BraintankAI ASIA

EXPLAINER · WORLD ACTION MODELS · 5 MIN

SINGAPORE · National University of Singapore

How world action models differ from VLAs, video generators and world models

A 2026 survey from the National University of Singapore defines robot models that forecast what happens next and act on it, and sorts them along four axes.

12 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Tara Winstead

Ask a robot arm to push a cup across a table and there are two broad ways to choose its next move. It can map what its camera sees, plus the instruction, straight to a motor command. Or it can first picture how the scene should look a moment later and let that picture shape the command. Models of the second kind have come to be called world action models, or WAMs. Qiuhong Shen, Xinchao Wang and colleagues at the National University of Singapore argue that the label has spread without an agreed meaning. Papers using it come from video generation, robot learning and language modelling, and two of them can share the name while sharing almost nothing in how they are built.

Their survey, posted to arXiv in 2026 as a preprint not yet peer reviewed, aims to give the field a common definition and vocabulary. It trains no model and runs no benchmark of its own. It offers a boundary, two ways of sorting the literature and a list of open problems.

The idea

The survey starts from two older families. The first is the vision-language-action model, or VLA. This is a robot controller built on a model pretrained on images and text, which takes a camera view and an instruction and outputs an action. RT-2, OpenVLA and the π series belong here. A standard VLA never models how the world will change because of what it does. The authors argue that this limits how well it can reason about consequences before acting.

The second family is the world model, which predicts what the environment will look like after a given action, as in the Dreamer and V-JEPA lines of work. A world model forecasts, but nothing obliges it to turn that forecast into a decision. The authors add that generating video is one way of modelling the world, not the definition of it.

A WAM sits where these two meet. It counts only if a predicted future stays in the path to the action, helping to produce, choose or check it. There are three admissible routes. A cascade forecasts first and then derives the action, for instance with an inverse dynamics model, which looks at a current frame and a future frame and infers the movement linking them. Action scoring imagines the outcomes of candidate actions and picks among them. Joint prediction generates future and action together. A video generator that makes plausible clips is not a WAM, and neither is a VLA that forecasts nothing.

How this paper puts it to work

The first way of sorting is by philosophy, according to where along the process of generating video the action is extracted. Render-and-Decode methods produce a full future video in pixels and read the action off it. In the authors' census this line begins with UniPi in 2023 and runs through systems such as GR-1, GR-2 and Gen2Act. Latent-Only methods stop inside the video model before any pixels are drawn and use its internal state instead. One example is the intermediate features produced by diffusion, a technique that generates images by removing noise step by step. Video-Generation-Free methods skip video generation entirely and forecast in other forms, such as the tokens of a language model or abstract embeddings of the JEPA kind.

The second way of sorting is an anatomy with four axes. Substrate is the format the future is written in: pixels, abstract features, geometry such as motion flow or point clouds, or task-specific maps such as value or contact maps. Coupling is how action connects to the forecast. A model can simulate the outcomes of proposed actions, generate both together, or forecast first and attach an action decoder afterwards. Backbone is the underlying architecture. It spans five families, among them diffusion, autoregressive models, which build output piece by piece, and language and vision-language models. Deployment ranges from forecasting a whole task once and following it without re-planning (open loop) to executing a short chunk of actions and then re-planning from a fresh view (chunked closed loop), or re-planning at every step.

Each method in the census tables gets a coordinate on all four axes. The deployment axis carries a practical constraint: chunked control works only if the model computes the next chunk before the robot finishes the current one. Two examples show how this plays out. π0.7 qualifies as a WAM because a separate world model refreshes intermediate goal images from several viewpoints every four seconds, and these steer its action module. DreamZero is a 14-billion-parameter video model

that reaches closed-loop control seven times a second. It stays anchored to reality by swapping imagined observations for measured ones in its key-value cache, the store of earlier computations that spares a transformer from reprocessing old inputs.

What it shows, and what it does not

Laid out this way, the census shows patterns. Substrates stay within the four categories, and backbones cluster around diffusion and hybrids. The sharper differences lie in coupling and in runtime tricks. These include varying chunk length, calling the model only when an event demands it, and pairing slow and fast loops. The costliest combination renders full pixels, generates them jointly with actions by diffusion and re-plans at every step. Few methods can use it directly in production. Systems that seem to use it usually give ground on one axis, relying on simulation, offline rollouts or a smaller model at run time.

From this the authors draw their central observation: the field is moving towards generating less of the future while keeping what control requires. Fast-WAM and GigaWorld-Policy drop the rendered forecast at deployment without losing control. The authors read this as a sign that training on video may matter more than drawing it at test time. Video pretraining itself matters: removing it from a policy called VPP sharply cuts task completion. So does its source: VidMan found that robot footage helps while mismatched first-person video can hurt.

The limits are substantial. The survey runs no new experiments and re-tests no methods side by side, and its claim that rendered forecasts can be dropped rests on a few cited methods. It offers a reading of a trend rather than a measured result, and most findings it cites are summarised in words rather than numbers. The authors note that speed is often reported apart from accuracy, and on different hardware, which blocks fair comparison. They call for reports that put success, latency, memory and the length of task sustained on the same footing. A convincing forecast is also no reliable stand-in for a working robot: MotuBrain reports only a weak link between world-model scores and task success. Generalisation remains limited, with Act2Goal dropping steeply when conditions are randomised more aggressively. Open problems include keeping forecasts from drifting over long tasks and transferring skills between robot bodies. Another is making forecasts a robot can physically execute. A flow-based world model conditioned on a single frame cannot capture an object's velocity, so pushing tasks fail when the object keeps moving after contact ends. The authors concede that the boundaries between these open problems are analytical rather than absolute, since progress on one shifts the trade-offs of the others.

PAPER	*World Action Models: A Survey*
AUTHORS	Qiuhong Shen, Shihua Zhang, Yue Liao, Qi Li, Zhenxiong Tan, Shizun Wang, Shuicheng Yan, Xinchao Wang et al.
VENUE	arXiv preprint — not yet peer reviewed · 18 June 2026
LINK	arxiv.org/abs/2606.20781
INDEX	stacks.braintank.ai
CODE	github.com/world-action-models/awesome-world-action-models