BraintankAI ASIA

EXPLAINER · VIDEO WORLD MODELS · 5 MIN

HONG KONG · Hong Kong University of Science and Technology

An open recipe for turning a video generator into a camera-steered world model

Researchers release every stage of a pipeline that teaches video models to follow camera moves and generate chunk by chunk, cutting the wait for a first frame from minutes to seconds.

14 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Mitch Nielsen

A video generator can produce a convincing clip of a walk through a forest, but nobody can take that walk. You type a prompt, wait, and receive a finished film. To behave like a place you can explore, a model would have to produce frames as you watch, respond when you turn the camera, and do it quickly enough that the delay does not break the spell. Researchers call such a system an interactive world model. The distance between one and an ordinary video generator is the subject of minWM, a 2026 preprint on arXiv that has not yet been peer reviewed.

Its authors, Min Zhao, Jun Zhu and 14 colleagues, do not claim a fresh algorithm. Their case is that the techniques for turning a video model into a steerable, real-time one already exist but sit in separate pipelines. They argue that no unified, reproducible framework covers the whole journey

from training data to a working interactive model. minWM is their attempt to supply one, openly.

The idea

Most high-quality video generators are diffusion models. They begin with random noise and remove it over many steps until a clip emerges. The models the authors start from are also bidirectional, meaning every frame is refined together and each can draw on frames before and after it. That helps quality, but nothing is ready until the whole clip is done, and there is no moment at which a user could step in.

An interactive world model needs a different shape. It must be causal, or autoregressive: it generates video a chunk at a time, with each chunk conditioned only on what came before. Playback can then begin before the whole clip is finished, and the next chunk can reflect what the user just did. It must also respond to actions. Here that means a camera trajectory, the path and orientation of a virtual camera moving through the scene. Finally, it must keep latency, the delay before output appears, low. For a diffusion model, that means getting by with very few noise-removal steps.

Methods for each requirement exist. The central one is autoregressive diffusion distillation, in which a fast, chunk-by-chunk student model is trained to imitate a slower, stronger teacher. The authors describe these methods as scattered across separate pipelines. As a result, building an interactive world model still takes substantial effort at every stage, from data to inference.

How this paper puts it to work

The pipeline has two phases. First, the original bidirectional model is fine-tuned to follow camera motion, using videos whose camera paths are known. The camera information enters through PRoPE. This technique folds each frame's camera parameters into attention, the mechanism that decides how much each patch of an image influences every other. Those parameters cover both lens properties such as focal length (the intrinsics) and the camera's position and orientation (the extrinsics). When patches from two different frames interact, the model therefore takes account of how the camera moved between them.

The second phase turns this camera-aware model into a fast autoregressive one in three stages, following existing recipes called Causal Forcing and Causal Forcing++. Stage one retrains the model to work chunk by chunk using teacher forcing, in which it learns to predict each chunk from real earlier frames rather than from its own output. On its own, that leaves the model slow and prone to exposure bias: errors that accumulate once it must build on its own imperfect frames. Stage two creates a few-step student in one of two ways. It can learn shortcuts along the paths the slower model takes from noise to image. Alternatively, consistency distillation trains it to give matching answers from neighbouring noise levels, which avoids generating and storing those paths in advance. Stage three applies distribution matching distillation, which pushes the student's output towards what the stronger bidirectional teacher would produce. Throughout this stage the student generates entire videos from its own previous frames, as it will in use. Every model in every stage receives the camera signal.

The authors run the whole pipeline on two architecturally different video models. One is Wan2.1-T2V-1.3B, a text-to-video model; the other is HY1.5-TI2V-8B, which takes text and an image. Output clips are 77 frames at 480 by 832 pixels, produced with four denoising steps. They come four compressed, or latent, frames at a time. Rather than a single finished model, the authors say they release scripts, intermediate checkpoints for every stage, documentation and inference code. Others can then pick up or alter the pipeline at any point.

What it shows, and what it does not

The one measured result is the wait for the first frame. The tests ran on a single A800 graphics chip and left out time spent in the autoencoder that translates between the model's compressed frames and pixels. On that basis, the original HY1.5 took about 771 seconds to produce its first frame. The chunk-by-chunk version after stage one took about 81 seconds, and the final four-step model about 3.4 seconds, roughly 224 times faster than the original. For Wan2.1 the figures were about 269, 29 and 1.1 seconds, a speed-up of about 237 times. The authors note that part of the gap is structural. A bidirectional model cannot show any frame until it has finished the whole sequence, while an autoregressive one emits its first frame first. The figure captures the wait before playback, not the pace once it starts.

Everything else rests on example images. The authors show samples in which the distilled model follows camera paths and copes with changing camera instructions. However, they report no scores for picture quality or for how closely the camera obeys. Every comparison is between versions of the same base model rather than against other world models or distillation recipes. The clips are 77 frames long, so the work says little about whether a world stays coherent over a long session. Only camera control and two base models are supported; other controls and further models are left for future work.

The most practical material lies in the ablations, experiments that vary one ingredient at a time. Training on SpatialVid, real videos with camera poses estimated by software, did not give reliable control on either model, even after extra filtering. The authors suspect noisy poses and call this a limitation of their attempt, not proof that the dataset is unsuitable. What worked was data with effectively exact camera paths. One source was scenes from DL3DV, reconstructed in 3D and re-rendered along chosen trajectories. The other, used for the open release, was videos generated by another world model, WorldPlay. On HY1.5, control was absent after 1,000 to 2,000 training steps, emerged around 5,000 and was strong by 8,000. On Wan2.1, batches of fewer than four examples often failed, and 16 was enough. The step-count and batch-size findings each come from a single model, and all of these results are shown through example outputs rather than measured scores.

PAPER	*minWM: A Full-Stack Open-Source Framework for Real-Time Interactive Video World Models*
AUTHORS	Min Zhao, Hongzhou Zhu, Bokai Yan, Zihan Zhou, Yimin Chen, Wenqiang Sun, Kaiwen Zheng, Guande He, Xiao Yang, Chongxuan Li, Fan Bao, Jun Zhu et al.
VENUE	arXiv preprint — not yet peer reviewed · 28 May 2026
LINK	arxiv.org/abs/2605.30263
INDEX	stacks.braintank.ai
CODE	github.com/shengshu-ai/minWM