
AI's Impact on Video and Game Generation

Mahesh Ramesh
University of Wisconsin-Madison

With
Apart Research

Abstract

This report offers a policy-style forecast for the adoption of AI video generation and AI world or game models in film, television, advertising and games through 2030. We synthesize studio disclosures, platform policies, recent model capabilities and industry surveys. We focus on frontier video models such as OpenAI Sora 2, Google Veo 3.1, Kuaishou Kling 2.5 Turbo, Alibaba Wan 2.1, ByteDance Seedance and Luma Dream Machine, and on world models such as DeepMind Genie 3 that produce interactive environments usable for game prototyping and simulation. Evidence includes Netflix's first confirmed use of generative AI VFX to final, a low-budget national TV ad produced with Veo 3, rapid growth in Steam titles disclosing generative AI, rising M&E maturity for AI use, and wide AI workflow penetration in game development. We combine these signals with a simple adoption-curve model and scenario analysis. Our base case projects that by 2028 at least one major feature or season with a majority of AI-generated footage will release on a top streamer, that by 2030 most scripted productions will include AI in final pixels and that half of AA or AAA games will ship with disclosed generative assets or systems. We report three measurable outcomes: A) probability of a majority-AI feature or season, B) share of scripted productions with AI in final pixels, and C) share of AA or AAA games disclosing generative AI at launch. We discuss risks, labor and IP governance, and benchmark limitations.

Keywords: Video and Game Generation, Economy, model evaluations

1. Introduction

Text-to-video and world models have moved from research to production pilots. Veo 3.1 is offered through the Gemini API with expanded controls and is being integrated into YouTube Shorts. Sora 2 adds synchronized audio and improved physics and steerability with a new social media app. Chinese entrants such as Kling 2.5 Turbo and Wan 2.1 compete on human-preference arenas and VBench, while ByteDance’s Seedance targets multi-shot 1080p video and will probably be added to TikTok. On the world-model side, Genie 3 generates interactive 3D environments in real time at 720p that users can navigate for minutes, which is relevant for simulation and game design.

Studios and platforms have begun limited but public deployments. Netflix confirmed its first final generative AI VFX in *El Eternauta* with a reported ten-fold speed gain. Runway partnered with Lionsgate and launched a film fund to seed AI-augmented projects. Steam requires AI disclosure and one in five 2025 releases reports some generative use. These are early but concrete adoption signals.

2. Research question

When and to what extent will AI video generation and AI world or game models be adopted in professional film, television, advertising and games through 2030. We evaluate three outcomes: A) probability of a majority-AI feature or season, B) share of scripted productions with AI in final pixels, and C) share of AA or AAA games disclosing generative AI at launch.

3. Related work

3.1 Frontier models and capabilities

- Veo 3.1: state-of-the-art text-to-video and image-to-video with richer audio, narrative control and lower pricing; vertical 9:16 and 1080p support for social formats.
- Sora 2: improved physical realism and steerability with synchronised dialogue and sound effects; app and system card describe multi-shot workflows via stitching.
- Kling 2.5 Turbo: top human-preference rankings on the Artificial Analysis arena at release.
- Wan 2.1: reports overall VBench around 86.22 percent and emphasizes motion and spatial relations.
- Seedance 1.0: ByteDance model for multi-shot 1080p generation, positioned for cinematic coherence.
- Dream Machine and Ray 2: Luma’s production-oriented models with published commercial licensing tiers and widely available API access via Fal.
- Runway Gen-3 Alpha: foundation model focused on creative control, deployed with a studio partnership.
- Genie 3: DeepMind’s interactive world model that generates playable environments in real time at 720p, building on Genie and Genie 2.

3.2 Evaluation and risk

- VBench and VBench-2.0: multi-dimension suites covering motion, flicker, identity, spatial relations, physics and controllability with human preference validation.
- Human-preference arenas: crowdsourced pairwise rankings used by vendors and reviewers to compare closed models.
- Misuse and bias: reporting on deepfake risks in Veo 3 content and bias analyses of Sora outputs highlight governance challenges for production deployment.

3.3 Adoption and policy context

- Studios and streamers: Netflix VFX case, Lionsgate-Runway model partnership and Runway’s film fund indicate real pipelines, not only previews.
- Game platforms: Steam AI disclosure policy with rapid growth in titles reporting generative use.
- Labor and law: SAG-AFTRA template deals on digital replicas, California AB 2602 and AB 1836 on performer replicas and post-mortem likeness protection.
- Ad markets: a nationally aired NBA Finals ad made with Veo 3 for about 2,000 dollars shows cost deflation in commercial production.
- Industry surveys: Autodesk M&E 2025 shows high maturity for use of AI in content creation in Film and TV and in Gaming. Google Cloud and Harris Poll report broad AI agent use by game studios. Unity’s 2024 report shows majority AI tool usage. GDC 2025 tracks personal and organizational GenAI usage.

4. Methodology

4.1 Design

We use a policy synthesis approach without new data. Steps:

1. curate recent public adoption signals,
2. map signals to consuming sectors,
3. define three measurable outcomes,
4. fit a two-point logistic curve for each outcome,
5. provide optimistic and regulation-drag scenarios,
6. report year-by-year values.

4.2 Outcomes

- A. Probability that by year t a major studio or streamer releases a feature or full season with at least 50 percent AI-generated or AI-composited screen time.
- B. Share of scripted productions in year t that include AI in final pixels.
- C. Share of AA or AAA games launched in year t that disclose generative assets or live generative systems.

4.3 Calibration anchors

Anchors set the level and slope and are fixed from recent public signals.

A: 2026 = 0.35, 2030 = 0.85

B: 2027 = 0.45, 2030 = 0.75

C: 2025 = 0.20, 2030 = 0.55

4.4 Adoption model

For each outcome we use a logistic curve:

$$f(t) = 1 / (1 + \exp(-k * (t - t_0)))$$

Given two anchors (t1, p1) and (t2, p2) we solve:

$$\text{logit}(p) = \ln(p / (1 - p))$$

$$k = [\text{logit}(p_2) - \text{logit}(p_1)] / (t_2 - t_1)$$

$$t_0 = t_1 - \text{logit}(p_1) / k$$

We then evaluate f(t) for years 2024 to 2032.

4.5 Scenarios

Baseline: use (k, t0) from the anchors.

Optimistic: $k \rightarrow 1.3k$, $t_0 \rightarrow t_0 - 0.5$

Regulation-drag: $k \rightarrow 0.7k$, $t_0 \rightarrow t_0 + 0.5$

4.6 Outputs and reproducibility

We produce annual values for A, B and C for 2024 to 2032 and a milestone table for 2026, 2028 and 2030. The curves are fully reproducible from the anchors and equations above.

4.7 Scope

This is a small-signal, transparent adoption model. It captures smooth diffusion and structures forecasts rather than inferring causal effects.

5. Results

5.1 Near-term capability trajectory

Frontier systems now address key blockers. Sora 2 supports multi-shot stitching and character reuse with synchronized audio and improved physical plausibility. Veo 3.1 adds vertical and 1080p formats and lower pricing, improving accessibility for social and commercial use. Kling and Veo show strong human-preference results on public arenas. Wan 2.1 reports high multi-axis scores on VBench. World models such as Genie 3 provide controllable scene dynamics in real time, which is useful for rapid previsualization and interactive prototyping.

5.2 Adoption signals in 2024–2025

- Netflix moved generative AI from previs into final VFX with a reported tenfold speed gain on a released title.
 - A national TV ad used Veo 3 to air at low cost on a major sports broadcast, demonstrating production deflation and fast turnarounds.
 - Steam’s disclosure dataset indicates about 20 percent of 2025 releases report some generative use, from pre-generated assets to live systems.
 - Industry surveys show high “use of AI in content creation” maturity in Film and TV and in Gaming, and broad AI tool usage in game development workflows.
- These signals justify an adoption curve that is already rising.

5.3 Sector forecasts to 2030

Measure A (probability a majority-AI feature film or full season lands on a major streamer or in wide theatrical by end-year)
Base case: 35 percent by 2026, 65 percent by 2028, 85 percent by 2030. The optimistic case pulls 2028 toward 80 percent, contingent on replica consent clarity and continued API price declines.

Measure B (share of scripted productions that include AI in final pixels, not only previs)
Base case: 40–50 percent by 2027, at least 75 percent by 2030. This follows from 2025 maturity levels, public final-pixel case studies and rapid integration into editorial and DCC toolchains.

Measure C (share of AA or AAA games with disclosed generative assets or systems at launch)
Base case: 35–45 percent by 2027 and at least 50 percent by 2030, consistent with disclosure trends and enterprise workflow surveys.

Table 1. Forecast milestones (baseline)

Table 1. Forecast milestones (baseline)

Year	Measure A: P(majority-AI feature/season)	Measure B: share with AI in final pixels	Measure C: share of AA/AAA with disclosed gen-AI
2026	0.35	0.35	0.26
2028	0.64	0.56	0.39
2030	0.85	0.75	0.55

These values are generated by the logistic fits defined in Section 4.4. Scenario brackets at 2030 are approximately: $A = 0.73\text{--}0.93$, $B = 0.65\text{--}0.85$, $C = 0.51\text{--}0.61$.

5.4 Economic implications

Cost and schedule. On-air and studio examples show large cost and time reductions, including tenfold VFX speed improvements on specific sequences and low-budget national-air spots. Continued API price cuts and faster inference point to further marginal-cost declines.

Market growth. AI in media and entertainment is forecast to approach 100 billion dollars by 2030. AI in gaming is projected to grow at roughly 20–36 percent compound rates over the decade, depending on scope.

Labor and skills. Digital-replica agreements and state laws make consent and reporting standard. Studios and unions are converging on disclosure and compensation norms, while hiring shifts toward AI-native roles in post, localization and asset pipelines.

6. Conclusion

AI video generation and world models are moving from demonstrations to production-relevant tools. Public studio deployments, rising game disclosures and maturing evaluation indicate a clear adoption path. Under a conservative reading of current signals, majority-AI releases are likely by 2028 and broad integration into final pixels and shipped games is likely by 2030, subject to governance around replicas, training data and disclosure. Procurement and broadcast standards should reference multi-axis evaluations and require crediting, consent and provenance so adoption scales safely and transparently.

References

1. DeepMind. Genie 3 announcement and capability details. [Google DeepMind](#)
2. DeepMind. Genie 2 foundation world model. [Google DeepMind](#)
3. OpenAI. Sora 2 announcement. [OpenAI](#)
4. OpenAI. Sora 2 system card. [OpenAI](#)
5. Google. Veo models page and benchmarks. [Google DeepMind](#)
6. Google developer blog. Veo 3.1 updates in Gemini API. [Google Developers Blog](#)

7. The Verge. Veo 3 vertical format, 1080p and pricing changes. [The Verge](#)
8. Artificial Analysis. Video arena human-preference leaderboard. [Hugging Face](#)
9. Alibaba Cloud. Wan 2.1 VBench overview. [Alibaba Cloud](#)
10. Reuters. Alibaba to open source Wan 2.1. [Reuters](#)
11. Kuaishou Kling 2.5 Turbo coverage and arena standing. [aastocks.com](#)
12. ByteDance Seedance 1.0 model page. [seed.bytedance.com](#)
13. Luma. Dream Machine licensing and plans. [lumalabs.ai](#)
14. Fal. Dream Machine API access. [Fal.ai](#)
15. Runway. Gen-3 Alpha research page. [runwayml.com](#)
16. Autodesk. 2025 Spotlight on AI in Media and Entertainment. [damassets.autodesk.net](#)
17. Google Cloud. 2025 Games report summary page. [Google Cloud](#)
18. Game Developer on Unity 2024 report. [Game Developer](#)
19. GDC 2025 State of the Game Industry articles. [gdconf.com](#)
20. Netflix and El Eternauta AI VFX reports. [The Verge](#)
21. The Guardian on Netflix El Eternauta. [The Guardian](#)
22. Business Insider on the Veo 3 NBA Finals ad. [Business Insider](#)
23. The Verge on the Veo 3 NBA Finals ad. [The Verge](#)
24. Tom's Hardware on Steam AI disclosures in 2025. [Tom's Hardware](#)
25. VBench original paper and CVPR page. [arXiv+1](#)
26. VBench-2.0 paper. [arXiv](#)
27. TIME reporting on Veo 3 misuse risks. [TIME](#)
28. WIRED on Sora biases. [WIRED](#)
29. SAG-AFTRA Replica Studios agreement. [SAG-AFTRA](#)
30. California AB 2602 and governor release. [Governor of California](#)

31. California AB 1836 commentary. [Davis+Gilbert LLP](#)
32. AP News on SAG-AFTRA video game performers contract. [AP News](#)
33. Lionsgate-Runway partnership. [investors.lionsgate.com](#)
34. Runway Hundred Film Fund. [runwayml.com](#)
35. Moonvalley Marey announcement and positioning. [Business Wire](#)
36. The Guardian on Lionsgate AI partnership. [The Guardian](#)
37. Business Insider on Moonvalley funding. [Business Insider](#)
38. DeepMind Genie paper. [arXiv](#)

APPENDIX

7. Discussion

Quality and safety constraints

Benchmarking is improving. VBench decomposes video quality into motion, identity consistency, spatial relations and related axes, and VBench-2.0 adds physics and controllability. Human-preference arenas complement these suites when unified scores are unavailable. These are better policy gatekeepers than caption-similarity metrics alone.

Misuse and bias

Frontier models can produce highly realistic but misleading content, and stereotyped portrayals can persist in some outputs. Watermarks and provenance signals help but can be degraded. This argues for clear disclosure, provenance metadata and pre-release impact assessments in news and political contexts.

Governance and IP

The regulatory environment is maturing. Contract templates and state laws address consent and post-mortem likeness for digital replicas. Studio initiatives that train on licensed or owned data provide a compliance path for production use. Procurement policies that require disclosure, consent logging and provenance metadata can scale safely across productions.

Limits of the forecast

This is a small-signal adoption model. It assumes smooth diffusion once technical and legal thresholds are crossed. Adoption could lag with trust shocks, litigation or procurement pauses. Adoption could accelerate with deeper toolchain integration, reliable controls and further cost declines.