BraintankAI ASIA

DISPATCH · VIDEO CAPTIONING · 3 MIN

BEIJING · Peking University

Benchmark asks AI models to caption videos as code that can redraw them

CodeCap-Bench scores how faithfully multimodal AI models write programs that re-render 323 short 2D videos, and finds motion consistently harder to capture than static detail.

12 November 2026 · arXiv preprint, not yet peer reviewed

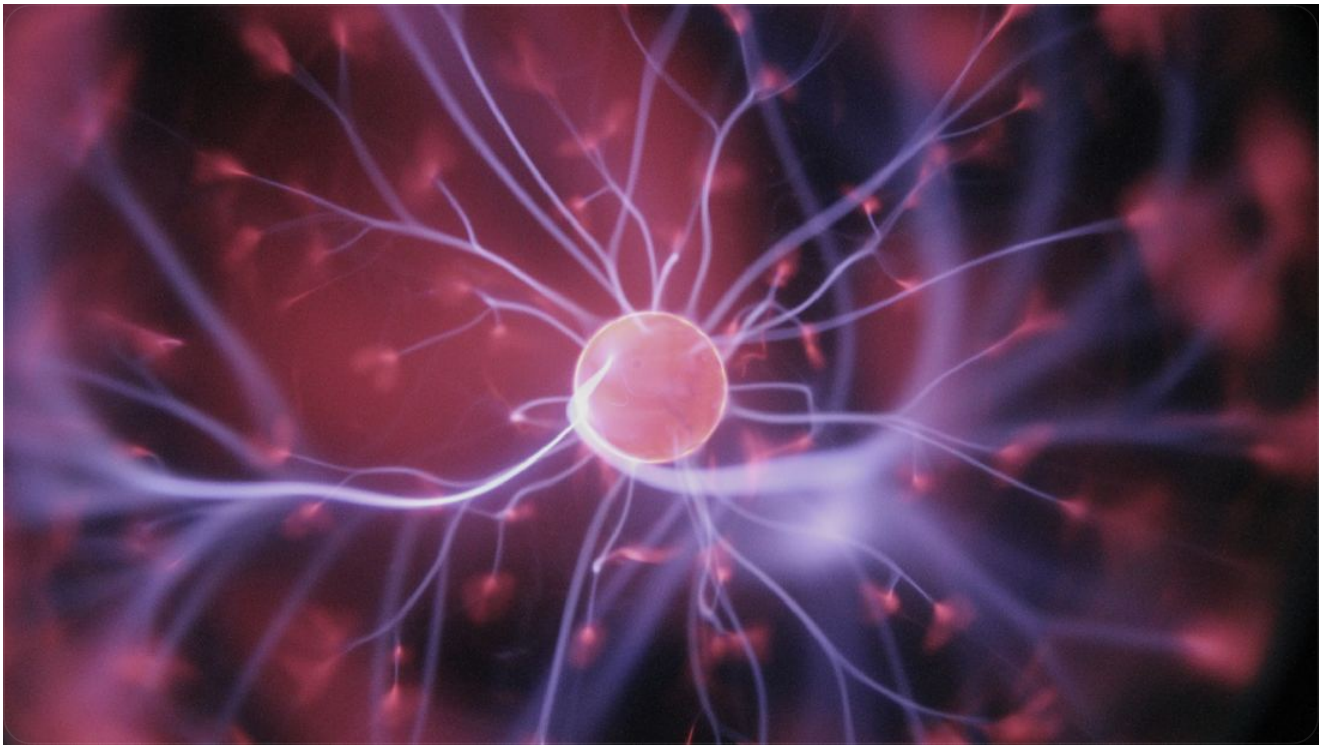


Photo: Hal Gatewood

A video caption is meant to preserve what a clip shows, but ordinary sentences are clumsy at recording exactly where things sit, how they move and when they change. A 2026 preprint, which has not yet been peer reviewed, tests an alternative: describing a video with an executable program that draws its shapes, text and layout and expresses motion as functions of time, so that running the code should reproduce the clip. One illustrative example in the paper uses a smoothing curve to govern a scroll, showing how code can fix an element's position at any instant. Because such a caption can be rendered and set beside the original, its fidelity can be checked directly. Until this work, the authors say, there was no dedicated test of whether multimodal large language models (MLLMs, systems that read images or video as well as text) can write such captions.

Their benchmark, CodeCap-Bench, draws 323 videos from platforms including Bilibili, YouTube, Behance and Vimeo, spread across five categories of structured 2D content such as desktop and app interfaces and data visualisations. Clips dominated by complex 3D viewpoints, fine textures or dense pixel-level motion were filtered out. The videos average 19.0 seconds, against 5.1 seconds for an earlier animation-to-code benchmark, Animation2Code. Each carries on average 16.8 annotated objects and 8.6 human-reviewed focal points: details worth checking, either static (appearance, layout, text) or dynamic (state changes, movement, timing). Models must work in a pure-code setting, drawing everything themselves rather than pasting in frames from the original.

Scoring has three layers. A pretrained video model, V-JEPA 2.1, compares the overall look of original and reconstruction. An object-level measure checks, frame by frame, how well each object's bounding box overlaps between the two, using GPT-6-sol to locate objects and the segmentation model SAM 3.1 to outline them. Finally GPT-6-sol grades every focal point from 0 to 5. Code that fails to render scores zero on everything. On outputs from four models for 50 reference videos, the AI judge's scores tracked human ratings with a Pearson correlation of 0.942, where 1 would mean perfect agreement.

Seven models were tested, either in one shot from frames sampled at 2 per second, or as agents inside tool harnesses (Pi Agent, Codex, Claude Code) that let them inspect frames, render drafts and revise. GPT-6-astra led on overall retention in both settings; with Codex it reached an overall retention score of 73.24 and a 99.4 per cent render success rate, and the authors report a 14.40-point lead over the runner-up, Opus-5 with Pi Agent at 58.84. Rendering is not the same as fidelity: GPT-6-sol with Codex rendered every video yet scored 51.34. In every configuration dynamic details fared worse than static ones, 68.61 against 80.72 even for the leader. Agent harnesses raised success rates but did not reliably raise fidelity; GPT-6-sol scored 53.08 without one, above its Codex and Pi Agent results. Raising GPT-6-sol's thinking effort from medium to xhigh lifted retention from 47.53 to 63.29 while average cost per video rose from USD 0.315 to 1.409. More tool calls did not mean better results: Gemini-3.8-flash made 99.6 tool calls per video and scored 47.73, while GPT-6-astra made 47.4 and scored 72.87.

To test the premise itself, the authors compared code and prose captions. On 100 videos from three existing question-answering datasets, with Gemini-3.8-flash answering from either the rendered code or the text, Opus-5's code captions gave 83.02 per cent accuracy against 72.64, close to the 83.96 obtained from the original videos; GPT-5.6-sol's gap was narrower, 76.42 against 73.58. In reconstruction, where prose captions were turned into video by the generator Seedance-2.5, code won clearly on appearance but only narrowly on a motion measure.

The design limits what these comparisons show. The benchmark deliberately keeps content suited to programmatic drawing; GPT-6-sol is both a contestant and the main judge, and swapping judges was checked on only two models; no human check of the automated object-overlap pipeline on rendered videos is reported; and no repeated runs or significance tests are reported. The question-answering test feeds the answering model video on one side and text on the other, and the reconstruction test measures prose captions together with a video generator rather than the captions alone.

The weakness on motion echoes another braintankAI report, in which video models trained on spatial data handled some relations better than others until their training was restructured (earlier story). The idea of a single executable description also sits alongside Tencent's effort to parse documents, images and video into one shared output format (earlier story).

PAPER	*CodeCap-Bench: Can Code Serve as a Good Video Caption?*
AUTHORS	Jiafu Tang, Yifan Xu, Eric Tang, Junjie Zhao, Shipeng Luo, Te Sun et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.37341
INDEX	stacks.braintank.ai