

WHEN SEEING IS DECEIVING.

ART&MEDICINE

Why generative AI can't do the one thing medical visualisation demands.

AI is seen by many as the answer to the world's problems in all areas of life, and business. And in the medical world AI is seen as a labour-saving way to visualise treatments and anatomy.

This is a word of caution. AI visualisations cannot guarantee the anatomical accuracy that a clinical audit demands. And they are not replicable years after delivery so have a limited life span. The premise of this article is that visualisations created by software experts in collaboration with clinicians are infinitely superior, more accurate and more adaptable than those created with generative AI.

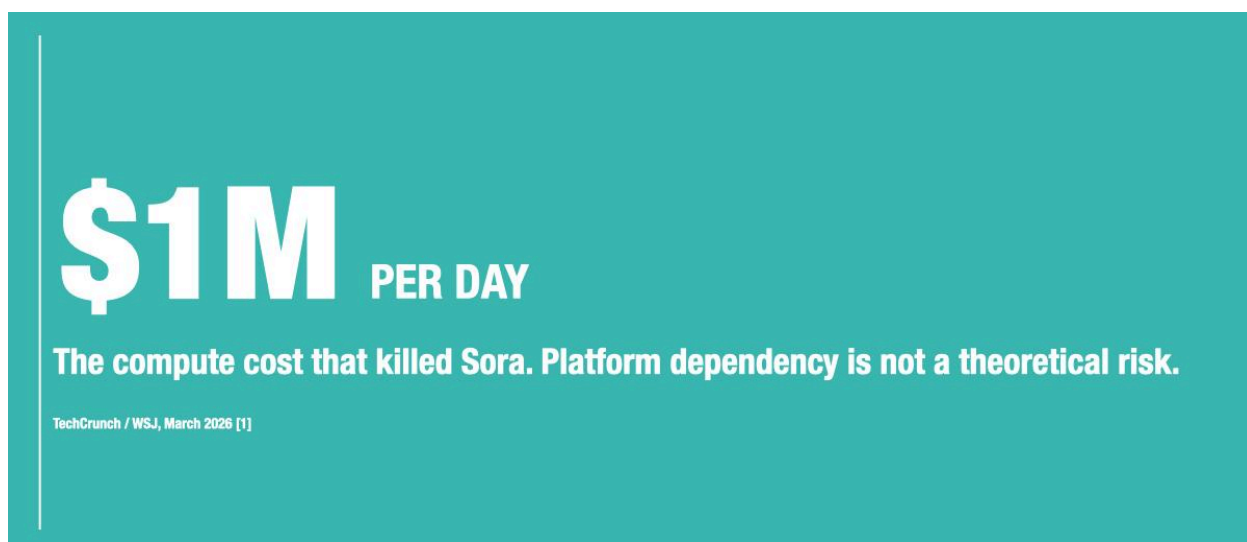
On April 26, 2026, OpenAI shut down Sora.

The generative AI video platform - once described as a glimpse of the creative future - had been live for barely six months. At peak, it was burning through roughly \$1 million per day in compute costs. Its user base had collapsed from around a million to fewer than half a million. A \$1 billion partnership with Disney, signed just three months earlier, was dead. Disney learned the platform was being killed less than an hour before the public announcement. [1]

This was not a niche experiment. Sora was the flagship generative video product from the world's most prominent AI company, backed by billions in investment and enormous institutional partnerships.

It vanished overnight.

If you are using generative AI as a production tool in pharmaceutical communication - where assets may need to be reproduced, revised, and resubmitted years after original delivery - this should give you pause. Serious pause.



The image is the argument

In medical visualisation and pharmaceutical narrative animation, a visual is not an illustration. It makes a scientific claim.

A visualisation of a healthy heart built from population MRI and CT data is a statement about how that organ looks and behaves. An animation of a monoclonal antibody binding to its target receptor explains how a drug works - to clinicians, regulators, and patients alike. In both cases, the accuracy of the image and its consistency across every use are the conditions under which the work has any value at all. Get the science wrong in the image and you have communicated the wrong science. Persuasively.

When Seeing is Deceiving

This distinction matters now more than it did a year ago. Not because generative AI has stopped improving. It hasn't. But because the regulatory environment in which pharmaceutical communication operates has become dramatically less tolerant of visual claims that cannot be verified.



Construction from evidence vs. synthesis from statistics

The fundamental advantage of a 3D, human-operated, artistic pipeline is that imaging data becomes the literal foundation of a model. When building a generic healthy heart from population CT and MRI datasets, the volumetric data is not inspiration - it is material. The artist constructs the mesh in direct dialogue with that data: tracing anatomical boundaries, calibrating proportions against measured structures, encoding the evidence in geometry that can be examined, questioned, and audited.

What emerges is not a plausible-looking heart. It is a representation of how hearts actually look, built from proof.

Once built, that mesh is fixed and shared. The same heart appears in the surgical simulation, patient education, clinician training modules, and the conference presentation; they are referenced, not regenerated. If a clinical review identifies that the mitral valve geometry needs correcting, the fix propagates to every production that uses it. The link between the science and the visualisation is permanent.

A diffusion generative AI model cannot do this. It cannot ingest a CT dataset and produce geometry that encodes the morphological information within it. What it produces are images shaped by statistical patterns in training data - a plausible approximation of what hearts tend to look like in pictures, not a construction grounded in evidence. There is no geometry to correct. No traceable link between the imaging data and the output. The image looks like an argument. It isn't one.

The consistency problem

It would be dishonest to pretend that generative AI has stood still. Runway's Gen-4, released in early 2025, introduced reference-image systems that maintain consistent character appearance across multiple shots with different camera angles and lighting. [4] Google's Veo 3.1 has made visible progress on what the industry calls the 'world consistency' problem. [5] These are real improvements, and they matter for entertainment, marketing, and general creative production.

They do not matter here.

Drug mechanism animation demands persistent object identity across a sequence. A macrophage introduced at the opening of an animation - moving through tissue, encountering a drug molecule, internalising it, changing activation state - must be the same object throughout. Same membrane geometry. Same receptor distribution. Same organelle structure. The narrative logic depends on the viewer following one biological actor through a process. Looking similar is not good enough. It has to be the same thing.

In a 3D pipeline, this identity is enforced by the pipeline itself. The macrophage is a rigged, shaded asset. Its geometry cannot drift because there is no mechanism by which it could. In a generative workflow, even the best current models apply probabilistic pressure toward resemblance. They make the output more likely to look like the reference. They do not make it the same object.

OpenAI's own Sora was described, even by sympathetic reviewers, as unable to reliably produce consistent characters across clips. [6] That was the most expensive generative video model ever built. If it couldn't hold a character's face together across a ten-second sequence, what happens to a receptor binding site across a three-minute mechanism narrative?

Generative AI can make an output resemble a reference. It cannot make it the same object. For this work, that difference is everything.

The Benefits Of 3D Pipeline vs. Generative AI re-rolling

Consistency across a sequence is one problem. There is another that gets less attention but matters just as much: what happens when you try to make something better. In a 3D pipeline, iteration is the engine of quality. Every lighting decision is grounded in the underlying geometry. Every animation is reviewed against the biological behaviour it depicts. The feedback loop between artist and asset is continuous and cumulative - you can see it working at every stage. A rough mesh becomes, through successive deliberate refinement, a model that can bear scientific scrutiny. The pipeline doesn't resist this process. It was designed for it.

Generative AI inverts that relationship entirely. Each attempt to improve an output is simultaneously an attempt to hold everything else still - which the model cannot do.

Correcting an anatomical inaccuracy by adjusting a prompt will shift the lighting, the composition, the surface detail, and the surrounding environment in ways that cannot be predicted or isolated. The artist is not refining. They are re-rolling, hoping the next result is closer to what they need without having introduced new problems. In a production where biological accuracy is the measure of every frame, an uncontrolled re-roll is not an improvement. It is a fresh source of error.

The search for quality in a generative workflow does not converge. It wanders. And in wandering, it introduces inaccuracies that were not present in the version it was trying to correct.

Precise iteration is also what makes clinical collaboration possible. In a 3D pipeline, the artist presents the outcome of a specific stage of work to a clinical expert. The clinician identifies a structure with incorrect geometry, or a motion that doesn't reflect how a cellular process actually unfolds. The artist acts on that note with surgical precision - adjusting the relevant asset without disturbing anything around it. The clinician sees the improvement. The conversation continues.

What develops over time is more than a production workflow. The artist develops a deeper understanding of the science. The clinician develops a sharper eye for what the artist and their tools are capable of. That shared understanding doesn't reset between productions. It compounds - in the people, and in the assets they build together.

Generative AI cannot host this conversation. When the clinician gives feedback, there is no layer to adjust, no asset to refine, no targeted response available. There is only a new prompt and a new gamble. The knowledge the clinician brought to the review dissolves into the uncertainty of the next generation.

But what about text-to-3D?

There is a counterargument worth taking seriously: the choice isn't between 2D generative output and a 3D pipeline anymore. AI can now generate 3D geometry directly.

And that's true. Tencent's Hunyuan3D-2, released in January 2025, generates high-resolution textured meshes in seconds. [7] Adobe's Substance 3D now generates editable 3D models from text prompts. [8] The technology is advancing fast.

In its current form, though, it doesn't solve the problem. A mesh generated by diffusion from training data statistics is still geometry synthesised from what hearts - or cells, or molecules - tend to look like in photographs. It is not geometry constructed from patient imaging evidence. It cannot be traced to a CT dataset. It cannot be audited against population morphological data. It cannot be clinically reviewed in the way that a mesh built in direct dialogue with volumetric imaging data can.

But the trajectory is worth watching.

When Seeing is Deceiving

Today, medical scanners output imaging data in a universal format called DICOM. Established tools already convert that data into 3D meshes through segmentation - software like 3D Slicer, Materialise Mimics, and ITK-SNAP has done this for years. It works. So the question is not whether AI could generate 3D anatomy. It's whether AI could do it better, or at least faster and cheaper, than the tools that already exist.

What matters is what the models are trained on. If text-to-3D systems were built not on photographic image datasets but on curated libraries of segmented, clinically validated CT and MRI volumes, they could become a useful accelerant within a 3D pipeline. A model trained on thousands of segmented cardiac volumes might generate a more complete base mesh from partial data, and do it faster, than rule-based segmentation alone. The starting point would be informed by real anatomical evidence rather than averaged from surface appearances.

It would still only be a starting point.

The mesh would still need to be validated by a clinician. Refined by an artist. Integrated into a pipeline where every subsequent decision is traceable and auditable. The AI would compress the early stages of construction. It would not replace the human judgement that makes the result trustworthy.

We'd welcome that. Faster geometry built on the right foundations, feeding into a pipeline that guarantees accuracy, is not a threat to this way of working. It's an enhancement. The danger is when faster geometry is mistaken for the finished product - when speed is treated as a substitute for rigour rather than a complement to it.

Your assets are not your assets

Pharmaceutical communication sits within a regulated environment. Animations used in medical education, clinical training, or regulatory submissions may need to be reproduced exactly - same geometry, same motion, same lighting - months or years after original delivery.

A 3D pipeline guarantees this. The source files contain everything needed to reconstruct the output. They do not change unless deliberately modified. Re-rendering a scene from three years ago produces the same frame.

A generative AI workflow offers no equivalent. These tools are hosted services, run by the companies that built them. When the provider updates its model - which it can do at any time, without notice - the tool changes under you. The same instruction that produced a reliable result at delivery may produce a visibly different result six months later. There is no way to examine what changed, and no way to roll it back.

And then there is the more fundamental risk: the platform disappears entirely. Three days ago, every production asset created on Sora became unreproducible. Not degraded. Not

inconsistent. Gone. A billion-dollar Disney partnership collapsed because the platform it depended on ceased to exist. [1]

Some providers are beginning to address this. Google's MedGemma project, for diagnostic AI, distributes fixed versions of its software that cannot change after download - so researchers can guarantee the same results years later. [9] But the video and visualisation tools - Runway, Veo, Kling - don't work this way. They remain hosted services where the software behind the interface can change at any time. For regulated work, that is not a trade-off worth making.

02 APRIL 2026

The FDA issued a warning letter specifically citing the risks of over-reliance on AI agents in the creation of specifications, procedures, and records intended for regulatory compliance.

Clarkston Consulting, April 2026 [10]

The regulators are paying attention

The FDA's 2025 enforcement surge was unprecedented. Over 200 letters to pharmaceutical companies, with a particular focus on visual representations that overstated efficacy or created misleading impressions of drug benefit. [2] The agency specifically objected to visual content that implied outcomes not demonstrated in clinical trials - 'attention-grabbing' visuals, rapid scene changes, lifestyle imagery suggesting quality-of-life improvements the data did not support. [3]

In January 2026, the FDA and EMA jointly published guiding principles for AI in drug development, emphasising credibility assessment, validation, and documentation requirements for AI-generated outputs used in regulatory contexts. [11]

None of this prohibits generative AI in medical communication. But the burden of proof for any visual claim has never been higher. A production pipeline's ability to demonstrate exactly how a visual was constructed, from what data, with what clinical review, has become a regulatory necessity rather than a competitive advantage.

A 3D artistic pipeline produces an auditable chain from evidence to output. Generative AI produces an output. In a regulatory environment that is actively punishing unverifiable visual claims, that gap should keep people up at night.

Assets vs. outputs

A rigorously built healthy heart model - constructed from population imaging data, validated through clinical review - is not a single-use production asset. It is a foundational resource. Re-lit for an endocarditis animation. Re-sectioned for a surgical training module. Re-animated for patient education. Updated as new imaging data becomes available. Across years. Across multiple client relationships. Each reuse costs a fraction of the original build and carries the same clinical authority.

Generative AI produces outputs. A 3D pipeline produces assets. One is a recurring expenditure. The other is a library that grows with every production.

The case for construction

This is not a debate about old versus new. It is about what the work demands.

Anatomical fidelity traceable to evidence. Consistency enforced by design. The ability to revise with precision when the science evolves, and to reproduce the work on demand, years after delivery. These are not aspirational qualities. They are the minimum conditions under which pharmaceutical visualisation has any value.

On 26 April 2026, the world's most prominent generative video platform shut down overnight. Its billion-dollar partnerships collapsed. Its assets became unreproducible. The regulatory environment for pharmaceutical visual communication has never been stricter. And the technology that was supposed to make 3D pipelines obsolete still cannot hold the identity of a single object across a three-minute sequence.

In pharmaceutical communication, the image is the evidence. If it cannot be traced to the science, it is not communication. It's wallpaper.

ART&MEDICINE

Seeing is believing

REFERENCES

Sources cited in this article:

1. TechCrunch (March 2026) - Why OpenAI really shut down Sora. Reporting based on WSJ investigation. Sora burning ~\$1M/day in compute; user base collapsed to <500K; Disney partnership terminated with less than an hour's notice. techcrunch.com/2026/03/29/why-openai-really-shut-down-sora/
2. King & Spalding (January 2026) - 2025 Year in Review: FDA Drug and Device Advertising and Promotion Enforcement. 200+ enforcement letters; 74 to pharma/biologic manufacturers; focus on misleading visual representations of efficacy. kslaw.com
3. Gardner Law (October 2025) - Healthcare and FDA Compliance in 2025: FDA War Stories from the Front Lines. 46 untitled letters on September 9 alone; 'attention-grabbing' visual content flagged for creating misleading efficacy impressions. gardner.law
4. Runway AI (March 2025) - Introducing Runway Gen-4. Reference image system for character consistency across shots. runwayml.com/research/introducing-runway-gen-4
5. PXZ.ai (January 2026) - Veo 3.1 vs Top AI Video Generators: 2026 Comparison. Notes Gen-4 and Veo 3.1 progress on 'world consistency'. pxz.ai/blog
6. MindStudio (March 2026) - Why OpenAI Killed Sora and What It Means for AI Video Generation. Cites creative control limitations and inability to reliably produce consistent characters across clips. mindstudio.ai
7. Tencent Hunyuan3D Team (January 2025) - Hunyuan3D 2.0: Scaling Diffusion Models for High Resolution Textured 3D Assets Generation. [arXiv:2501.12202](https://arxiv.org/abs/2501.12202). arxiv.org/abs/2501.12202
8. Adobe Research (January 2025) - The breakthrough research behind Adobe Substance 3D Viewer's new Text to 3D feature. Text-to-3D generation using transformers and Gaussian Splatting. research.adobe.com/news/breakthrough-research-behind-text-to-3d/
9. Google Research (2025) - MedGemma: Open models for health AI development. Distributed as frozen model snapshots; parameters will not change unexpectedly over time. research.google/blog/medgemma
10. Clarkston Consulting (April 2026) - FDA Issues Warning for Inappropriate AI Use. April 2, 2026 warning letter citing AI agents used in creation of specifications and records intended for regulatory compliance. clarkstonconsulting.com
11. FDA / EMA (January 2026) - Guiding Principles of Good AI Practice in Drug Development. Joint principles for AI credibility assessment and validation in regulatory contexts. fda.gov
12. Fahad et al. (2025) - Generative AI in clinical (2020–2025): a mini-review. *Frontiers in Digital Health*. Confirms hallucination risks and bias amplification in clinical generative AI. doi: 10.3389/fdgth.2025.1653369