BraintankAI

EXPLAINER · VIDEO WORLD MODELS · 5 MIN

AlayaWorld builds a playable game world one second of generated video at a time

A 2026 preprint fine-tunes a video model to draw each next moment of a world from the player's camera moves and typed actions, and shows its results only as example footage.

21 November 2026 · arXiv preprint, not yet peer reviewed

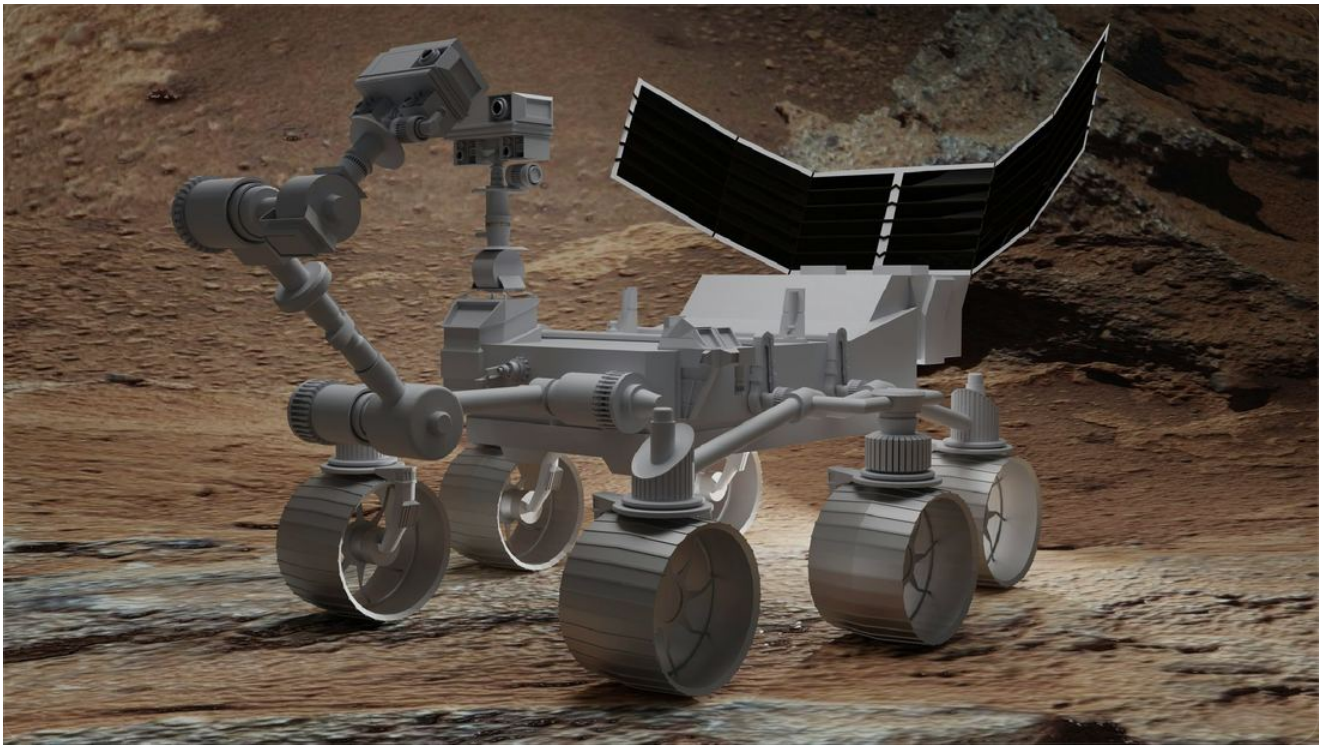


Photo: Sufyan

Every conventional video game world is built by hand. Artists model the buildings, designers place the doors and enemies, and programmers write the rules that decide what happens when a player turns left or swings a sword. That work is expensive, and once a game ships, the world is hard to change. A different idea has been gathering attention: skip the hand-building, and have a model simply draw the next moment of the world in response to what the player has just done.

That raises some obvious questions. Can a model that invents pictures as it goes let a player walk anywhere and do anything? Will a street look the same when the player wanders off and comes back? Will the images hold together for minutes rather than seconds, and can all of it happen fast enough to feel like play? A 2026 preprint from a group calling itself the AlayaWorld Team, posted to arXiv and not yet peer reviewed, is a useful way to see what answering those questions involves.

The idea

The system at the heart of this approach is a video world model: a model that produces the next frames of video in response to user input, so that it acts as both the game engine and the renderer at once. There are no hand-written rules or physics code underneath. The model looks at the current state of the scene and the player's latest command, and generates what should come next. It works autoregressively, meaning it produces video piece by piece, with each new piece conditioned on what it has already made.

The authors set out four problems that stand between this idea and something genuinely playable. The first is control: whether players can move without limit and take arbitrary actions. The second is consistency: whether a place looks the same on a return visit. The third is stability: whether the model can keep generating for a long time without drift, the gradual decay in quality that sets in as small errors feed back into later frames and compound. The fourth is runtime: whether generation can happen in real time with little delay. The authors note that one well-known system, Genie 3, achieves real-time play with coherence over minutes but is closed-source, while open or related efforts such as Matrix-Game, Hunyuan-GameCraft and Lingbot-World tackle parts of the problem.

The paper's own survey of earlier work is organised around those four problems, and it frames each as a trade-off. Feeding the camera position into a model as a simple signal is flexible but only as accurate as the model's ability to interpret it. Showing the model a rendering of what the new viewpoint should look like is more accurate but needs a heavier pipeline that depends on estimating geometry. AlayaWorld is pitched as a way of combining complementary approaches rather than as a single invention.

How this paper puts it to work

AlayaWorld starts from an existing video generator, LTX-2.3, and fine-tunes it together with modules of the authors' own design. The result is an autoregressive diffusion transformer, a type of neural network that creates images or video by gradually removing noise from a random starting point. It generates video in chunks of roughly one second each, at 720p resolution and 24 frames per second, using four denoising steps per chunk; each step is one pass of noise removal, so fewer steps mean faster output.

For camera control, the system keeps a 3D cache, a three-dimensional reconstruction built from earlier frames, following an earlier method called GEN3C. When the player moves, that cache is re-rendered from the new camera angle, giving the generator concrete visual evidence of what should be visible there. Alongside this, a compact description of the camera's position is fed into the network through AdaLN modulation, a lightweight technique that nudges the scale and offset of the network's internal normalisation layers and adds little extra cost. For actions beyond moving the camera, such as fighting or casting a spell, the player changes the text prompt, and the switch takes effect at the next chunk boundary without regenerating anything already shown.

Memory comes in two kinds. The 3D cache supplies spatial memory of places already visited. Recent frames are also compressed into a small summary, following a method called Frame Preservation, to capture recent motion and changes that a static reconstruction cannot hold. To fight drift, the model is trained on histories that have already degraded, following a method called

Helios, so it learns to continue sensibly from imperfect input. It also uses an error bank, a store of residual artefacts that build up during generation, an idea earlier used in Stable Video Infinity; the authors reinject these artefacts during training into both the memory the model reads from and the segment it must produce. Finally, to cut the number of denoising steps, they use DMD distillation, which trains a fast student model to match the output distribution of a slower teacher model in far fewer steps.

What it shows, and what it does not

Every result in the paper is qualitative: example sequences shown as figures, accompanied by the authors' descriptions. They show the camera following varied commands, prompts switched mid-generation, trajectories that leave a place and return to it, rollouts lasting one minute, and a single route through a scene rendered in seven styles, including ink painting, pixel art, cyberpunk and a Minecraft look. The authors report that revisited areas keep their geometry, layout and texture, and that the one-minute runs hold up without pronounced build-up of artefacts. In the return-visit comparison, they report that earlier interactive world models show degraded images, inaccurate camera control and inconsistency.

Those claims are the authors' own reading of their examples. The paper reports no numerical results of any kind: no quality metrics, no benchmarks, no user studies. The text does not name the models it compares against in the return-visit test, and the authors say matching conditions were used only where the other systems' public code allowed. There are no ablations, the experiments that remove one component at a time to show which pieces matter. The system is described as real-time, yet the paper gives no hardware, generation speed or end-to-end delay. The longest demonstration is one minute. Most of the building blocks come from earlier published work, so the contribution lies chiefly in how they are combined.

The paper's survey also flags weaknesses in the kinds of techniques AlayaWorld relies on. Spatial caches mainly capture static structure, struggle with moving objects and depend on estimating depth and geometry. Aggressive distillation to a few steps can weaken the fidelity and variety of motion, and generating in a stream can leave visible seams between chunks. Prompt changes apply only at chunk boundaries, and the authors say commands take effect after a short delay. The authors present the paper as a preview, with full technical details, results and code promised in a later release, which makes it best read as a clear map of the problems in this field rather than proof that they are solved.

PAPER	*AlayaWorld: Long-Horizon and Playable Video World Generation*
AUTHORS	AlayaWorld Team, Kaipeng Zhang, Chuanhao Li, Yifan Zhan, Yongtao Ge, Yuanyang Yin, Jiaming Tan, Kang He, Liaoyuan Fan, Ruicong Liu, Xiaojie Xu, Xuangeng Chu et al.
VENUE	arXiv preprint — not yet peer reviewed · 7 July 2026
LINK	arxiv.org/abs/2607.06291
INDEX	stacks.braintank.ai