DAEJEON · KAIST

Compact 3D memory helps a video world model stay consistent when the camera returns

A KAIST AI preprint stores a generated scene as tens of thousands of 3D blobs carrying the video model's own features, cutting memory and rendering costs.

7 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Eli Alvarez

An interactive world model is a video generator that shows what a virtual camera would see as a user steers it. A central challenge is memory. A camera that returns after a long excursion should find the same scene it left, a test called loop closure. A 2026 preprint from KAIST AI, not yet peer reviewed, describes L3G-World. It builds an explorable world from a single image and keeps that world consistent with a compact 3D memory the generator can read directly.

The authors argue that earlier systems take one of two routes. Some feed the generator a growing history of past frames. Others keep an explicit 3D memory made of points or Gaussians, which are small fuzzy 3D blobs with position, shape, colour and opacity that render quickly from any viewpoint. The fast versions of these memories pile up redundant copies of the same regions and do not update what they have stored, so bad geometry can persist. These explicit memories also

hand the scene over in a slow way. They render an ordinary image and re-encode it through a VAE, the autoencoder that compresses frames into the compact latent form a video diffusion model works on. The authors say this step rules out real time.

L3G-World alternates a memory updater with a video generator. For each new window of frames, the updater renders its memory into the new camera views. Pixels where the memory is too transparent or gets the colour wrong form a 'spawn mask'. New Gaussians appear only there, seeded from geometry features of the Depth Anything 3 model. Gaussians already in view receive predicted corrections. Each Gaussian also stores a VAE latent, so rendering the memory yields input in the generator's own working space. The generator's output then flows back into the memory without re-encoding. The generator is a diffusion transformer, a transformer that turns noise into video step by step. It is built on Cosmos-Predict2.5 and fine-tuned with LoRA, a cheap method that trains small add-on weights. A 'Fast' variant is distilled, meaning a student model is trained to mimic the full one in only a few denoising steps.

The generator was trained on two scene datasets, DL3DV and RealEstate10K, and the memory updater on DL3DV alone. Testing used the same two datasets. The test runs used 481-frame camera paths that travel away from the opening view and then retrace their route. The comparison covered eight systems, including Yume 1.5, Matrix-Game 3.0, GEN3C, Lyra 2.0 and the Gaussian-memory method NeoVerse. Every explicit-memory method was given the same Depth Anything 3 depth estimates, and all were timed on one A100 GPU. On DL3DV, the authors report 33K primitives against Lyra 2.0's 1,273K, and 11.88 MiB of memory against 38.16 MiB for the leanest baseline, SPMem. Rendering memory into the generator's latent space ran at 1,086.58 frames per second, against 3.11 for NeoVerse. On loop closure, L3G-World scored a PSNR, a pixel-level similarity measure, of 16.00 against 13.96 for the best baseline, Spatia. An ablation isolates the latent trick: direct rendering ran at 943.6 frames per second against 2.97 for rendering and re-encoding.

The lead is not uniform. MEt3R measures how consistent generated frames are in 3D, and lower is better. L3G-World is not best on this measure on either dataset. On DL3DV it scored 0.078 against 0.050 for Yume 1.5, although the text concedes a shortfall only on RealEstate10K. On RealEstate10K, Lyra 2.0 and GEN3C also follow the requested camera path more accurately. The paper's real-time figure of 9.02 frames per second comes from the Fast variant running on an H100, not the A100 used for the comparisons. That variant also gives back some quality. FID measures how far generated images sit from real ones, and on RealEstate10K the Fast variant scored 49.94 against the full model's 36.09, putting it behind three baselines. The setup leaves other gaps too. Testing is in-domain on static scenes, and the only action is camera movement. The 'infinite' of the title rests on 481-frame runs, and memory still grows with the area explored. No run-to-run variance is reported.

braintankAI has covered Tencent's Hy-Nomas, a video world model that likewise has to keep track of what has already appeared on screen as the camera moves (earlier story). Another KAIST AI preprint steered video generation with geometry maps rebuilt from reference photos, keeping a subject consistent as the camera circles it (earlier story). L3G-World makes the same bet: that explicit 3D structure helps a video generator stay coherent.

PAPER	*L3G-World: Real-Time Infinite World Generation with Latent 3D Gaussian Memory*
AUTHORS	Jinhyeok Choi, JoungBin Lee, Jongmin Lee, Minkyeong Jeon, Jiyoung Kim, Kaeun Han et al.
VENUE	arXiv preprint — not yet peer reviewed · 29 October 2026
LINK	arxiv.org/abs/2610.40614
INDEX	stacks.braintank.ai
CODE	github.com/cvlab-kaist/L3G-World