KAMX CHROMA STUDIOS
PRESENTS

SYNTHETIC CINEMA

JOURNAL OF AI-ASSISTED FILMMAKING

ISSUE 001
OCT 2026

FILM
TECHNOLOGY
CREATIVITY
REAL WORLDS
NEW POSSIBILITIES

"New tools. New worlds. Same human obsession."

CINEMA
BEYOND
LIMITS

COVER STORY

CAN AI MAKE A BLOCKBUSTER?

THE RISE OF THE DIRECTOR AS PRODUCTION SYSTEM

KAMX
CHROMA STUDIOS

REAL IDEAS / SYNTHETIC WORLDS / HUMAN STORIES

A MONTHLY PUBLICATION BY
KAMX CHROMA STUDIOS
EDITOR ENRICO R. F. VALENTINO

Synthetic Cinema

A monthly journal of AI-assisted filmmaking

Publisher
KAMX Chroma Studios

Editor
Enrico Rudolph Famulari Valentino

Editorial and creative direction
KAMX Chroma Studios

Issue
001, October 2026. Published 1 October 2026.

Editorial cutoff
20 September 2026. Items dated later were checked on 21 September 2026 and are flagged where they matter.

Frequency
Monthly

Cinema is not becoming automatic. It is becoming programmable.

Synthetic Cinema covers the changing craft of film: cinematography, directing, performance, production design, visual effects, sound, rights and the new generation of AI production tools. Our rule is deliberately simple. Cinema first, technology second. A tool earns space in these pages only when it changes what a filmmaker can actually do on a real production, with a real schedule and a real budget.

We write for directors, producers, VFX supervisors, editors and the people who will soon hold job titles that do not exist yet. We assume you know what a lens does. We do not assume you have time to read every model launch.

Editorial charter

We separate verified fact from analysis on every page. Product capabilities and legal developments are checked against first-party sources wherever they exist, and dated. Model comparisons are production observations with a shelf life, not permanent rankings. A benchmark winner is not automatically the best filmmaking tool. Human authorship, direction, consent and creative intent remain central to this publication.

How to read the evidence marks

● **Verified.** Confirmed against a first-party or primary source listed in the source notes.

◐ **Vendor claim.** Stated by the company that sells the product. Reported accurately, not independently tested.

○ **KAMX analysis.** Our editorial judgement or a hypothesis we intend to test. Treat it as argument, not fact.

[07] **Source.** Numbered references resolve to the source notes on page 33.

> The quality of the machine matters. The quality of the person directing it matters more.

Inside issue 001

A field journal for filmmakers working in the synthetic production era.

On the cover: London, 2098. The cover story, “Can AI make a blockbuster?”, begins on page 05.

About the section tabs

The coloured tabs on each page edge follow the film industry’s script-revision colour sequence: white, blue, pink, yellow, green, goldenrod, buff, salmon, cherry, then double-blue onwards. Flip the edge to find a department.

> Direction remains the act of knowing what belongs in the frame, why it belongs there, what should happen next, and what must be rejected.

Editor's letter

The person behind the wheel

AI is changing the production apparatus. It is not removing the need for direction. It is making direction more consequential.

For a long time the debate around AI filmmaking has been staged as a contest: machine against camera, synthetic actor against human actor, prompt against set. I think that framing misses the most interesting possibility.

I do not imagine the future director confined to a desk. I imagine a director still walking London at dawn, studying a location, watching the light change, photographing textures, testing camera positions, casting performers, building objects, recording movement and sound, and then using AI to extend those choices far beyond what a conventional budget would allow.

The production question is not simply "Can AI make a film?" It is this: can a director manufacture a feature film using AI as part of the production apparatus? My answer is increasingly yes, with one condition.

The condition is control. A powerful model in the wrong hands does not create good cinema. It creates more output. Direction remains the act of knowing what belongs in the frame, why it belongs there, what should happen next, and what must be rejected.

This first issue is built around that idea. The cover story treats a feature as a system of shots rather than one long generation. The camera war asks which engine should carry which shot, instead of which engine is "best". The test bench sets out how we will measure usable footage rather than beautiful demos. And the rights and platform pages deal with the parts of production that no model can solve for you: consent, contracts and whether your tools will still exist when you reach the final mix.

We have marked what is verified, what is a vendor's claim and what is our opinion. When our tests contradict a vendor, or contradict us, the test wins. That is the territory this magazine will explore.

Enrico

Enrico Rudolph Famulari Valentino, Editor

Cover story

Can AI make a blockbuster?

The real question is no longer whether a model can generate a beautiful shot. It is whether a director can manufacture an entire film from thousands of controlled decisions.

Every few months a new video model is greeted as the moment AI "can make movies". The demos get longer, the faces get steadier and the physics get more convincing. None of this answers the question a producer actually asks, which is whether a team can deliver ninety to one hundred and twenty minutes of coherent, rights-clean, finishable cinema on a schedule.

This cover story argues that the answer depends less on any single model than on the production system wrapped around it. It proceeds in four parts: what a feature actually is, what a director must now own that used to be owned by the camera, why continuity is a data problem, and where the machine still bends.

By Synthetic Cinema editorial, KAMX Chroma Studios

A feature is not one generation

A 100-minute movie is a system of shots. The production challenge is coordination, not asking a model to remember two hours of story.

When people ask whether AI can make a Hollywood-scale feature, they usually imagine a single model being asked to “generate the movie”. That is almost certainly the wrong mental model, and it produces the wrong investments.

Conventional cinema is already modular. A feature is assembled from performances, plates, set extensions, inserts, transitions, sound effects, dialogue, music and hundreds or thousands of separately produced shots. Nobody on a live-action set shoots the movie in one take; the continuity script supervisor, the art department and the editor exist precisely because the film is manufactured in pieces and reassembled. Editorial hides the machinery.

Run the arithmetic. If a 100-minute film averaged five seconds per shot, it would contain roughly 1,200 shots. Cut faster, as much contemporary action cinema does, and the count climbs towards two thousand. The number itself is not the point. The point is that Shot 646 does not need a mysterious memory of the previous 645 shots if the production system can supply the correct actor, costume, prop, location, damage state, lens language and lighting state.

That reframes the problem. AI does not have to be the film’s memory. The production does.

It also reframes the economics. A model’s maximum clip length matters, but far less than its **usable take rate**: how many of its generations survive editorial review. A thirty-second generation that is usable once in ten attempts may cost more screen time per approved second than a ten-second engine that lands every third try. We return to this in the test bench on page 25.

Seen this way, the blockbuster question becomes a question about production management. Can the team define the film precisely enough, in assets and states, that any competent engine renders the next shot correctly? If so, the ceiling on scale is set by budget and discipline, not by a context window.

How many shots is a feature?

Shots in a 100-minute film at a given average shot length (ASL). Illustrative arithmetic, not a survey.

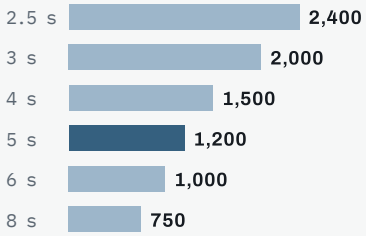

Highlighted: the 5-second ASL used throughout this story. ○ KAMX calculation.

> The blockbuster problem is not “generate 100 minutes”. It is “control the next five seconds, 1,200 times”.

Scale

The same frozen identity can feed hundreds of shots. Scale comes from repeatable asset control, not one impossibly long generation.

Failure mode

Continuity collapses when every shot silently invents a new version of the world. Canonical masters stop that drift becoming invisible.

Directorial advantage

Attention goes to story, blocking, framing and performance because production state is explicit, not remembered informally.

The director as production system

The director remains the coordinating intelligence. AI expands production capacity; it does not remove authorship.

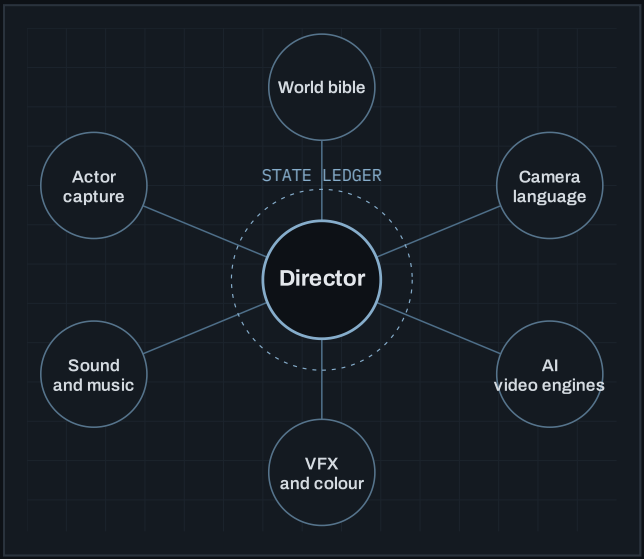

Every department reports to one intent. The dashed ring is the state ledger: the shared record of what is true in the film at any moment.

On a conventional production the camera enforces a great deal of consistency for free. The actor has one face. The set exists in one place. The prop is a physical object that stays cracked once it has been cracked. A generative pipeline removes those free guarantees. Every shot is, in principle, a fresh invention.

That shifts work upstream to the director and the people around them. Decisions that used to be implicit in the physical world must now be made explicit and written down: which version of the character is canonical, which lens family the film uses, what the light is doing in scene 71, which engine is allowed to render which kind of shot.

The result is a role that looks less like “prompt writer” and more like a combination of director, production designer and pipeline supervisor. The creative choices are the same ones cinema has always required. What changes is that the choices must be recorded precisely enough for a machine to follow them.

What the director now owns

The canon: approved identity masters, location anchors, prop states and the look of the film. Nothing becomes true in the movie until it is approved into the canon.

What the pipeline owns

Routing, versioning and records: which engine produced which plate, from which references, at what cost, and whether the result was approved.

What the engines own

Rendering. The model is a camera and a performer-for-hire, not the continuity department and not the author.

Who decides what: a working split

Decision	Director	Pipeline supervisor	Engine	Recorded in
Approve a character identity master	Decides	Versions and stores	—	`CHARACTERS/<name>/STATE_nn`
Change a story state (wet, cracked, destroyed)	Decides	Updates the ledger	—	`STATE_LEDGER`
Choose the engine for a shot class	Approves	Proposes from test data	—	`ROUTING_TABLE`
Render a take	—	Submits with references	Executes	`SHOTS/<id>/takes`
Accept a take into the edit	Decides	Locks it as evidence	—	`SHOTS/<id>/APPROVED`

○ KAMX working model. Job titles vary between productions; the separation of decision, record and execution is the point.

Continuity is a systems problem

Generative drift becomes dangerous the moment the previous generation is allowed to become the new source of truth.

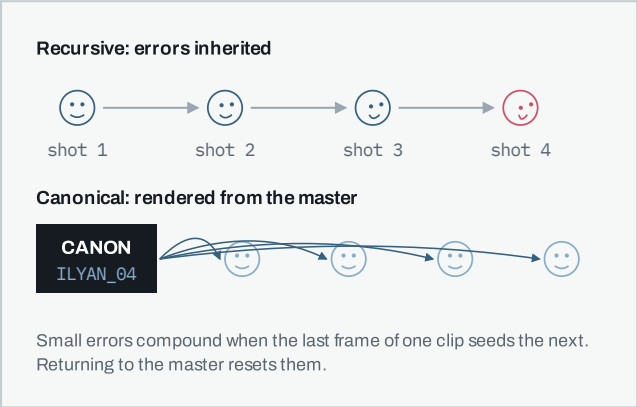

Small errors compound when the last frame of one clip seeds the next. Returning to the master resets them.

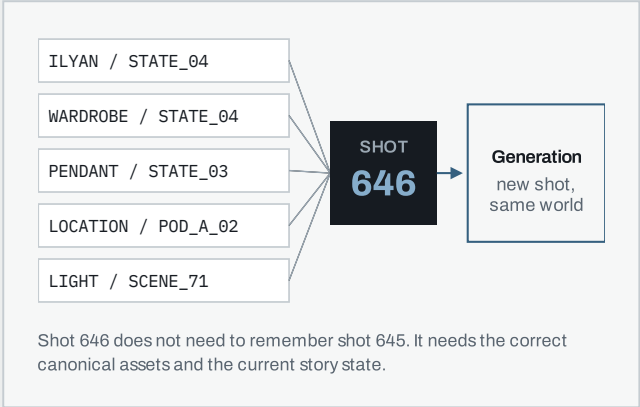

Shot 646 does not need to remember shot 645. It needs the correct canonical assets and the current story state.

Recursive generation is seductive. Take the last frame of one clip, use it as the first frame of the next, repeat. It feels like continuity because each join is smooth. It is also a near-perfect mechanism for slow error accumulation. A face changes slightly; an accessory moves; a vehicle loses a panel; an architectural detail migrates two metres to the left. No single join looks wrong, so nobody notices until the reel is screened and the hero no longer resembles the poster.

The stronger approach is to send every important shot back to **canonical masters**. The character exists independently of any particular generation. So does the hero prop. So does the location. Story events then create versioned states of those assets: clean, wet, damaged, transformed. A shot requests a state; it never inherits one by accident.

This resembles asset management in VFX, animation and software far more than it resembles casual prompting, and that is exactly why it scales. The disciplines already exist. Animation studios have run model sheets and turnarounds for a century. VFX facilities version every element and track which plate fed which composite.

The current generation of engines is moving in the same direction. ByteDance's Seedance 2.5 accepts up to 30 reference images, 10 video clips and 10 audio clips in a single pass, and can take a textureless 3D "clay render" as a blocking reference.[01] Google's Veo 3.1 accepts reference images of characters, objects and scenes.[02] These are ways of handing the model production context rather than prose.

Reference capacity does not remove the need for canon. It raises the value of having one. A model that can read thirty references will faithfully reproduce thirty inconsistent references if that is what the production feeds it.

○ Our working rule: **approved shots are evidence, never sources**. A later shot may be checked against an approved shot, but it is generated from the canon.

Where the machine still bends

Photorealism is no longer the only threshold. Performance, physical interaction and deterministic control are the harder frontier.

01

Interaction

Hands, prop exchanges, contact, fights and tightly coupled physical events remain disproportionately hard. ByteDance itself lists complex-motion physics and multi-subject interaction as open problems for Seedance 2.5.[01] Breaking action into coverage is often more reliable than demanding a continuous miracle.

02

Performance

A convincing face is not a sustained performance. Eyelines, breath, hesitation and reaction are acting. Google notes that natural spoken audio in short speech segments is still an area of active work for Veo.[02] Motion-transfer systems matter because they import human timing instead of inventing it.

03

Determinism

Generation is probabilistic. Cinema often requires exactness. A logo, a scar, a prop, a costume seam or a subtitle may need classical compositing rather than another expensive reroll. Every reroll also risks breaking what was already right.

> The winning pipeline is not ideological. Generate what AI does brilliantly. Capture what reality does brilliantly. Fix what deterministic tools do brilliantly.

That hybrid position matters because “100% AI” is not necessarily the artistic goal, and it is rarely the cheapest route to a finished shot. A director may photograph London because the real street contains geometry, weathering, reflections and spatial accidents worth preserving. They may record an actor because human timing is worth preserving. They may build a physical pendant because its glass and metal respond beautifully to light.

AI can then amplify those elements beyond budget: turn London into 2098, place the actor inside a kilometre-scale station, transform the pendant, or extend a ten-metre set into a city. The question is not whether the source is synthetic or physical. The question is whether the final shot belongs to the movie.

Editorially, this suggests a discipline: **decide the source of truth before generation.** Is the actor’s movement the truth? The photographed street? A concept-art master? A stunt rehearsal? A shot becomes more controllable when the model is not asked to invent every layer at once.

It is also why conventional cutting remains so powerful. A complex action can be divided into close-up, reaction, insert, wide and impact. Cinema already owns a century of techniques for hiding discontinuity. AI production should use them rather than pretending every scene must be one unbroken demonstration of model intelligence.

Capture	Generate	Cut	Finish
Reality, when material truth matters.	Scale, transformation and impossible worlds.	Editorial grammar simplifies difficult action.	Exact continuity, made deterministic in post.

Eight shifts worth knowing

Issue 001 tracks changes that alter production capability, workflow or rights, not every model release that crosses the feed.

01

31 Jul 2026

Seedance 2.5 moves towards long-form work

ByteDance's update doubles single-pass generation from 15 to 30 seconds of joint audio and video, supports multiple rounds of extension, and accepts up to 30 images, 10 video clips and 10 audio clips as references. It adds clay-render referencing, where untextured 3D blocking drives composition and camera, plus timestamp-level editing. At launch, it rolled out on Jimeng and Doubao, with API access via BytePlus ModelArk announced as coming soon.[01]

02

9 Jun 2026

Luma Ray3.2 starts speaking the language of post

Up to 16 keyframes per clip, performance tracking that follows expressive detail on up to eight faces, native HDR generation with 16-bit EXR export, and a Reframe tool for changing aspect ratio or background after the take. Clips run up to 20 seconds at 1080p, and the full control surface is exposed through an API for the first time.[07]

03

Feb to Jun 2026

Kling builds out the 3.0 family

Kling 3.0 arrived in February with multi-shot storyboards, native audio and 15-second generations. Motion Control 3.0 followed in March, a native 4K output tier in April, and Turbo and Omni upgrades in June. The app now advertises extension to three-minute videos.[04][05]

04

Jan 2026, current

Veo 3.1 remains Google's flagship

Google's model page lists reference "ingredients" for characters, objects and scenes, style references, first and last frame, scene extension, camera and motion controls, object insertion and removal, and 1080p and 4K output. All Veo output carries a SynthID watermark. Google flags short spoken lines as an area still being improved.[02][03]

The pattern

The significant launches of 2026 are not simply "better video". They add **production controls**: references, longer temporal windows, editing after the take, camera and motion direction, native audio and post-friendly output. That is the migration from generative novelty towards a usable department.

For filmmakers the consequence is strategic. Engine choice now belongs in pre-production. The team should know which engine is expected to carry which class of shot, what its failure modes are, and how a shot will be repaired when the first generation is close but not exact.

Editorial note

Fast-moving product claims age quickly. Every specification in this issue uses an editorial cutoff of 20 September 2026 and a first-party source where one exists. Confirm current capabilities, terms and pricing before committing a production.

The landscape is changing around the models

Contracts, regulation and platform stability now belong in filmmaking strategy.

05

Apr 2026, current

Wan 2.7 consolidates Alibaba's video line

The Wan 2.7 family covers text-to-video, image-to-video with an optional driving audio track for lip-sync, reference-to-video from one to five images or clips with voice reference, and instruction-based video editing, at 720p or 1080p and 2 to 15 seconds. Alibaba's published open weights stop at Wan 2.2, which matters for archiving (page 30).[09][10]

06

26 Apr and 24 Sep 2026

Sora leaves the market

OpenAI closed the Sora web and app experiences on 26 April 2026 and is ending the Sora API on 24 September 2026, three days after our cutoff. Data is deleted once any final export window passes.[13]

07

Ratified June, in force 1 Jul 2026

SAG-AFTRA tightens the rules on synthetics

The 2026 TV/Theatrical agreement runs to 30 June 2030. It adds a principle favouring human performance, a "significant additional value" test before a synthetic replaces a human role, a requirement for an articulable business reason before scanning a performer, protections for minors' replicas and a bar on replicas crossing a picket line.[11]

08

Signed 16 Sep 2026

California requires synthetic-performer labels in ads

SB 1050 (Chapter 246, Statutes of 2026) requires a clear disclosure, in wording such as "no human performer is depicted", when an audio or video ad prominently features a synthetic performer. Advertising for films, series and games is exempt when the synthetic performer is used as in the work itself. California follows New York.[12][16]

2026 at a glance

13 Jan	●	Google ships Veo 3.1 "Ingredients to Video" update
5 Feb	●	Kling 3.0 launches
4 Mar	●	Kling Motion Control 3.0
24 Mar	●	OpenAI announces Sora discontinuation
Apr	●	Wan 2.7 family goes live on Model Studio
23 Apr	●	Kling adds native 4K output
26 Apr	●	Sora web and app close
9 Jun	●	Luma releases Ray3.2 and its API
Jun	●	SAG-AFTRA members ratify 2026 TV/Theatrical
1 Jul	●	New SAG-AFTRA terms take effect
31 Jul	●	Seedance 2.5 launches
15 Sep	●	Kling retires six legacy models and 119 effect templates
16 Sep	●	SB 1050 signed in California
20 Sep	●	**Synthetic Cinema editorial cutoff**
24 Sep	●	Sora API ends

Red marks a capability or service being withdrawn. Sources: [01]–[14]. Launches continue after our cutoff; check current documentation before committing a production.

There is no universal winner

A production does not need the “best AI model”. It needs the right camera for each shot, and one dominant visual language across the film.

Model discourse collapses everything into a leaderboard. Filmmaking punishes that simplification. An engine that excels at long, reference-heavy shots may not be the one you want for performance transfer. An engine with beautiful physics may not offer the editability or duration a particular sequence needs. Public arena rankings measure viewer preference on isolated clips; they say little about continuity across a scene, compositing tolerance or cost per usable second.

The practical answer is a **primary camera plus specialists**. Keep most of the movie inside one dominant engine so the film develops a coherent visual interpreter: one way of rendering skin, haze, highlights and motion. Then route genuinely difficult shot types to tools that solve that specific bottleneck.

The cover of this issue asks for Seedance against Kling against Veo against Wan. The honest answer on the next two pages is that they are built for overlapping but different jobs, and that the most useful comparison is the one you run on your own shots. Our protocol for doing that is on page 25.

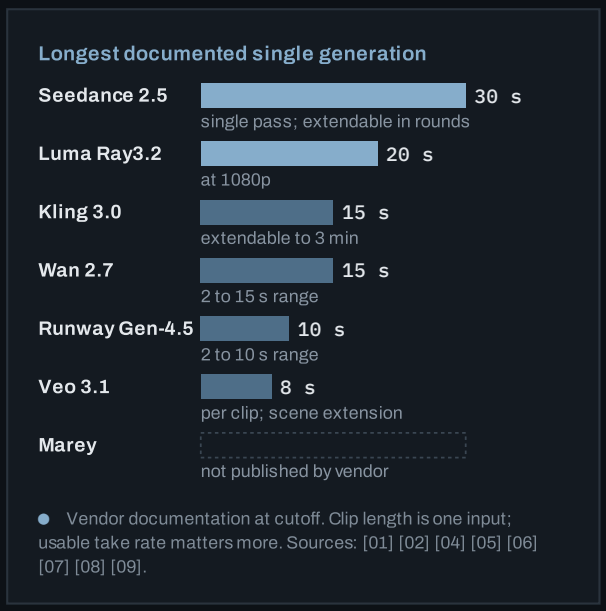

Primary camera

Choose one dominant generator for most narrative shots, so the film develops a coherent look.

Specialist bench

Switch only when a shot presents a capability problem: performance, physical interaction, transformation or exact editability.

Common finish

Grade, grain, sound and compositing pull every source back into one cinematic language.

What each engine is documented to do

Capabilities as published by each vendor at our 20 September cutoff.

Engine	Duration	References	Motion and performance input	Editing and output	Production caution	Src
Seedance 2.5 ByteDance	30 s per pass; multi-round extension	Up to 30 images, 10 video, 10 audio; clay-render blocking	Motion and creative references	Timestamp edits; green-screen, camera-perspective and reference-based editing	Vendor names complex-motion physics and multi-subject interaction as open problems.	01
Kling 3.0 Kuaishou	15 s; extension to 3 min	Image references; element binding (Omni)	Motion Control 3.0	Video editing (Omni); native audio; 1080p and native 4K tiers	Kling 1.0 to 2.1 Master were retired on 15 Sep. Pin and record versions.	04 05 14
Veo 3.1 Google	8 s per clip; scene extension	Character, object and scene ingredients; style reference	Character controls (body, face, voice); motion paths	First/last frame, outpainting, object add and remove; 1080p and 4K	SynthID watermark on all output. Short speech sync still improving.	02 03
Runway Gen-4.5 Runway	2–10 s	Text or single image	Sequenced camera choreography in the prompt	720p base (1280×720); 24 or 25 fps; ProRes or PNG sequence on top plans	Base resolution makes the finishing plan part of engine choice.	06
Luma Ray3.2 Luma AI	20 s at 1080p	Up to 16 keyframes per clip	Performance tracking, up to 8 faces	Native HDR; 16-bit EXR; Reframe; API	Built to be a VFX element, not only a finished clip. Test in your comp.	07
Wan 2.7 Alibaba	2–15 s at 720p or 1080p	1–5 image or video references, voice reference	Driving audio for lip-sync	Instruction-based video editing	Newest open weights are Wan 2.2; 2.7 is API only.	09 10
Marey Moonvalley	Not published	Separate reference image per subject; timeline keyframes	Motion transfer; pose control; trajectory drawing	Camera moves from a single image; shot extension	Vendor says it trained only on licensed footage. Ask for documentation.	08

● / ◐ Every entry is taken from vendor documentation; it describes what the vendor says the tool does, not how well it does it on your shots. “Not published” means we found no figure, not that none exists. Availability, pricing and interfaces change; confirm before committing a production.

> Route the problem, not the brand. If the problem is human timing, use a motion-centric tool. If the problem is a wrong sign, edit the sign.

Reading the table. Duration is the ceiling for one generation, not the length of a usable shot. Reference capacity tells you how much production context you can hand the engine. The output column is where post-production begins: EXR and ProRes drop into a pipeline, while an 8-bit 720p file needs a finishing plan. Watermarks and provenance tags such as SynthID should be recorded in your shot records, because distributors increasingly ask.

What we could not verify

Seedance 2.5 output resolution. The launch post describes duration, references and editing but not a resolution or frame-rate ceiling. Check the platform you access it through.

Marey duration and output format. Moonvalley’s product page describes controls and training data but publishes no clip length or file format.

Pricing, everywhere. We deliberately omit prices. They change monthly, vary by platform and plan, and are the wrong basis for comparison without usable take rates.

Which engine should carry which shot?

A starting routing table built from documented features. It is a hypothesis for testing, not a result.

Shot class	Seedance 2.5	Kling 3.0	Veo 3.1	Runway Gen-4.5	Ray3.2	Wan 2.7	Marey	Documented feature that qualifies it
Long continuous take over 15 s	●	·	·	·	●	·	·	Seedance 30 s pass; Ray3.2 20 s.
Ensemble with locked identities	●	●	●	·	·	●	·	Multi-reference and element binding.
Performance-driven (actor timing)	·	●	●	·	●	·	●	Motion control, character controls, performance tracking, motion transfer.
Blocking from a 3D layout	●	·	·	·	·	·	●	Clay-render reference; 3D camera from a still.
VFX plate for compositing	·	·	·	●	●	·	·	ProRes or PNG sequence; HDR and 16-bit EXR.
Camera move from an approved still	·	·	●	●	●	·	●	Camera controls, prompt choreography, keyframes.
Revising an approved take	●	●	●	·	●	●	·	Timestamp, object, reframe and instruction edits.
Dialogue with lip-sync	●	●	●	·	·	●	·	Native audio or driving-audio lip-sync.

○ KAMX analysis. A dot means the vendor documents a feature that directly addresses the shot class. It does not mean the engine performs that shot well, and a missing dot does not mean it cannot. Issue 002 will replace dots with measured usable take rates.

How to use it

Start by choosing a primary camera from the engines that cover your most common shot classes, usually dialogue coverage and establishing shots. Then look down the table for the classes your script leans on unusually hard. An action film with heavy stunt work needs a performance-driven specialist. A film built on long, unbroken takes needs an engine with a long single pass. A VFX-heavy film needs output that composites cleanly.

Keep the specialist bench small. Every additional engine brings its own colour science, grain, motion texture and failure modes, and each one adds work in the grade.

Five questions before you add an engine

Which shot class fails in the primary engine, and how often?

Does the specialist accept the same canonical references?

Can its output match the primary look in the grade?

What does it cost per usable second, not per generation?

What happens to the shots if the vendor withdraws the model?

Tool of the month

Marey

Why performance transfer could matter more to cinema than another increment in text-to-video beauty.

Moonvalley describes Marey as a foundational model built for professional cinematography. Its documented controls read like a camera department's wish list: camera moves generated from a single image, which the model treats as a 3D scene; **motion transfer**, which pulls movement from a reference video and applies it to a new subject or scene; trajectory control, where you draw the path an object should follow; timeline keyframing between reference images; per-subject references; pose control; and shot extension. [08]

In filmmaking terms the most interesting of these is motion transfer. It changes the directing conversation. Instead of asking a model to invent a specific physical beat, a director can stage the action with a human performer, refine the timing in the real world, and then use the captured movement as production input.

Marey is also sold on its training data. Moonvalley states the model was trained only on licensed footage, with no scraped material or user submissions. [08] That is a vendor claim we have not audited, but it points at a question every producer should now ask of every engine: can you document where the training data came from?

Editorial illustration: a performer directed on a motion-capture stage. It represents the workflow, not a Marey interface.

The return of rehearsal

AI cinema becomes more interesting when the director works with actors before generation rather than after it.

Imagine a spacecraft scene that would normally need a set, a stunt rig, a previsualisation team and VFX replacement. The director can still rehearse the beat in an empty studio: where the actor looks, how fast they turn, where the hand reaches, how the body reacts to impact.

That studio footage is not the finished film. It is the **performance blueprint**. The generated character may wear a different body, suit or species design. The environment may be a kilometre-long alien chamber. But the temporal intelligence, the part audiences read as acting, originated in directed human movement.

This is why performance-transfer systems deserve their own category. They are not simply image generators. They are a bridge between conventional directing and synthetic production, and they bring back something generative video had quietly removed: rehearsal.

Marey is not alone. Luma's Ray3.2 advertises performance tracking across up to eight faces,[07] Kling shipped Motion Control 3.0 in March,[05] and Google lists character controls that drive a character with your body, face and voice.[02] The category is forming quickly. What is missing is evidence about which system preserves timing best when the character, costume and world all change at once. That is what our protocol measures.

A caution for producers: a transferred performance is still a performance. The actor who provided the timing is contributing work that can be reused, and that has consequences for consent and pay. See the rights desk on page 29.

KAMX test protocol

1 Record one actor performing a 12-second beat with a turn, a reach and a reaction.

2 Transfer the motion to three character designs: human, suited human, non-human.

3 Score timing, hand fidelity, centre of gravity, facial continuity and contact with a fixed prop.

4 Repeat with camera motion locked, then with camera motion transferred.

5 Report usable take rate, not the best result.

> A synthetic actor does not have to mean synthetic acting.

Stunts

Capture safe rehearsal motion, then transform scale, environment and character design.

Creatures

Human timing can drive non-human bodies without surrendering the director's blocking.

Reshoots

A preserved performance reference supports controlled changes without rebuilding the physical set.

Real + synthetic

London today, London 2098

The future director can still scout locations, choose light, pick lenses and build props. AI simply makes reality extensible.

A common vision of AI filmmaking puts the director alone at a desk. That is one possible workflow, and perhaps not the most interesting one. Location scouting remains valuable because real places are dense with information that is expensive to invent convincingly: perspective, weathering, reflections, signage, light spill, material imperfection and the spatial relationships between things.

Those details can become the physical substrate for a generated world. Photograph the Thames at dusk and the water, the stone and the sky already agree about where the light is coming from. Extend the skyline into 2098 and the new towers inherit that agreement. Generate the whole city from prose and every element must be persuaded to agree from scratch, shot after shot.

The hybrid method is simple to state. Capture a real location, preserve its light and spatial truth, then extend only what the story requires. The capture becomes the location master; the extension becomes a versioned state of it.

It also changes who is on the crew. A director of photography who understands how light behaves on a real embankment is exactly the person to judge whether a generated extension is lying.

Editorial illustration of the hybrid method: a crew captures a real London location while a tablet previews the extended future skyline.

What to capture before you generate

The real world becomes a reference library. The aim is not to photograph everything; it is to capture what is expensive to reinvent badly.

Character package

Front, three-quarter, profiles, back, full body, hands, expression range, costume states, hair states, hero accessories and movement reference.

Prop package

Turntable views, macro detail, scale in hand, transparent and refractive angles, wear states and interaction reference.

Location package

Wide masters, cardinal directions, entry and exit logic, ceiling and ground detail, day and night, reflections, room tone and ambient audio.

Camera package

Lens family, camera height, intended depth of field, framing rules, movement language, exposure target and practical-light positions.

Rights package

Identity release, digital-replica permissions, permitted uses, territory, term, compensation, vendor limits and deletion obligations where applicable.

Version package

Every hero asset carries a version and a story state. “Pendant.jpg” is not enough. `PENDANT_STATE_03_CRACKED_v02` is production information.

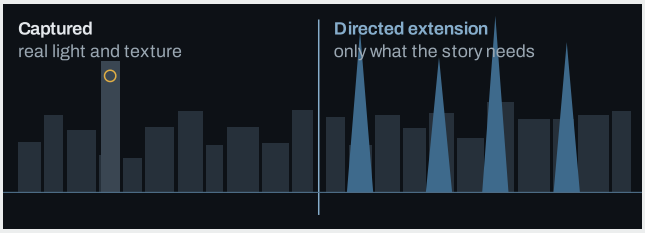

London captured, London extended. The seam is a decision, not an accident.

> The benefit of a high-quality camera image is not just more pixels. It is better physical information: texture, geometry, light, imperfection and material response.

1
Scout
Choose the real spatial anchor and the time of day.

2
Capture
Record plate, texture, ambience and reference movement.

3
Transform
Create and approve the future-world master still.

4
Animate
Generate motion from the frozen canon.

5
Finish
Match grade, grain, sound and continuity.

○ Step 3 is where image models earn their place: a still master is cheaper to iterate than a moving shot, and it can be approved before any video is generated. ByteDance’s Seedream 5.0 Pro is one example of the design-oriented image models aimed at this stage.[15] Next month we test whether still-master-first beats direct video transformation (page 32).

AI actors

The digital double becomes a production asset

Identity, performance, wardrobe, story state, rights and provenance need to be managed as separate layers.

Current video systems increasingly accept character references, and performance-focused systems can separate identity from motion.[02][08] Together they create a powerful but demanding production object: a digital double whose appearance, behaviour and permissions must be tracked throughout the film.

Technical consistency is only one part of the challenge. A face can remain perfectly recognisable and still fail as a character if the eyeline, breathing, timing or emotional arc drifts from shot to shot. Audiences forgive a costume seam long before they forgive a performance that stops belonging to one person.

That is why we treat the double as a stack of layers rather than one file: an **identity master** (what the person looks and sounds like), a **performance source** (whose timing drives the shot), **wardrobe and story state** (what they are wearing and what has happened to them), and a **rights record** (what the scan is permitted to become).

Each layer has a different owner, a different approval path and, increasingly, a different legal status.

Editorial illustration of facial capture and expression reference. The performer depicted is fictional and AI-generated.

Performance is not just data

The most valuable digital-actor pipeline may combine real casting and rehearsal with synthetic scale.

A director could cast a performer conventionally, photograph and scan them with explicit permission, record dialogue and performance references, then place that performance inside scenes that were never physically built.

That does not reduce the actor to a texture. It preserves the very thing generative video often struggles to invent consistently: intentional human timing.

At the same time, a digital replica is a reusable production asset, and its legal state matters. SAG-AFTRA's 2026 TV/Theatrical agreement, in force since 1 July 2026, requires producers to have an articulable business reason before scanning a performer for a replica, adds protections for minors' replicas, and sets a principle strongly favouring human performance, with synthetics used in a human role only where they bring "significant additional value".[11]

The distinction between a **digital replica** (based on a real, identifiable performer) and a **synthetic performer** (a human-seeming figure that is not a recognisable person) is now carried in contracts and, in advertising, in state law. Keep the two categories separate in your asset records from day one.

Production rule

> Scan the actor only after you know what the scan is allowed to become.

Document purpose, term, territory, kinds of generation, permitted models and vendors, security, compensation, and deletion or retention. A clean rights package is part of a clean asset package.

Who owns the performance?

Separate the actor's recorded performance from the producer's generated derivative, and define permitted reuse for each.

Who can process the data?

Vendor choice becomes rights management when biometric or likeness material is involved. Name permitted vendors in the release.

What can be changed?

Age, language, body, wardrobe, sexuality, violence and context may each need separate permission or restriction.

Asset record: the four layers of a double

Layer	Canonical file or record	Approved by	Changes when
Identity master	`CHARACTERS/ILYAN/ID_MASTER_v03`	Director, with performer consent on file	Casting or design changes
Performance source	`PERF/ILYAN/SC71_TAKE04`	Director	A new take is chosen
Wardrobe and story state	`CHARACTERS/ILYAN/STATE_04`	Director; ledger updated by pipeline	A story event changes the character
Rights record	`RIGHTS/ILYAN_PERF_RELEASE_v3`	Producer and legal	Scope, term or vendor list changes

The world build

A feature film needs a memory

Persistent worlds are the difference between a beautiful clip and a believable feature.

Film continuity is usually described as a visual problem. In an AI feature it is better understood as a data problem. The world has state. A door is open or closed. A tower is intact or damaged. A costume is clean or wet. A character has or has not acquired an object.

If those facts live only inside previous generations, the model becomes the continuity department, and it is a poor one: it has no reason to remember and no way to be held accountable. If they live in a production database, shot list or asset tree, the model becomes what it should be: a camera.

Worldbuilding still begins where it always has: with research, architecture, material samples, maps and production design. The generator enters as one tool inside that larger design discipline, not as a substitute for it.

The payoff is practical. A world that is written down can be handed to a second unit, a second engine or a second season without being reinvented.

Editorial illustration: a production designer's atelier. Research, maps and physical models come before generation.

Freeze the universe before you shoot it

The more ambitious the film, the more valuable a production bible becomes.

World bible

Geography, architecture, technology, weather, materials, scale, and the rules the world is not allowed to break.

Character bible

Identity masters, proportions, costume states, props, injuries, movement habits and story-state transitions.

Location masters

Approved spatial anchors from multiple angles. Never regenerate “the city” from prose each time.

Camera bible

Lens families, aspect ratio, movement restrictions, focus behaviour, exposure language and the shots the movie avoids.

Light bible

Key-source logic, colour temperature, contrast range, practicals and atmosphere. Consistent light can unify generations from different engines.

Sound bible

Persistent worlds have an acoustic identity: room tone, machinery, transport, weather, communication systems and silence.

State ledger

Every scene that changes an asset updates the production state. The ledger is what lets shot 646 know what happened in shot 211.

> The generator does not need to remember the film if the film remembers itself.

A state ledger in practice

Scene	Story event	Asset state before	Canonical from here	Valid until
34	North tower destroyed in the raid	`CITY_NORTH_STATE_01`	`CITY_NORTH_STATE_02`	End of film
51	Pendant cracks when Ilyan falls	`PENDANT_STATE_02`	`PENDANT_STATE_03_CRACKED`	Scene 88 (repaired)
62	Coat soaked crossing the flooded deck	`ILYAN_COAT_DRY`	`ILYAN_COAT_WET`	Story logic allows a change
71	Pod A power fails; emergency light only	`POD_A_01 / LIGHT_A`	`POD_A_02 / SCENE_71_KEY_B`	Scene 74

○ Illustrative ledger for a fictional production. A shot that requests a superseded state should fail review automatically.

Post is not dead

Generation is the beginning of the shot

As AI imagery becomes more spectacular, deterministic finishing becomes more valuable, not less.

AI video often arrives with the psychological authority of a finished image. That can be misleading. Different engines have different sharpness, motion texture, highlight behaviour, noise, colour and compression. Cut two of them together and the audience feels the join even if they cannot name it. A feature needs those differences reconciled.

Luma's Ray3.2 is notable precisely because it emphasises frame-level control, native HDR generation and 16-bit EXR output.[07] Runway offers ProRes and PNG-sequence output on its higher plans.[06] Those features speak the language of post-production rather than pretending the first generated file is the end of the process.

The finishing room is where a synthetic film becomes one film. Edit, compositing, colour, sound and mastering are not optional polish; they are how material from several engines, a real location and a motion-capture stage is made to belong together.

Editorial illustration of a finishing suite. Generated plates still pass through edit, compositing, colour, sound and mastering.

4K is not the same as cinema

Resolution is one dimension of image quality. Audiences experience motion, texture, dynamic range, consistency and finishing as a system.

An engine can offer 4K output or upscaling and still produce a weaker theatrical image than a lower-resolution source that survives motion, grading and compositing gracefully. The variables that matter include motion coherence, edge stability, texture persistence, highlight roll-off, skin behaviour, compression, temporal noise, depth cues, and whether the image tolerates reframing or VFX work.

The current engines make very different propositions. Veo 3.1 offers 1080p and 4K output;[02] Kling added a native 4K tier in April;[05] Runway Gen-4.5's documented base output is 720p at 24 or 25 fps;[06] Ray3.2 generates 20-second clips at 1080p but in HDR with EXR export.[07] These are not rungs on one ladder. A 1080p HDR plate with clean highlights may out-grade an 8-bit 4K file that clips.

Two practical consequences follow. First, test on your finishing pipeline, not your laptop screen: push the grade, pull a key, reframe by twenty per cent and project it large. Second, agree a delivery target before generation starts, so every engine's output is judged against the same master format rather than its own best face.

What to measure

Temporal stability. Do textures and edges survive motion?

Grade tolerance. Can the shot take contrast and colour changes without banding?

Compositing tolerance. Can you key, matte and integrate elements cleanly?

Continuity tolerance. Can adjacent shots be made to feel like one scene?

Projection tolerance. Does the shot still hold when it is large?

Myth

"4K means theatrical."

Resolution alone cannot repair unstable texture or weak motion.

Reality

Temporal quality is visible.

A frame can look perfect while the shot breaks as soon as the subject moves.

Finish

Consistency beats isolated sharpness.

Matched colour, grain, motion and sound make different sources feel like one film.

A conform checklist for mixed-engine timelines

- Lock one project frame rate before generation; do not mix 24 and 25 fps sources without a plan.
- Record each plate's engine, version, resolution, bit depth and colour handling in the shot record.
- Prefer EXR, ProRes or image sequences over compressed delivery files wherever offered.
- Build a common grain and texture pass so engines converge in the grade.
- Upscale as a deliberate finishing step, tested per shot, not as a default.
- Log watermarks and provenance tags, such as SynthID, for distribution disclosure.

KAMX test bench

Ten shots. Same references. No excuses.

A useful model comparison should resemble the production you are actually trying to make.

Benchmarks are valuable, but filmmakers need an additional question: **how many usable shots did the system give us before we lost patience, continuity or budget?** The bench below is one short scene broken into ten shots, each chosen to stress a different failure mode. Every engine receives the same canonical references and the same written beat.

Shot	Shot class	What it demands	What it exposes
01	Establishing wide	Location master, time of day, atmosphere	Architecture drift, scale errors, light direction
02	Dialogue two-shot	Two identity masters, blocking, eyelines	Eyeline mismatch, identity swap, lip-sync
03	Emotional close-up	Held performance over 6 to 8 s	Dead eyes, breath and timing, micro-expression
04	Prop hand-off	Hero prop in a defined state, two hands	Hand fidelity, occlusion, prop state
05	Walk-and-talk tracking	Moving camera, moving subjects, dialogue	Camera obedience, background continuity
06	Crowd and ensemble	Six or more consistent extras	Duplicated faces, identity bleed
07	Physical impact	Contact, fall or fight beat	Physics, weight, interaction errors
08	Performance transfer	Captured actor motion onto a new character	Timing loss, centre of gravity, face
09	Night practicals	Low light, practical sources, haze	Noise, banding, light logic
10	Insert with state change	Macro of the pendant before and after cracking	Wrong state, detail migration

○ KAMX protocol. Results will be published in Issue 002 with the full reference package, so readers can rerun it. Header image: an editorial illustration of one face rendered four ways; it is not output from any engine named in this issue.

KAMX test bench, scorecard

Measure usability, not beauty alone

Every engine gets the same canonical references and the same story beat wherever technically possible.

What we record

Metric	Definition
Usable take rate	Takes accepted into the edit ÷ total generations
Rerolls to first usable	Generations submitted before the first acceptable take
Cost per usable second	Total generation spend ÷ seconds accepted into the edit
Continuity defects	Count per take, by class: face, wardrobe, prop, location, light
Interaction errors	Hands, contact, occlusion, exchanged objects
Performance fidelity	Eyeline, timing, body weight, emotional beat, scored 0 to 3
Camera obedience	Lens, framing and move respected as written, scored 0 to 3
Finish load	Minutes of classical clean-up per usable second

What we do not do

We do not cherry-pick one beautiful generation and call it a winner. We do not write different prompts to flatter each engine. We do not treat published resolution as final cinematic quality. We do not accept vendor-curated showcases as evidence.

The result should answer a production question: which engine gives this director the most controllable route to this particular shot?

> The best model is the one that helps finish the film, not the one that wins a beauty contest.

Pass A

First pass

Five generations per engine per shot. No prompt changes unless a technical format requires it.

Pass B

Director pass

One revision round, using the same written notes a human department would receive.

Pass C

Finish pass

Measure how much classical clean-up is needed before the shot can enter editorial.

Why cost per usable second beats cost per generation

Take two hypothetical engines. Engine A produces 30-second clips; one take in ten is usable. Engine B produces 10-second clips; one in three is usable. To bank 30 usable seconds, Engine A needs about ten generations, or 300 generated seconds. Engine B needs about nine generations, or 90 generated seconds.

Unless A is more than three times cheaper per generated second, B is the cheaper route to finished screen time, despite the shorter clip. The numbers here are illustrative. The arithmetic is the point, and it is why we publish usable take rates rather than clip lengths. ○

Build the film like software

Version control is a better metaphor for AI feature production than an infinitely long prompt.

The strongest objection to feature-length generative cinema is continuity. The counterargument is procedural: a later shot can reference the same frozen character and object assets as an earlier one. If fifteen minutes of film can preserve its characters, objects and locations, there is no magic wall at minute sixteen.

What changes over time is not the identity master. It is the **story state**. A character becomes wet, injured or older. A vehicle loses a door. A pendant cracks. Those changes become approved canonical states of their own.

A disciplined project keeps folders or database objects for characters, locations, props, vehicles, creatures, camera, light and sound, each holding masters plus versioned states. Shots reference those objects. If a jacket tears in scene 42, later shots request the torn state; if a window breaks in scene 51, the intact version is invalid from then on.

The same records help everyone. A compositor can find the exact source asset. An editor can see which engine produced a plate. A producer can track reroll cost. A director approves a new state instead of relying on memory.

This is where AI filmmaking stops looking like prompting and starts looking like production engineering.

Minimum shot record

```yaml
# SHOTS/0646/shot.yaml
shot_id: "0646"
scene: 71
beat: "Ilyan finds the pendant cracked"
characters: [ILYAN_STATE_04]
props: [PENDANT_STATE_03_CRACKED_v02]
location: POD_A_02
light: SCENE_71_KEY_B
camera:
  lens: 35mm_equiv
  height: eye
  move: slow_push_in
engine:
  name: seedance
  version: "2.5"
  mode: reference_to_video
references: [CANON_118, ILYAN_04_front,
             ILYAN_04_34, PENDANT_03_macro]
prompt_sha256: "9f2c…e41a"
seed: null  # not exposed by engine
takes: 7
usable: 2
provenance: [vendor_watermark_logged]
rights: [RIGHTS/ILYAN_PERF_RELEASE_v3]
status: APPROVED
approved_by: director
approved_on: 2026-09-14
```

Illustrative record for a fictional production. Hash the exact prompt text so it can be proven later without exposing it.

Asset tree

```
FILM/
├─ CHARACTERS/ILYAN/
│  ├─ ID_MASTER_v03
│  └─ STATE_04
├─ PROPS/PENDANT/STATE_03_CRACKED_v02
├─ LOCATIONS/POD_A/STATE_02
├─ LIGHT/SCENE_71_KEY_B
├─ CAMERA/CAMERA_BIBLE_v02
├─ RIGHTS/
├─ LEDGER/STATE_LEDGER
└─ SHOTS/0646/
   ├─ shot.yaml
   ├─ takes/
   └─ APPROVED/
```

Names are production information

A naming scheme and a state lifecycle cost an afternoon to agree and save weeks of forensic work in post.

State lifecycle

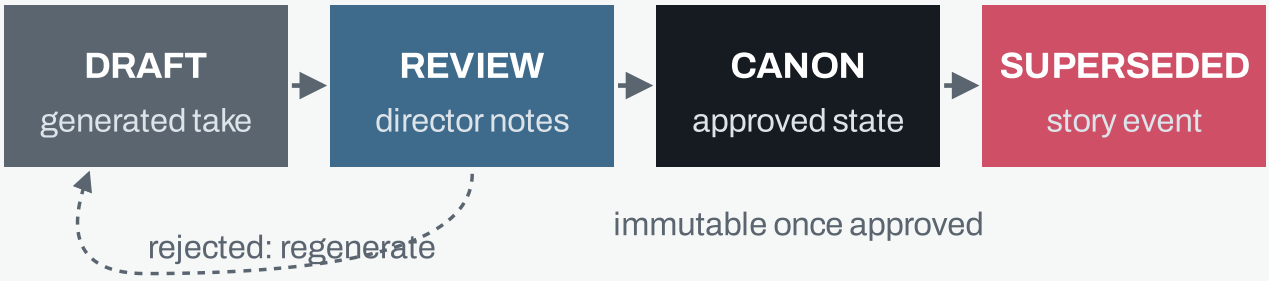

Nothing becomes true in the film until it reaches canon. Canon is never edited; a story event creates a new state and supersedes the old one.

Naming convention

```
<CLASS>_<NAME>_STATE_<nn>[_<DESCRIPTOR>]_v<nn>

PROP_PENDANT_STATE_03_CRACKED_v02
CHAR_ILYAN_STATE_04_WET_v01
LOC_POD_A_STATE_02_EMERGENCY_v03
VEH_SHUTTLE_STATE_02_NO_DOOR_v01
```

STATE changes when the story changes the asset.

v changes when the same state is redrawn or improved.

DESCRIPTOR is for humans; the state number is what software checks.

Never reuse a number, even for a rejected version.

Rules

Approved shots are immutable evidence. New generations reference canon; they never silently redefine it.

State changes need an owner. When a story event changes an asset, the director approves the new state and the ledger is updated in the same session.

Superseded states fail review. A shot after scene 51 that requests PENDANT_STATE_02 should be rejected automatically, not caught in the edit.

Record the engine version, not just the engine. Vendors retire versions; “Kling” is not reproducible, a dated version string might be.

Keep rights on the same tree. A shot record that cannot point to its performer release is not finished.

> Shot 646 is not “continue whatever the AI did in shot 645”. It is “render this exact state of the movie, with these exact references, under these exact scene rules”.

Consent. Compensation. Control.

Synthetic performance is now a contractual production issue, not a theoretical ethics sidebar.

SAG-AFTRA members ratified the 2026 TV/Theatrical agreement in June by 91.42 per cent. It runs from 1 July 2026 to 30 June 2030 and builds on the 2023 digital-replica framework.[11] The provisions producers will meet first:

A preference for human performance. Producers do not intend to use a synthetic in a role a human would otherwise play unless it brings "significant additional value", a term the agreement does not define numerically.

A reason to scan. A producer needs an articulable business reason before scanning a performer for a digital replica.

Minors and strikes. New protections cover minors' replicas, and a performer's replica cannot be used to cross a picket line.

California's SB 1050, signed on 16 September as Chapter 246 of the Statutes of 2026, adds a regulatory layer.[12] [16] An audio or video advertisement that *prominently* features a synthetic performer (demonstrating the product, narrating, or reacting to the narration) must carry a clear and conspicuous disclosure worded substantially like "this performance features a synthetic performer" or "no human performer is depicted". The chaptered text sets no special operative date, so under California's default rule for ordinary statutes it takes effect on 1 January 2027.

The exemption filmmakers need to know. The law does not apply to advertising for expressive works, including motion pictures, television, streaming content and video games, provided the synthetic performer is used in the trailer or promotion the same way it is used in the work. It also exempts AI used only for translation or accessibility. A trailer for your film is covered by the exemption; a brand spot built from your synthetic character is not.

A third question is arriving quietly: training data. Moonvalley markets Marey as trained only on licensed footage.[08] Whether or not a distributor asks today, "what was the model trained on?" is now a reasonable diligence question for every engine on your bench.

Producer checklist

- What exactly is being captured?
- Can new performances be generated from it?
- For which project or franchise?
- For how long, and in which territories?
- Which vendors and models may process the data?
- What compensation applies to reuse?
- How is the asset secured, and who can access it?
- What happens when the agreement ends?
- Can the vendor document its training data?
- Where will synthetic performers need disclosure?

This page is editorial information, not legal advice. Productions should obtain jurisdiction-specific legal and union guidance.

Archive the film, not the vendor

AI filmmaking moves so quickly that model availability has itself become a production risk.

September 2026 made the point twice. On 15 September, Kling retired six legacy video models (Kling 1.0, 1.5, 1.6, 2.0 Master, 2.1 and 2.1 Master), 119 video-effect templates and its Virtual Try-On API.[14] On 24 September, OpenAI ends the Sora API, five months after closing the Sora app and website; user data is deleted once any final export window passes.[13] A production that began in 2025 on either would now be unable to regenerate a single shot in its original engine.

Whatever one thinks of those models, the lesson is general. No filmmaker should assume that a commercial interface, model version or price will exist unchanged for the length of a feature, which can easily exceed eighteen months from development to delivery.

Archive approved outputs locally. Keep source frames, prompts, reference packages, engine and version names, dates, seeds and settings where exposed, edit histories and rights documentation. Preserve final masters in ordinary professional formats that no generator is needed to open.

Open weights are the one partial hedge: a model you can run yourself cannot be withdrawn. They are scarce at the frontier. Alibaba's published Wan weights stop at version 2.2,[10] and every other engine in this issue is available only as a service. Plan on the assumption that you will never be able to reproduce a take exactly, and keep the take.

Production archive

Canon. Approved characters, props, locations and story states, with the ledger.

Shot records. Prompt or prompt hash, references, engine and dated version, cost, status.

Media. Original generations as delivered, plus post-processed masters.

Rights. Licences, performer releases, vendor terms as accepted on the date of use.

Checksums. SHA-256 for every master and evidence file, verified on each copy.

The 3-2-1 rule

Three copies of everything that matters, on two different kinds of storage, with one copy off site. A vendor's cloud library is not one of the three.

Weekly

Export every approved take and its shot record. Verify checksums.

At each engine change

Snapshot the vendor's terms, version string and documentation page.

At picture lock

Render masters to a standard mezzanine format. Freeze the ledger.

At delivery

Package canon, records, rights and checksums as one evidence set.

> A film should own its memory.

Glossary

Twenty terms, defined for production

The vocabulary of this issue, in the sense we use it. Vendors sometimes use the same words differently.

Canonical master
The approved reference version of a character, prop or location. Every shot is generated from it, never from a previous shot.

Story state
A versioned condition of an asset caused by a story event: wet, cracked, destroyed, older.

State ledger
The record of which state of each asset is canonical from which scene onward.

Recursive generation
Seeding each clip from the last frame of the previous one. Smooth joins, cumulative drift.

Drift
Gradual unintended change in identity, wardrobe, props or architecture across generations.

Reference-to-video
Generation guided by supplied images, clips or audio of the subjects, not only by text.

Clay render
An untextured 3D scene used as a blocking and camera reference for generation.

Keyframe control
Fixing chosen frames within a clip so the engine interpolates between approved images.

Motion transfer
Taking movement from a reference video and applying it to a different subject or scene.

Performance tracking
Capturing an actor's posture, gesture and facial expression to drive a generated character.

First and last frame
Supplying the opening and closing images and letting the engine generate the motion between.

Extension
Generating further seconds that continue an existing clip.

Usable take rate
Accepted takes divided by total generations. The core KAMX efficiency measure.

Cost per usable second
Total generation spend divided by seconds accepted into the edit.

EXR
OpenEXR, a high-dynamic-range image format standard in VFX and compositing.

HDR
High dynamic range: more latitude in highlights and shadows, and more room in the grade.

Digital replica
A digital reproduction of a real, identifiable performer's voice or likeness.

Synthetic performer
A human-seeming generated figure that is not a recognisable real person.

Provenance marking
Watermarks or metadata, such as Google's SynthID, that identify content as AI-generated.

Open weights
A model whose parameters are published, so it can be run and archived independently of the vendor.

Next frame

What we are testing next

Issue 002 will be driven by evidence from controlled experiments, not by whichever launch gets the loudest marketing.

Test 01

Real camera to future city

Capture one London location with a high-quality camera and compare direct video transformation against a still-master-first workflow.

Test 02

Performance transfer

Direct one actor through the same 12-second beat and measure how well timing survives character replacement and world transformation.

Test 03

The ten-shot sequence

Run the KAMX bench across the leading engines with identical references. Report usable take rate, rerolls and continuity failures.

Coming in issue 002

Can synthetic actors actually act?

A performance issue: not likeness, not lip-sync, not beauty. Acting. We test timing, eyelines, reaction, physical weight and emotional continuity across a complete dialogue scene.

The scene is two people, one table and one secret, about ninety seconds long. We will stage it with real actors, then attempt it with performance transfer, with reference-driven generation and with text alone, and show the three versions side by side with their costs.

Watch list

Longer multimodal reference windows.

Performance transfer with stronger facial continuity.

Video editing that preserves the successful parts of a shot.

Any successor to Veo 3.1 or Wan 2.7, and what happens to open weights.

First enforcement of SB 1050 after it takes effect on 1 January 2027.

Editorial promise

We will separate documented capability from marketing, and model preference from production evidence. When our tests contradict a vendor claim or our own previous assumptions, the test wins, and we will say so in print.

Source notes

Primary references

First-party and official sources throughout.
All accessed 20–21 September 2026.

[01] **ByteDance Seed.** One-take creation, flexible referencing: introducing Seedance 2.5. 31 Jul 2026.
https://seed.bytedance.com/en/blog/one-take-creation-flexible-referencing-introducing-seedance-2-5

[02] **Google DeepMind.** Veo 3.1 model page: capabilities, controls, limitations, safety. current.
https://deepmind.google/models/veo/

[03] **Google.** Enhanced Veo 3.1 capabilities in the Gemini API. 13 Jan 2026.
https://blog.google/innovation-and-ai/technology/developers-tools/veo-3-1-gemini-api/

[04] **Kuaishou Technology.** Kling AI 3.0 launch announcement. 5 Feb 2026.
https://ir.kuaishou.com/news-releases/news-release-details/kling-ai-launches-30-model-ushering-era-where-everyone-can-be

[05] **Kling AI.** Release notes and version history. current.
https://kling.ai/release-note/release-history

[06] **Runway.** Creating with Gen-4.5: specifications and guidance. current.
https://help.runwayml.com/hc/en-us/articles/46974685288467-Creating-with-Gen-4-5

[07] **Luma AI.** Luma introduces Ray3.2 model and API. 9 Jun 2026.
https://lumalabs.ai/news/introducing-ray-3-2

[08] **Moonvalley.** Marey product page: controls and training-data statement. current.
https://www.moonvalley.com/marey

[09] **Alibaba Cloud.** Model Studio documentation: Wan 2.7 video generation. current.
https://www.alibabacloud.com/help/en/model-studio/

[10] **Wan-Video.** Official open-weight repositories. current.
https://github.com/Wan-Video

[11] **SAG-AFTRA.** 2026 TV/Theatrical contracts: AI and digital replica terms. Jun 2026.
https://www.sagaftra.org/contracts-industry-resources/contracts/2026-tvtheatrical-contracts

[12] **Governor of California.** Governor Newsom signs new law requiring disclosures on AI-generated advertising (SB 1050). 16 Sep 2026.
https://www.gov.ca.gov/2026/09/16/governor-newsom-signs-new-law-to-protect-workers-require-disclosures-on-ai-generated-advertising/

[13] **OpenAI Help Center.** What to know about the Sora discontinuation. updated Sep 2026.
https://help.openai.com/en/articles/20001152-what-you-need-to-know-about-the-sora-app-discontinuation

[14] **Kling AI.** API documentation, Quick Start: retirement of legacy models, effect templates and Virtual Try-On API. notice for 15 Sep 2026.
https://kling.ai/document-api/guides/get-started/quick-start

[15] **ByteDance Seed.** Introducing Seedream 5.0 Pro. 8 Jul 2026.
https://seed.bytedance.com/en/blog/beyond-generation-it-understands-design-introducing-seedream-5-0-pro

[16] **California Legislature.** SB 1050, False advertising: synthetic performers. Chaptered text, Chapter 246, Statutes of 2026. 16 Sep 2026.
https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260SB1050

Editorial cutoff 20 September 2026. Specifications, pricing, availability and contractual terms change. Verify current official documentation before relying on any product or legal detail in production. Editorial analysis throughout remains KAMX and Synthetic Cinema commentary.

How this issue was made

A magazine about synthetic cinema should be explicit about its own synthetic material.

Images

Every photograph-style image in this issue, including the cover, is an AI-generated editorial illustration produced by KAMX Chroma Studios. No person depicted is real, and no image documents a real event, product interface or location shoot. Images were cropped and graded for layout; none was altered to misrepresent a product's capability.

Diagrams and charts were drawn by the editorial team from the sources listed on page 33. Where a chart shows illustrative arithmetic rather than measured data, it says so.

Text

Written and edited by Synthetic Cinema editorial under the direction of the editor. AI tools assisted with research, fact-checking against primary sources and layout. Every factual claim was checked against the source cited beside it; responsibility for errors rests with the editors.

Corrections

We publish corrections in the next issue, on this page, stating what was wrong and what is right. We do not silently edit published issues.

Typography

Archivo, headlines and labels, set wide to echo the masthead.

Source Serif 4, body text, using optical sizes for small settings.

IBM Plex Mono, shot records, file names and data.

Independence

No company named in this issue paid for coverage, supplied free access in exchange for coverage, or reviewed any text before publication. KAMX Chroma Studios makes films with some of the tools discussed; where we judge a tool, we say how we judged it.

Format

A4, 36 pages. Section tabs follow the script-revision colour sequence. Evidence marks: ● verified, ◐ vendor claim, ○ KAMX analysis.

Contribute

Send us the shot that broke.

Synthetic Cinema welcomes production experiments, tool research and AI-assisted film case studies, and we are most interested in the failures.

Case studies

A finished short or sequence, with the pipeline that made it and what it cost in rerolls and time.

Bench results

Run the ten-shot protocol on page 25 with your own references and share the scorecard.

Corrections

Anything in this issue that your production experience contradicts. The test wins.

Submissions for editorial consideration by KAMX Chroma Studios.

KAMX Chroma Studios presents

Synthetic Cinema

Real ideas. Synthetic worlds. Human stories.

Issue 001, October 2026
Published by KAMX Chroma Studios
Editor: Enrico Rudolph Famulari Valentino