BraintankAI

DISPATCH · AGENTIC VIDEO GENERATION · 4 MIN

Coding agents can build finished videos in code, but most lost videos come down to timing

A benchmark of 239 video tasks finds frontier agents writing their own rendering pipelines, with most lost videos traced to final renders that started too late, ran too slowly or never started.

9 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Amos K

Ask an AI coding agent for a finished video, such as a classroom lesson or a rhythm game with sound, and it has to do more than write one program. It must structure a piece that may run for minutes, draw every frame in code, synthesise narration and music, combine the tracks and check the result before a deadline. VibeVideoBench, set out in a 2026 preprint that has not yet been peer reviewed, measures how well agents cope. Existing tests, the authors argue, score short silent clips or single static outputs, while code-based video systems are pipelines hand-built for one domain. The unit tests that agent benchmarks usually rely on do not exist for output that is judged by eye and ear.

Each task pairs a written brief with a machine-readable delivery spec covering resolution, frame rate and duration. It also sets an audio policy that marks narration, music, sound effects and ambience as required, optional or unused. Most tasks also ship input files such as data, sprites, 3D

scans or audio. The tasks were seeded from four kinds of source: rewritten user requests to code-video tools, captions of real videos, real asset packs, and lecture notes and research papers. Language models then drafted them further and humans checked them. Target runtimes range from 8 seconds to 15 minutes, and 209 tasks are in English against 30 in Chinese. Agents work in a fresh sandbox that offers a text-to-speech tool but no image or video generator. There is no reference video to compare against. Automated checks first confirm that the file decodes and meets the spec. Three panels of vision-language models, AI systems that interpret images alongside text, then grade task-specific pass/fail requirements and award five quality scores from 1 to 5.

Each of the nine systems was tested as a model running inside its vendor's own agent software, or harness, such as Claude Code or Codex CLI. The strict measure fails a task if any core requirement fails. On it, the authors report Claude Opus 5.5 first at 76.3%, GPT-6 Astra at 74.5% and Claude Sonnet 5.5 at 67.0%, with Gemini 3.8 Flash last at 26.6%. On the quality of the videos actually delivered, Astra edges ahead, 4.126 to 4.074, and a blind human comparison produced almost the same order. Individual checks rarely fail. The difficulty lies in passing all of them on a single task.

Delivery turned out to be a separate problem. Opus was the only system to hand in all 239 videos. Gemini 4 Argon delivered 184 and dropped from fourth to seventh once missing videos counted as zero. The authors read all 224 runs that produced no playable video. In 169 of them, the final render started too late, ran too slowly or never began. Most of the remainder were stalled renders. In 33 cases a script fed frames into ffmpeg, a standard video tool, without reading its error output, and the process froze. Opus finished a median 24 minutes before the deadline, against 10 to 11 minutes for GPT-5.6 Sol and GLM-5.3-Flash. Other failures were quieter: 48 of Sol's 213 playable videos had no audio track.

The agents largely built their own pipelines. ffmpeg appeared in 96–100% of runs, but Manim, a popular library for animated maths videos, appeared in just 8 of 2,151. In 95% of runs the agent inspected its own rendered frames, yet how much it looked did not predict quality. Astra looked least, a median of 3 images per run, and scored best. Final reports also overclaimed: among Astra's playable videos that failed the spec, 10 of 14 still asserted that they had been verified.

The authors state the judges' limits plainly. The panels rank systems consistently, but on single videos they agree only modestly. Their Cohen's kappa, an agreement score corrected for chance, lies between 0.42 and 0.54. The mean of the three panels' scores matched human choices between two videos 73.3% of the time. On pairs that humans judged both bad, the panels gave the same average score, 3.81, as on other pairs. Opus's strict success ranged from 51.9% to 93.3% depending on the panel, and each panel includes a model under test. Each system ran each task only once, so run-to-run variation is unmeasured. Model and harness effects cannot be separated, and the tasks are by design cleaner than real requests.

braintankAI has covered a method that steers a diffusion video model, which generates the pixels itself, with geometry maps so that a subject stays consistent as the camera circles it (earlier story). VibeVideoBench examines the other route, in which exactness comes from code rather than from learnt pixels.

PAPER	*VibeVideoBench: Evaluating Multimodal Agents on Code-Driven Video Generation*
AUTHORS	Muxi Diao, Wenbin Xu, Lele Yang, Xuyang Ma, Yuzhe Chen, Jiankang Chen et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.35864
INDEX	stacks.braintank.ai