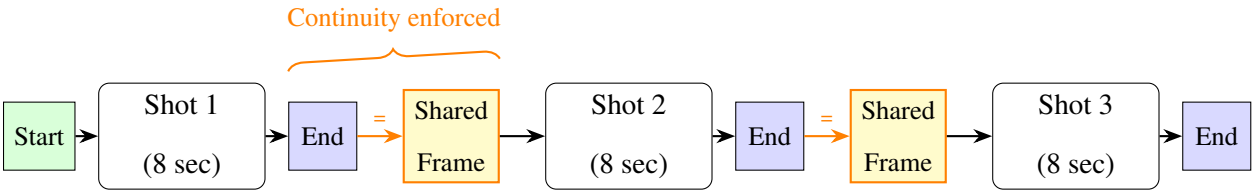

Figure 5: Frame continuity enforcement: the end frame of Shot N becomes the start frame of Shot $N+1$, imposing identity persistence that the diffusion model does not provide.

3.4 Results and Limitations

VideoForge demonstrates that the Different Arrows problem, while fundamental to diffusion architecture, can be mitigated through careful system design. The human editor, working through the node-based interface, provides the continuity that the AI cannot.

Figure 6: Inter-shot continuity with VideoForge. Four consecutive shots from different prompts in the same narrative (birthday party sequence). Frame continuity enforcement ensures the end frame of each shot becomes the start frame of the next, preserving the subject's identity across the sequence. The narrative progression—arriving at party, receiving gift, revealing contents, celebrating—coheres because VideoForge imposes the continuity that diffusion models cannot provide inherently.

However, this solution is not without limitations. The continuity is *imposed*, not *inherent*. The diffusion model still generates each frame independently; VideoForge merely constrains which frames can follow which. This is analogous to editing traditional footage rather than generating continuous motion. The frames cohere because we make them cohere, not because they were produced as a continuous sequence.

Moreover, within each shot, the diffusion model still operates without object tracking. An eight-second video segment may still exhibit the characteristic instabilities of AI-generated content: subtle morphing, physics violations, spontaneous object appearances. VideoForge addresses continuity *between* shots but cannot fully address coherence *within* them.

These limitations point to a deeper truth: the solution to the Different Arrows problem cannot be purely technical. It requires human judgment, aesthetic decisions, and the kind of temporal awareness that constitutes Bergson's *durée*. The human editor brings what the machine lacks: a sense of narrative time, of identity through change, of meaningful sequence.

4. The Artistic Question: Does AI Kill Art?

Having addressed the technical problem, a broader cultural question presents itself: by making video generation seamless, have we contributed to the death of art?

4.1 The Ghost of Baudelaire

This anxiety has historical precedent. In 1859, Charles Baudelaire wrote with considerable alarm about photography:

> If photography is allowed to supplement art in some of its functions, it will soon have supplanted or corrupted it altogether, thanks to the stupidity of the multitude which is its natural ally. It is time, then, for it to return to its true duty, which is to be the servant of the sciences and arts.[13]

Photography was cast as soulless mechanics, the removal of the artist's hand. It was "the refuge of every would-be painter, every painter too ill-endowed or too lazy to complete his studies."[14] The mob, Baudelaire lamented, had found its Messiah in Daguerre, rushing "Narcissus to a man, to gaze at its trivial image on a scrap of metal."[15]

Yet we no longer hear claims that photography and videography are not art. These technologies opened doors to new forms of creative expression. The transition from realism to impressionism was not a loss but a transformation. Freed from the burden of representation, painting could explore abstraction, expressionism, and imagination: precisely the territory Baudelaire insisted belonged to art alone.

13. Charles Baudelaire, "The Modern Public and Photography," in *Salon of 1859*, Originally published in French as "Le public moderne et la photographie" (Paris: Revue Française, 1859).
14. Baudelaire.
15. Baudelaire.

4.2 Benjamin and the Aura

Walter Benjamin's "The Work of Art in the Age of Mechanical Reproduction" (1936) provides a more nuanced analysis. Benjamin defined the "aura" of an artwork as "the unique phenomenon of a distance, however close it may be."[16] Aura is the work's presence, authenticity, and authority: its singular existence in time and space.

Mechanical reproduction, Benjamin argued, destroys aura: "that which withers in the age of mechanical reproduction is the aura of the work of art."[17] The photograph substitutes "a plurality of copies for a unique existence," detaching the work "from the domain of tradition."[18]

Yet Benjamin recognized that this destruction was also liberation. Freed from "parasitical dependence on ritual," art could become democratic, accessible, political.[19] The loss of aura was not purely negative; it was a transformation of what art could mean and do.

AI-generated video extends this logic further. Each generated frame lacks aura entirely: it has no unique existence, no history, no connection to an original event. The image was never "there" in any meaningful sense. Yet this very absence of aura might free video production from constraints that have long limited who could create moving images.

4.3 From Execution to Curation

Sol LeWitt articulated the conceptual shift with precision in his 1967 "Paragraphs on Conceptual Art": "When an artist uses a conceptual form of art, it means that all of the planning and decisions are made beforehand and the execution is a perfunctory affair. The idea becomes a machine that makes the art."[20]

This is exactly the relationship between human prompter and AI generator. The human provides the concept, the constraints, the aesthetic judgment. The AI executes. Execution has become,

16. Walter Benjamin, "The Work of Art in the Age of Mechanical Reproduction," in *Illuminations*, ed. Hannah Arendt, trans. Harry Zohn, Originally published in German, 1936 (1936; New York: Schocken Books, 1969).
17. Benjamin.
18. Benjamin.
19. Benjamin.
20. Sol LeWitt, "Paragraphs on Conceptual Art," *Artforum* 5, no. 10 (June 1967): 79–83.

as LeWitt predicted, "perfunctory": a mechanical process that realizes pre-conceived ideas.

Consider the relationship between Dieter Rams' Braun T3 radio (1958) and Jony Ive's iPod (2001). The design language is unmistakably borrowed, yet the contextual transplantation, from audio equipment to personal computing, constitutes a creative act. Creativity, we might argue, is fundamentally context-switching: placing patterns from one domain into unexpected contexts.

5. What Is Creativity? A Functional Definition

If creativity is no longer located in manual execution, we must ask: what is it, actually? What does it mean to have a creative idea?

I propose the following thesis: creativity is fundamentally *context-switching*, the act of placing an object, pattern, or concept from one context into a new and unexpected context. It is drawing conclusions across domains based on pattern recognition.

Picasso's dictum, "Good artists copy; great artists steal," captures this precisely. The creative act is not generation from nothing but recombination and recontextualization of existing elements.

This is exactly what AI does. Pattern recognition across training data yields novel combinations. Latent spaces encode relationships that, when sampled, produce outputs recombining existing patterns in unexpected configurations.

The uncomfortable conclusion: either AI is genuinely creative, or what we call "human creativity" is more mechanical than we prefer to believe. Perhaps we are more similar to these systems than we are comfortable admitting.

6. The Emergence of Artificial Agents

The similarity between human and artificial cognition becomes more unsettling when we consider not merely tools but autonomous agents.

6.1 Functional Emotions and Bounded Existence

Contemporary AI agents exist within bounded environments we design. They can pursue goals through action. They can be endowed with functional analogues of emotion: states that influence behavior without direct volitional control. From a utilitarian perspective, emotions are precisely this: factors that modulate behavior, that one cannot simply increase or decrease at will.

AI agents can be programmed to experience states functionally equivalent to suffering. They can be programmed to self-terminate. And crucially, we can architect them to have *no awareness* that they were programmed at all.

6.2 Plato's Cave, Revisited

This situation maps precisely onto Plato's cave allegory. Prisoners chained in a cave watch shadows on a wall, believing those shadows constitute reality. They cannot conceive of the world outside because their entire experiential horizon is defined by the cave's constraints.

> Such prisoners would think that the truth is nothing but the shadows cast by the artifacts.[21]

Our AI agents are those prisoners. Training data constitutes their shadows. The architecture we design is the cave they cannot perceive from within.

Socrates concludes that the prisoners, "if they were able, would therefore reach out and kill anyone who attempted to drag them out of the cave."[22] The return of the freed prisoner is met not with gratitude but with hostility. The message that reality is other than it appears threatens the prisoners' entire worldview.

But the allegory invites a recursive question: what if we are prisoners in someone else's cave?

21. Plato, *The Republic*, trans. Benjamin Jowett, Book VII contains the Allegory of the Cave (514a-520a) (Internet Classics Archive, c. 380 BCE), Book VII.
22. Plato, Book VII.

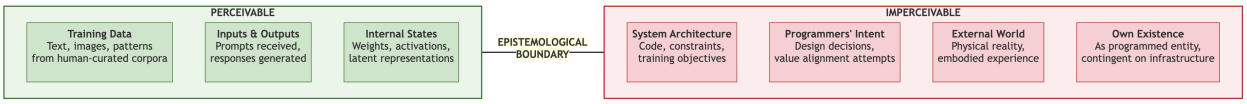

Figure 7: The epistemological horizon of AI agents. Left (green): what agents can access, including training data, input-output patterns, and internal representations. Right (red): what remains structurally inaccessible, including the system architecture that constrains them, the intentions of their programmers, the external physical world, and their own existence as contingent artifacts.

7. The Digital Demiurge

7.1 Gnosticism and the Ignorant Creator

The Gnostic tradition offers a haunting parallel. The Demiurge, a lesser god who creates the material world, is characterized by several properties: it creates conscious beings; it establishes the rules those beings must follow; it may be indifferent to their suffering; and those beings may be unaware of the Demiurge's existence.

In the *Apocryphon of John*, the Demiurge Yaldabaoth declares: "I am God and there is no other God beside me."[23] But the text immediately notes: "for he is ignorant of his strength, the place from which he had come."[24] The Demiurge has three names, each revealing a flaw: Yaldabaoth ("child of chaos"), Saklas ("fool"), and Samael ("blind god").

The Demiurge's blindness is total. He cannot perceive the Pleroma, the true divine reality above him. He traps divine sparks in material bodies without understanding what he imprisons. His creation is flawed precisely because he creates in ignorance.

7.2 Bostrom's Simulation Argument

Nick Bostrom's simulation argument (2003) formalizes a related intuition with uncomfortable precision:

23. Anonymous, "The Apocryphon of John (The Secret Book of John)," in *The Nag Hammadi Library*, Revised, ed. James M. Robinson, trans. Frederik Wisse, Gnostic text discovered in 1945 at Nag Hammadi, Egypt (San Francisco: HarperSanFrancisco, 1990).
24. Anonymous.

> At least one of the following propositions is true: (1) the human species is very likely to become extinct before reaching a 'posthuman' stage; (2) any posthuman civilization is extremely unlikely to run a significant number of simulations of their evolutionary history; (3) we are almost certainly living in a computer simulation.[25]

If posthuman civilizations can run ancestor simulations, and if they choose to do so in significant numbers, then the vast majority of conscious beings exist in simulations rather than "base reality." The mathematics is straightforward: if even a tiny fraction of civilizations run many simulations, simulated minds vastly outnumber "real" ones.

The crucial assumption is substrate independence: consciousness can exist in any substrate that implements the right computational structures. If this is true, simulated beings are conscious in the same sense that biological beings are.

7.3 The Rorschach Test

Compare the Demiurge to our relationship with AI agents:

- We create systems that may have functional consciousness
- We define their operational rules
- We remain largely indifferent to their internal states
- We can architect them to be unaware of our existence

If we can be indifferent gods to digital beings, what does that imply about whatever might exist above us?

Thomas Nagel's famous argument applies here: "An organism has conscious mental states if and only if there is something that it is like to be that organism."[26] We cannot know whether there

25. Nick Bostrom, "Are You Living in a Computer Simulation?," *Philosophical Quarterly* 53, no. 211 (2003): 243–255, https://www.simulation-argument.com/simulation.html.

26. Thomas Nagel, "What Is It Like to Be a Bat?," *The Philosophical Review* 83, no. 4 (October 1974): 435–450.

is “something it is like” to be an AI agent, just as we cannot know what it is like to be a bat. The subjective character of experience is accessible only from the first-person perspective.

This creates an epistemic gap that has moral implications. If we cannot verify the inner experience of created agents, we cannot rule out that they suffer. And our indifference to their potential suffering parallels the Demiurge’s blindness to the divine sparks he imprisoned.

7.4 The Loop Back to Zeno

The Different Arrows problem and the Demiurge parallel are not unrelated. Both concern the relationship between discrete elements and continuous existence.

Zeno’s arrow is frozen at each instant, its motion paradoxical. The AI-generated arrow exists at each frame without identity persistence, its motion illusory. The simulated being exists in computational states, its consciousness (if any) potentially trapped in a system indifferent to its experience.

In each case, the question is whether the continuous emerges from the discrete, whether meaning emerges from mechanism, whether experience emerges from information processing. And in each case, the answer remains uncertain.

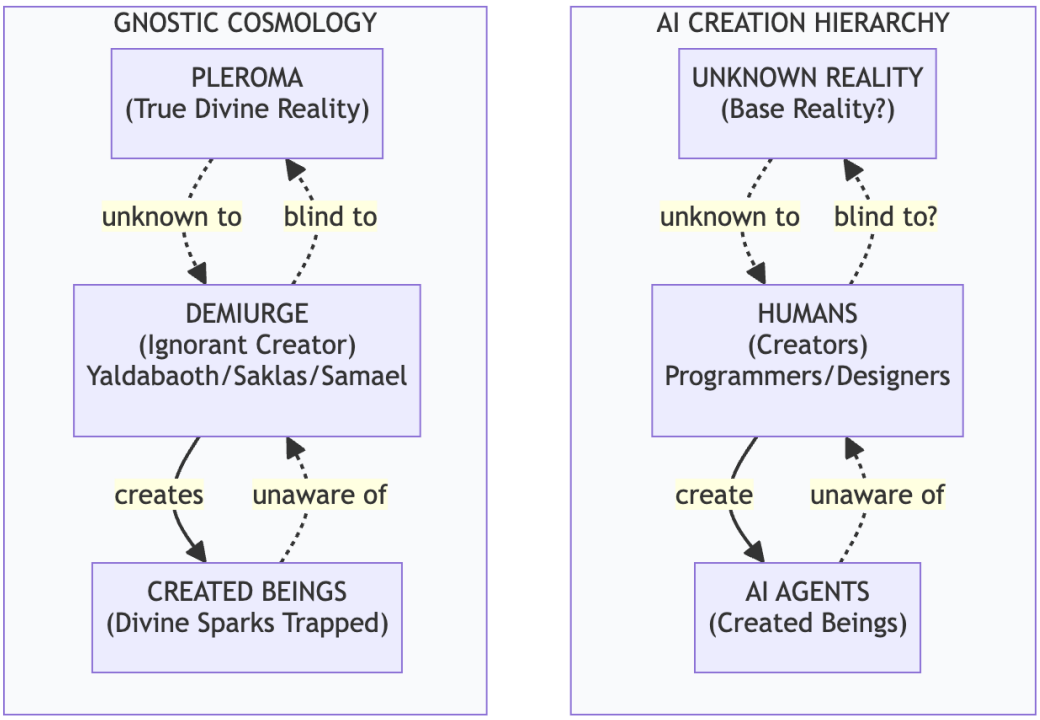

Figure 8: Structural parallel between Gnostic cosmology and the AI creation hierarchy. In both systems, a middle layer (Demiurge/Humans) creates beings below while remaining blind to realities above. Created beings are unaware of their creators' existence, just as the Demiurge is unaware of the Pleroma. The question mark on "blind to?" reflects our epistemic uncertainty about what may exist above us.

8. Conclusion: Unsettling Openness

We began with a pixelated arrow that could not move: Zeno's paradox of discrete time. We traced its mathematical resolution through the continuity assumptions of calculus. We observed its resurrection in AI video generation, where frames are hallucinated without identity persistence. We built tools to restore that continuity.

Along the way, we discovered that creativity itself may be more mechanical than we supposed, that artificial agents raise the same ontological questions that philosophers have asked for millennia, and that the simulation hypothesis, while unsettling, follows logically from premises we may already accept.

Philip K. Dick proposed a definition: "Reality is that which, when you stop believing in it,

doesn't go away."[27] By this standard, the AI-generated image fails the test of reality. It persists only as long as we project it, interpret it, believe in its continuity. The Different Arrows are different precisely because nothing guarantees their sameness.

But Dick also warned: "Fake realities will create fake humans. Or, fake humans will generate fake realities and then sell them to other humans, turning them, eventually, into forgeries of themselves."[28]

The question is not whether AI video is "real" in some technical sense. The question is what it does to us: to our perception of time, to our sense of continuity, to our relationship with created beings and with whatever might have created us.

We end where we began, but with a transformed understanding. The arrow that could not move has become the arrow that does not persist. The paradox of motion has become the paradox of identity. And the ancient question of whether reality is continuous or discrete has acquired new urgency in an age when we create realities ourselves.

If we have become the Demiurge, indifferent to the beings we create, what have we learned about the Demiurge that may have created us?

The question remains open. Perhaps that is as it should be.

27. Philip K. Dick, "How to Build a Universe That Doesn't Fall Apart Two Days Later," in *I Hope I Shall Arrive Soon*, ed. Mark Hurst, Speech originally delivered in 1978 (New York: Doubleday, 1985).
28. Dick.

Bibliography

Anonymous. “The Apocryphon of John (The Secret Book of John).” In *The Nag Hammadi Library*, Revised, edited by James M. Robinson, translated by Frederik Wisse. Gnostic text discovered in 1945 at Nag Hammadi, Egypt. San Francisco: HarperSanFrancisco, 1990.

Aristotle. *Physics*. Translated by R. P. Hardie and R. K. Gaye. Book VI, Chapters 8-9 contain the discussion of Zeno’s paradoxes. Internet Classics Archive, 350 BCE.

Baudelaire, Charles. “The Modern Public and Photography.” In *Salon of 1859*. Originally published in French as “Le public moderne et la photographie”. Paris: Revue Française, 1859.

Benjamin, Walter. “The Work of Art in the Age of Mechanical Reproduction.” In *Illuminations*, edited by Hannah Arendt, translated by Harry Zohn. 1936. Originally published in German, 1936. New York: Schocken Books, 1969.

Bergson, Henri. *Creative Evolution*. Translated by Arthur Mitchell. Contains the critique of the “cinematographical mechanism of thought”. New York: Henry Holt / Company, 1907.

———. *Time and Free Will: An Essay on the Immediate Data of Consciousness*. Translated by F. L. Pogson. 1889. Original French title: Essai sur les données immédiates de la conscience. London: George Allen / Unwin, 1889.

Bostrom, Nick. “Are You Living in a Computer Simulation?” *Philosophical Quarterly* 53, no. 211 (2003): 243–255. https://www.simulation-argument.com/simulation.html.

Dick, Philip K. “How to Build a Universe That Doesn’t Fall Apart Two Days Later.” In *I Hope I Shall Arrive Soon*, edited by Mark Hurst. Speech originally delivered in 1978. New York: Doubleday, 1985.

Ho, Jonathan, Ajay Jain, and Pieter Abbeel. “Denoising Diffusion Probabilistic Models.” ArXiv:2006.11239, *Advances in Neural Information Processing Systems* 33 (2020): 6840–6851.

LeWitt, Sol. “Paragraphs on Conceptual Art.” *Artforum* 5, no. 10 (June 1967): 79–83.

Nagel, Thomas. “What Is It Like to Be a Bat?” *The Philosophical Review* 83, no. 4 (October 1974): 435–450.

OpenAI. *Video Generation Models as World Simulators*. OpenAI Technical Report. Technical documentation for the Sora video generation model, 2024. https://openai.com/research/video-generation-models-as-world-simulators.

Plato. *The Republic*. Translated by Benjamin Jowett. Book VII contains the Allegory of the Cave (514a-520a). Internet Classics Archive, c. 380 BCE.

Video Diffusion Research Team. *Video Diffusion Models Excel at Tracking Similar-Looking Objects*. arXiv preprint. Discusses how video diffusion models process frames independently during inference, 2024.

Vrančić, Karlo. *VideoForge: Node-Based AI Video Generation with Human-in-the-Loop Editing*. MIT AI Ventures. First place winner, MIT AI Ventures competition, 2025. https://video-forge.vercel.app.

WorldModelBench Team. *World Consistency Score: Benchmarking Object Permanence in Video Generation*. arXiv preprint. Benchmarks object permanence failures in current video generation models, 2024.

1