

World Action Models: A Survey

Dream Less, Act More

Qihong Shen, Shihua Zhang, Yue Liao, Qi Li, Zhenxiong Tan, Shizun Wang,
Shuicheng Yan, Xinchao Wang[†]

National University of Singapore

World Action Models (WAMs) are embodied predictive-action models that make a forecast of the future available to action. Recent WAMs repurpose large video generation models, and a parallel line relies on language or vision-language backbones without a video-generation core. This rapid expansion has blurred the boundary among broad world models, video generation models, action-grounded video world models, Vision-Language-Action policies, and WAMs. This survey gives the field a common account. It first clarifies these boundaries, then organizes existing works through two complementary views. The first view asks what each method is required to generate, spanning rendered futures, latent futures, and video-generation-free action reasoning. The second view decomposes each method by predictive substrate, backbone, action coupling, and deployment regime. This anatomy supports a unified discussion of interactability, causality, persistence, physical plausibility, and generalization, followed by data, evaluation, and open challenges. Across these axes, a consistent design pattern emerges: WAMs are not simply video generators with action heads, but predictive-action methods whose design choices trade representational richness against compute, memory, latency, and action-label cost. The field is moving toward methods that generate less of the future while preserving what control requires. The survey homepage is available at <https://world-action-models.github.io/>.

Contents

1	Introduction	2
2	The Emergence of World Action Models	4
2.1	Vision-Language-Action Models: Policies from Vision-Language Pretraining	4
2.2	World Models: Predictive Dynamics Beyond Video	4
2.3	World Action Models: Making the Future Action-Facing	5
3	Three Design Philosophies of World Action Models	6
3.1	Render-and-Decode World Action Models	7
3.2	Latent-Only World Action Models	10
3.3	Video-Generation-Free World Action Models	11
3.4	Where the three philosophies meet the formal anatomy	12
4	What Makes a World Action Model	12
4.1	A Unified Notation for World Action Models	12
4.2	Predictive Substrate: Where WAMs Dream	14
4.2.1	Pixel-grounded substrates	14
4.2.2	Feature substrates	15
4.2.3	Geometric primitive substrates	16
4.2.4	Affordance map substrates	16
4.2.5	Substrate and coupling are distinct axes	16

Corresponding author: Xinchao Wang (xinchao@nus.edu.sg).

4.3	Action Coupling: How Action Enters and Leaves	17
4.3.1	Action-conditioned rollout	18
4.3.2	Joint generation	19
4.3.3	Post-prediction action head	20
4.3.4	Action representation and chunk size	21
4.4	Architectural Backbone: How the Prediction Is Produced	21
4.4.1	Iterative-denoising backbones (video diffusion)	21
4.4.2	Autoregressive backbones (next-frame, next-token)	23
4.4.3	Joint-embedding-predictive backbones	24
4.4.4	Hybrid backbones (generative head plus action head)	25
4.4.5	LLM- and VLM-backbone WAMs	26
4.4.6	The five families together	26
4.5	Deployment Regime: Interactive, Rollout, Open Loop, Closed Loop	27
4.5.1	Open-loop rollout	27
4.5.2	Chunked closed-loop control	27
4.5.3	Single-step closed-loop control	28
4.5.4	Interactive simulator operation	29
4.6	Putting the Four Ingredients Together	29
5	Core Properties of World Action Models	32
5.1	Interactability	32
5.2	Causality	33
5.3	Persistence	34
5.4	Physical Plausibility	34
5.5	Generalization	35
6	Data and Evaluation	36
6.1	Data sources	36
6.2	Evaluation	38
7	Open Challenges	39
7.1	Dream More or Act More?	39
7.2	What Data Should Each Stage Learn From?	39
7.3	Can Memory Keep Up?	40
7.4	How Can WAMs Generalize?	40
7.5	What Grounds Abstract Action?	41
7.6	When Is a Future Physical?	41
7.7	What Should Evaluation Report?	41
8	Conclusion	42

1 Introduction

The long-standing goal of embodied AI is to build agents that perceive, reason, and act in unstructured physical environments. Over the past two years the field has moved toward policies that do more than react to the current observation: they also anticipate how the world may change before choosing what to do. The trend began with Vision-Language-Action (VLA) models, which repurpose pretrained vision-language models to map an observation and an instruction directly to an action. Works such as RT-2 [12], OpenVLA [78], and π_0 [11] showed that the semantic knowledge accumulated during internet-scale pretraining can be grounded into robotic behavior, generalizing across instructions, objects, and embodiments with modest fine-tuning. A standard VLA, however, never models how the environment changes under its own intervention. It learns an observation-to-action mapping with no explicit account of physics, contact, or perspective change. This omission limits the policy’s ability to reason about consequences before acting.

A rapidly growing body of work now equips such a policy with an explicit predictive component drawn from world models. Large-scale pretrained video generation models serve as a major source of such components for predicting future dynamics, but they are not the only choice among existing works. Despite their variety, these works fall into three recognizable streams, distinguished by how far the model carries its prediction before it decodes an action. The first stream adopts a video generation backbone all the way to pixels and then decodes action from the rendered future, as in UniPi [29] and the video-plan policies that followed [30, 81]. The second keeps the same backbone but halts before pixel decoding, recovering action from intermediate latents, flow fields, or masks, as in VPP [62] and later latent-space foresight [132, 181]. The third forgoes the video generator entirely and predicts in a language-token, audio, or joint-embedding space, as in DUST [171], Audio-WM [194], and latent-dynamics policies over frozen vision features [112]. Despite these differences, the three streams share a contract that a plain VLA does not. Each makes a future substrate action-facing, either by training future and action prediction together or by using a separate action module that consumes the predicted future. We name this family **World Action Models (WAMs)**, following recent usage of this terminology, and make its definition precise in Section 4.

This label has spread faster than the understanding behind it. Works that call themselves World Action Models arrive from the video-generation, robot-learning, and language-model communities, and two of them can agree on the name while sharing almost no implementation details. Two others can avoid the name entirely and yet build the same thing. Underneath the disagreement lies a common interface question: what predicted future is retained for action, and where along the path is the action decoded. Each answer buys predictive richness at a price in compute, memory, and latency inside a control loop, which is why the strongest WAMs tend to dream less of the future while still acting on what they need.

This survey organizes the resulting works so that any new method can be placed and compared, and it does so from two complementary viewpoints. The first is a design-philosophy-level taxonomy, presented in Section 3, that categorizes works by where action is decoded along the inference path and sorts every WAM into exactly one of three families: Render-and-Decode, which carries generation to pixels, Latent-Only, which stops at intermediate representations, and Video-Generation-Free, which is built without the video generator. The second viewpoint, developed in Section 4, is a component-level anatomy. It places each WAM on four axes, the predictive substrate that represents the future available to action, the architectural backbone that produces the prediction, the action coupling that joins prediction to control, and the deployment regime the model is built for. A unified notation, made precise in that section, expresses every WAM as a 4-tuple over these axes. Across both viewpoints we cover the works in our census, listed in Tables 1 and 2, the first holding the video-generation-based families of Render-and-Decode and Latent-Only and the second the Video-Generation-Free family that drops the video backbone. The two viewpoints are complementary, since the philosophy fixes what a model is required to generate while the 4-tuple fixes how it is built.

Organizing the design space is only half of what this survey provides. The other half is a critical account of what these models must do once they are deployed in embodied settings. We examine five properties that embodiment demands, treated in Section 5. A WAM must be interactable, accepting control signals during generation rather than only at the start. It must be causal, never letting the future leak into the action executed now. It must be persistent, holding long-horizon predictions together as the robot acts and replans. It must be physically plausible, predicting futures that the embodiment can actually realize. And it must generalize, keeping the same predictive-action contract useful when tasks, objects, scenes, cameras, and embodiments change. For each property we examine how existing methods meet or miss it, and we track what their solutions cost in compute, memory, data, and the effort to evaluate them.

The remainder of this survey is organized as follows.

- **From World Models to World Action Models (Section 2).** The discussion separates VLAs, broad world models, video generation models, video world models, and WAMs, then defines the action-facing future contract that turns a world model into a WAM.
- **Three Design Philosophies of World Action Models (Section 3).** The philosophy-level taxonomy places every WAM considered here into Render-and-Decode, Latent-Only, or Video-Generation-Free according to what each method is required to generate.
- **What Makes a World Action Model (Section 4).** The formal anatomy introduces the unified notation and

the four axes: predictive substrate, architectural backbone, action coupling, and deployment regime. It then records each WAM as a 4-tuple in Tables 1 and 2.

- **Core Properties of World Action Models (Section 5).** The property discussion examines interactability, causality, persistence, physical plausibility, and generalization, together with how existing methods address each requirement.
- **Data and Evaluation (Section 6).** The section organizes the data sources that fuel WAM training, from teleoperation and portable human demonstrations to simulation and internet-scale egocentric video, alongside the evaluation practices the field is settling on.
- **Open Challenges (Section 7).** The final discussion identifies the open problems that the four-axis decomposition brings into view, grouped by where they arise in the WAM stack.

Throughout, one observation persists through the survey. Every design choice in a WAM carries a practical consequence for compute, memory, and latency inside a control loop. The interesting question is rarely how much a WAM costs on its own. It is how that cost trades against the properties embodiment imposes. Viewed over time, the field has moved toward generating less of the future while holding on to what control requires.

2 The Emergence of World Action Models

2.1 Vision-Language-Action Models: Policies from Vision-Language Pretraining

Consider an embodied agent that observes o , receives an optional language instruction l , and must choose an action a . Vision-Language-Action (VLA) models became influential because they showed that pretrained vision-language and language backbones could be grounded into this action channel. In its simplest form, a VLA learns

$$\mathcal{L}_{\text{VLA}}(\theta) = \mathbb{E}_{(o,l,a)}[-\log p_{\theta}(a | o, l)]. \quad (1)$$

RT-2 [12] established the web-to-robot transfer pattern by co-fine-tuning a VLM to output action tokens. OpenVLA [78] made the recipe open and scalable across robot data mixtures, while the π series [11, 71, 72] moved the action side toward continuous flow matching and broader embodiment coverage. These methods are the closest neighbors of WAMs because they already bind vision, language, and action in one action-producing policy.

The VLA objective also states the boundary. Equation 1 does not require the model to predict what will be observed after the action is executed. A direct VLA can inherit semantic and spatial priors from a pretrained VLM, yet still act from the present observation alone. It becomes a WAM only when a predicted future observation helps produce, choose, or check the action.

2.2 World Models: Predictive Dynamics Beyond Video

A world model asks the complementary question. Rather than predicting the action directly, it predicts what future observation o' should follow from the current observation and an intervention. A recent robot-learning survey frames world models as predictive representations of how environments evolve under actions, with uses in policy learning, planning, simulation, evaluation, and data generation [61]. For the present survey, the minimal form is

$$\mathcal{L}_{\text{WM}}(\theta) = \mathbb{E}_{(o,l,a,o')}[-\log p_{\theta}(o' | o, a, l)], \quad (2)$$

where o' is written as an observation for clarity. The symbol is deliberately broad in this section. Depending on the method, it may be rendered pixels, a hidden feature, a geometric state, an affordance map, an audio cue, a symbolic state, or a token-level description. Section 4 later separates these cases. In latent model-based reinforcement learning, PlaNet [55], DreamerV3 [57], TransDreamer [18], and Dreamer 4 [58] instantiate this idea with compact dynamics states used for imagination and planning. V-JEPA [7] and V-JEPA 2 [4] show the same predictive idea in feature space, while iVideoGPT [174], RoboDreamer [210], EnerVerse [66], and InteractiveWorldSimulator [166] express it through token or video prediction.

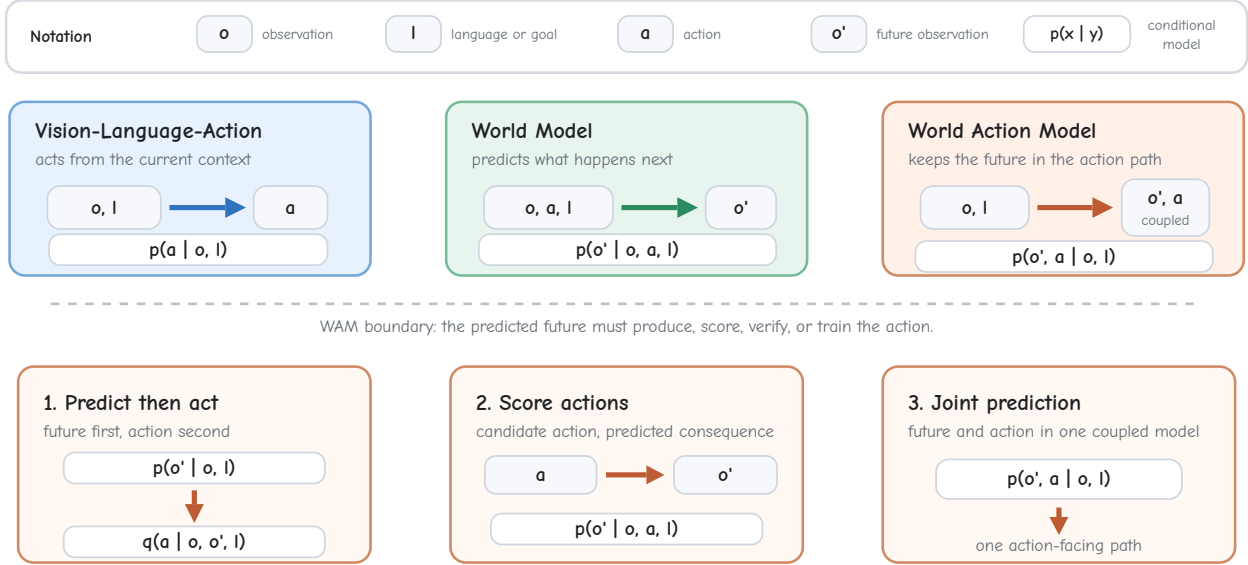


Figure 1 Definition of a World Action Model. A direct VLA predicts action from the present context, and a world model predicts a future observation. A WAM requires that future to stay in the action path, either through predict-then-act cascades, action-scoring rollouts, or joint future-action prediction.

Video generation is one visible instance of this broader idea, not its definition. A video generation model learns

$$\mathcal{L}_{\text{VGM}}(\theta) = \mathbb{E}_{(r, o')}[-\log p_{\theta}(o' | r)], \quad (3)$$

where r is a prompt, image, observation, or other conditioning signal. When the conditioning includes the agent’s action and the output is a future visual observation, the same family becomes a video world model:

$$\mathcal{L}_{\text{VWM}}(\theta) = \mathbb{E}_{(o, l, a, o')}[-\log p_{\theta}(o' | o, a, l)]. \quad (4)$$

Internet-scale video generators such as Wan, CogVideoX, Cosmos, NOVA, Sora, Latte, AnimateDiff, and VideoPoet [26, 53, 82, 114, 128, 130, 131, 156, 185] made this branch newly influential because they supplied transferable visual dynamics priors. The rise of WAMs follows from this broader world-model success. Video world models gave the field a high-capacity source of predictive structure, while feature, geometric, audio, symbolic, and token predictors showed that the useful future need not always be rendered video.

2.3 World Action Models: Making the Future Action-Facing

A World Action Model begins when the predicted future observation becomes part of how the action is obtained. The historically early form is a cascade:

$$p_{\Theta}(o', a | o, l) = p_{\theta}(o' | o, l) q_{\psi}(a | o, o', l). \quad (5)$$

The action module q_{ψ} may be an inverse-dynamics model, a pose tracker, a trajectory optimizer, a planner, or a separately trained policy. Presenting this cascaded form first is historically defensible. Early WAMs such as UniPi [29], VLP [30], and AVDC [81] generated or scored future visual trajectories before recovering executable actions with a separate module. Later cascaded WAMs keep the same logic while changing what o' looks like, as in video-to-trajectory, flow-to-policy, trace-to-control, point-track-to-policy, or future-audio-to-policy pipelines [28, 70, 73, 86, 87, 104, 108, 133, 194]. These works remain WAMs because the predicted future is used to obtain the action, even when the predictor and actor are trained separately.

A second form proposes an action first and then predicts its consequence:

$$p_{\Theta}(o', a | o, l) = q_{\psi}(a | o, l) p_{\theta}(o' | o, a, l). \quad (6)$$

This factorization covers candidate-action scoring and planning, where a proposed action is evaluated through its predicted consequence. It is still a WAM when the predicted consequence decides which action is executed [2, 33, 68, 142, 152].

The same idea can also live inside one trained model:

$$\mathcal{L}_{\text{WAM}}(\theta) = \mathbb{E}_{(o,l,o',a)}[-\log p_{\theta}(o', a | o, l)]. \quad (7)$$

Here one backbone or tightly coupled set of experts predicts the future observation and the action together. GR-1 and GR-2 [16, 172] helped establish this shared video-action direction with autoregressive image and action tokens, while PAD [52], UWM [213], WorldVLA [14], DreamZero [188], AIM [32], and newer joint designs [92, 141, 209] instantiate the same idea through diffusion, autoregression, or hybrid backbones. Feature-tap methods such as Fast-WAM [193] show that the future branch need not be fully rendered at inference as long as the action path still uses a representation shaped by future prediction.

The WAM definition is therefore direct and is summarized in Figure 1. A VLA models $p(a | o, l)$ without needing an explicit future. A world model models $p(o' | o, a, l)$ or $p(o' | o, l)$ without necessarily choosing the action. A WAM links the two: its predicted future observation helps produce the action, score the action, or train the action path inside one model. This definition uses o' as a Section 2 notational simplification. Section 4 instantiates that simplification as different predictive substrates and expands the equations into horizon-level notation. A direct VLA with an auxiliary future loss, a simulator used only as an RL environment, or a future head discarded before action use does not satisfy the definition.

3 Three Design Philosophies of World Action Models

Existing world action models share the WAM definition introduced in Section 2. They differ, however, in where the future predictor meets the action module. Some methods train a shared video-action backbone, while others keep a world model and a separate policy, tracker, inverse-dynamics module, or optimizer in a cascade. The training pipeline alone therefore does not expose the main design split. The decisive distinction is visible either in the inference forward pass or in how action supervision is wired during training: action prediction may pass through a rendered pixel future, stop inside a video backbone before rendering, or avoid the video-generation path altogether.

The resulting taxonomy has three mutually exclusive philosophies. **Render-and-Decode** runs the video-generation backbone all the way to pixel output before action is obtained. The action module may be trained jointly with the generator, trained separately, or supplied by tracking and trajectory optimization. **Latent-Only** keeps the video-world-model lineage but stops the inference path before pixel decoding. In many cases this means a video-generation trunk whose VAE decoder is bypassed, while in others it means a video-trained or video-derived latent predictor whose action-facing future has no decoded-pixel stage. Its action signal comes from intermediate latents, partially denoised features, flow fields, semantic masks, or value maps. **Video-Generation-Free** removes the video-generation backbone from the predictive path. The predictive component instead produces a future representation in the embedding or token space of a large language model, a vision-language model, a joint-embedding-predictive encoder, a deterministic regressor over a frozen vision foundation model, or a non-video diffusion or hybrid backbone over a compact non-pixel substrate. The taxonomy therefore tracks the requirement imposed during inference or action-supervision, rather than the training data, backbone family, or action interface alone.

This split is separable from the cascaded-versus-joint distinction introduced in Section 2.3. Cascaded-versus-joint asks how prediction and action are arranged inside the model. The philosophy-level taxonomy asks where the action prediction is grounded along the inference path: at the rendered pixel output for Render-and-Decode, at an intermediate latent or feature for Latent-Only, or outside the video-generation path for Video-Generation-Free. A Latent-Only WAM can be either cascaded or joint, and the same holds for a Render-and-Decode WAM. Figure 2 sketches the inference data flow under each philosophy.

Figure 3 adds the chronological dimension that the static taxonomy omits. Early WAMs largely adopt the direct path from video generation to control: generate a visual future, then decode action from it. Later Latent-Only methods preserve a video-derived prior but move the action signal earlier in the inference path,

Three Design Philosophies of World Action Models

What does the model commit to generating?

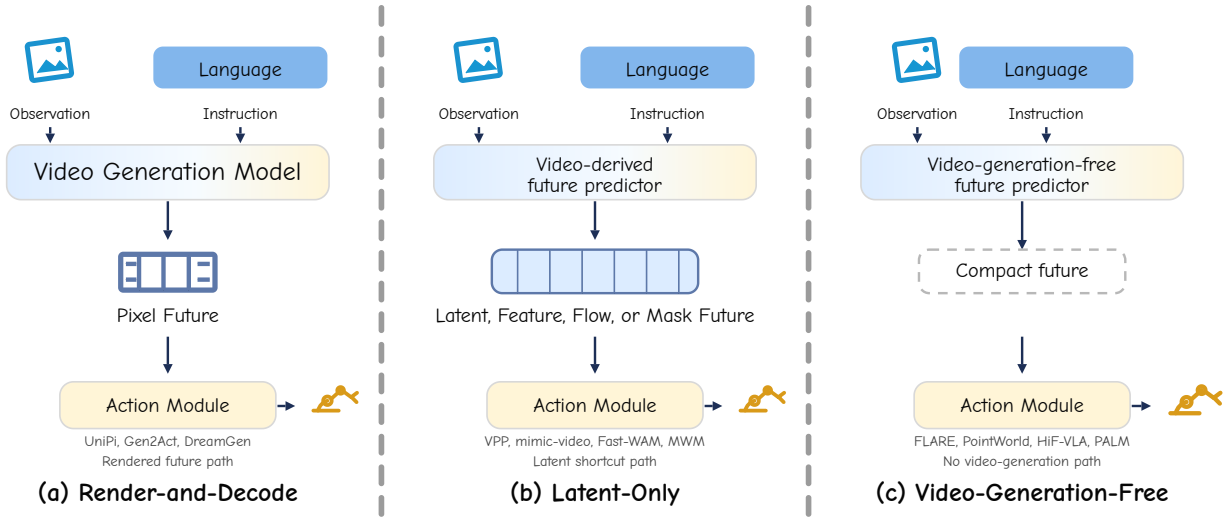


Figure 2 Three design philosophies of World Action Models. The columns separate the last future representation required before action is decoded: a rendered pixel future, an intermediate video-derived latent or feature, or a non-video-generation representation. The categorization is exhaustive over the WAM census and separable from the action-coupling and backbone choices treated in Section 4.

before pixel decoding. Video-Generation-Free methods appear as a more recent alternative, replacing the video generator with predictive supervision in the embedding or token space of LLM, VLM, JEPA, or non-video diffusion or hybrid backbones. Over the same period, the application domain has also broadened. Recent entries now cover tabletop manipulation, dexterous hand control, contact-rich tactile interaction, autonomous driving, and aerial manipulation. They instantiate the same three philosophies under substantially different latency requirements and observation regimes. With this chronological view in place, the remainder of this section explains how each philosophy instantiates this design choice across the WAM census. Section 4 then gives the component-level anatomy behind these choices through substrate, backbone family, action coupling, and deployment regime.

3.1 Render-and-Decode World Action Models

Render-and-Decode is the most direct path from video generation to action. Its premise is that a rendered future is worth producing inside the inference-time control loop because it preserves the full visual prior learned by the video backbone. By running that backbone all the way to pixel output, these methods expose appearance, motion, contact, and scene dynamics to the actor in an inspectable form. The inference pattern is simple: condition a video generator on observation history and instruction, produce a future video, and decode the next action from that rendered future. The output may be raw RGB, a multi-view or 4D RGB-D stream, or a sequence of visual tokens that decode back to images, while the action decoder may be inverse dynamics, 6-DoF tracking, dense correspondences, or a dedicated action head. This premise gives Render-and-Decode its conceptual clarity, but also makes visual synthesis part of the policy’s latency budget.

The founding branch treats rendered video itself as the plan. UniPi [29] establishes the template with text-conditioned video followed by inverse dynamics, VLP [30] adds search over language-level plans, and AVDC [81] shows that dense correspondences can replace action labels. The same interface then moves from simulated or task video to demonstrations with different sources of embodiment evidence. Dreamitate [99] tracks generated stereo video, Gen2Act [9] conditions policy learning on generated human motion, This&That [157] adds gesture coordinates, and GR-MG [89] turns a progress-aware goal image into a manipulation condition. Later variants refine how the rendered plan becomes executable. CreFlow [125] post-trains the video-as-policy generator with violation masks from a temporal-logic monitor, while MoLA [96] maps a frozen SVD rollout

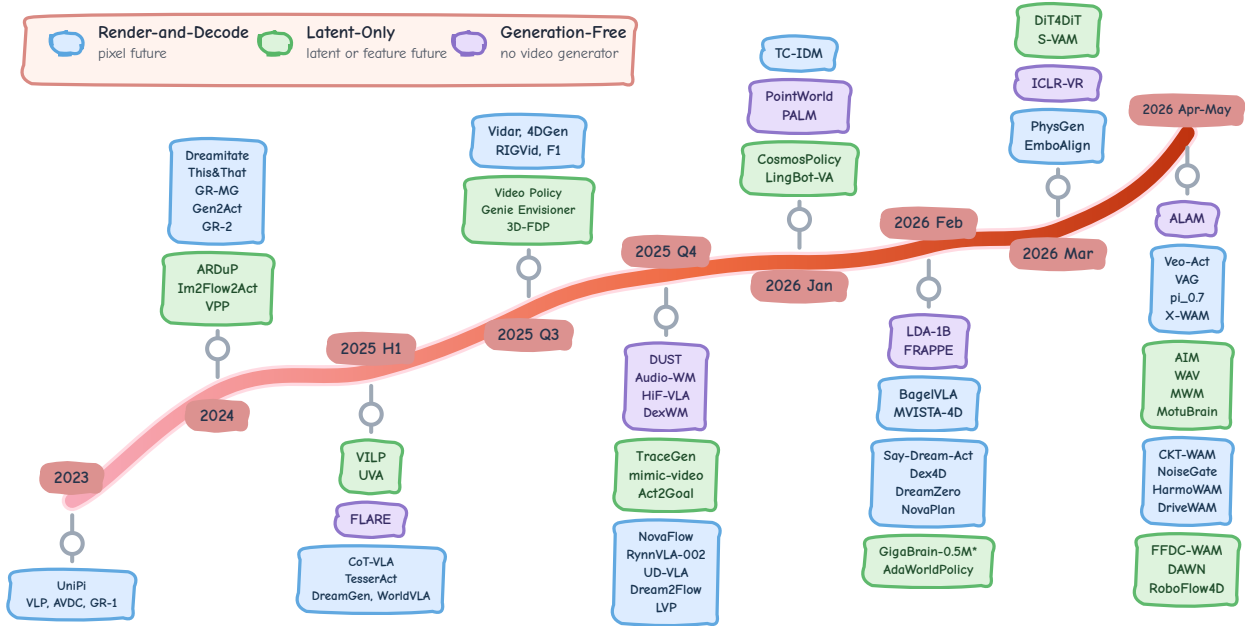


Figure 3 Chronological stream of representative WAMs grouped by design philosophy. The timeline uses coarse bins for 2023, 2024, and 2025 H1, then finer bins for the denser late-2025 and early-2026 period through May. Each date marker attaches a compact vertical stack to the curve, alternating above and below the stream. Dense periods retain early or central verified WAMs so labels remain legible. Stacks above the curve place the widest block nearest the stream, while stacks below the curve begin with the narrowest block. Render-and-Decode appears earliest, Latent-Only emerges as the field starts to drop pixel decoding from the control path, and Video-Generation-Free entries appear as LLM, VLM, JEPA, and non-video diffusion or hybrid backbones begin to carry the predictive component without a video-generation core.

into semantic, depth, and flow latent-action codebooks before action decoding. The shared lesson is that the rendered future can be a common planning currency even when the recovery module changes.

Modular Render-and-Decode WAMs make that separation explicit. DreamGen [73] uses generated robot futures to obtain pseudo-actions before downstream policy training, so the coupling is offline rather than inference-time imagination. 4DGen [108], RIGVid [133], TC-IDM [119], and GraspDreamer [147] keep the world model and action recovery separate through pose tracking, VLM filtering, tool-centric point-cloud conversion, or human-grasp retargeting. VERA [94] makes the split even sharper by leaving the video planner action-free and training a Jacobian inverse-dynamics translator. These works are useful because the world model can be reused across embodiments, but the downstream module must absorb the gap between visual plausibility and robot executability.

A related modular pattern inserts geometric or search structure between the generated future and the action. NovaFlow [87] and Dream2Flow [28] distill generated videos into 3D object flow, while EmboAlign [195] filters rollouts with compositional constraints before trajectory optimization. NovaPlan [38] places generated video inside a closed-loop manipulation planner, 3D-ALP [142] uses a 3D world model as a rollout oracle during test-time search, and AeroPlace-Flow [120] grounds a task-complete image into object flow for aerial placement. The pixel future remains the first prediction, but control increasingly depends on the structured object recovered from it.

Joint Render-and-Decode WAMs move action prediction inside the video backbone rather than attaching a separate decoder afterwards. GR-1 [172] is the anchor for this line, because future image tokens and action tokens are learned in one autoregressive stream. GR-2 [16] scales the same video-language-action recipe to a larger pretraining and robot-trajectory mixture. WorldVLA [14] adds masking to reduce action leakage while still predicting future VQ images, RynnVLA-002 [13] refines actions through future image-state prediction, PhysGen [144] adapts a continuous autoregressive video backbone, and DriveWAM [141] interleaves driving

video latents with ego-action tokens on a shared DiT. This group changes the coupling pattern, but the action token is still grounded by a pixel-decodable future.

Diffusion and hybrid variants make the same joint target less sequential. PAD [52] jointly denoises future images and actions in one DiT. UD-VLA [19], CoVAR [183], and VideoVLA [140] vary the bridge through discrete diffusion, cross-stream attention, or video-generator co-training. Motus [10] is placed here by its joint-prediction path, where sparse future video and actions are denoised together. Its scheduler can switch to action-only or inverse-dynamics use cases, but those shortcuts do not change the category assignment. These works keep the action and the future in the same generative state, but make the scheduling less rigid than autoregression.

Other joint variants expand what the generated state contains. ActionImages [203] encodes actions as multi-view RGB heatmaps so the observation and action videos share a generated channel grid. DriveVA [106] applies flow matching to driving video latents and action tokens, HarMoWAM [36] sends rendered frames and latent features to different action experts, and NoiseGate [67] learns per-latent denoising schedules around a frozen joint backbone. The family therefore grows through stronger coupling, not through removing the rendering obligation.

Other Render-and-Decode methods change what is rendered so that the future carries more task structure than RGB alone. TesserAct [202] generates RGB-D-Normal 4D video, MVISTA-4D [160] infers arbitrary-view RGB-D through test-time latent optimization, and X-WAM [51] combines multi-view RGB-D prediction with explicit 3D reconstruction. Dex4D [83] lifts generated frames into 3D object point tracks for dexterous control. DriveDreamer-Policy [212] threads depth, video, and action through modular experts, while MV-VDP [90] diffuses future multi-view RGB together with end-effector heatmap videos. These methods keep the visual future, but they make geometry or affordance structure easier for the action side to extract.

A parallel sparse-future line keeps the output pixel-grounded while reducing how much video must be produced. RoboEnvision [184] uses instruction-aligned keyframes as long-horizon plans. CoT-VLA [201] is the influential VLA-side example, because it generates a visual chain-of-thought before action prediction rather than relying on hidden language reasoning alone. pi0.7 [72] follows the same pixel-grounded logic: a BAGEL-class generator refreshes multi-view subgoal images asynchronously, and a PaliGemma action expert consumes the latest available goals. BagelVLA [63] interleaves subtask language, future keyframes, and action chunks, while SWEET [143] feeds sparse edited keyframes to a goal-conditioned diffusion policy. Here the design aim is not to avoid pixels, but to render only the visual milestones that the actor can use.

Foundation-scale Render-and-Decode WAMs focus on whether a large video prior can be made responsive enough for closed-loop control. DreamZero [188] is the anchor: it turns a Wan-class autoregressive video diffusion backbone into an online video-action policy, uses chunked prediction with observation replacement in the KV cache, and shows that a large rendered-future WAM can stay inside a real control loop. Say-Dream-Act [48] follows the same practicality question from the distillation side, compressing Cosmos-style imagination into a few-step generator before action fusion.

Other foundation video planners keep the rendered future as a planning object and differ mainly in how they ground it. Vidar [37], LVP [17], and Veo-Act [199] use masked inverse dynamics, Wan-style image-to-video generation, or a precision VLA executor. VAG [85] and F1 [111] attach action experts to generated or pooled foresight, while EVA [162] aligns the video planner with an inverse-dynamics reward so the rendered future stays executable.

The newest foundation-scale variants spend the video prior more selectively. CKT-WAM [75] transfers context between heterogeneous WAMs, DreamAvoid [33] invokes video imagination only on critical phases, and Pelican-Unify 1.0 [198] links VLM reasoning to a Wan-style future generator. This group keeps the rendered future available when it carries task information, while shifting routine action selection toward cheaper context transfer, critical-phase triggering, or grounded VLM guidance.

The defining limitation of Render-and-Decode is the price of producing pixels. Each prediction step pays for the full denoising or autoregression schedule, even though the actor rarely consumes the rendered output in its entirety. Moreover, visual-quality metrics inherited from video generation only weakly predict downstream task success. This practical tension links the category closely to the open-loop and chunked closed-loop regimes unpacked in Section 4.5.

3.2 Latent-Only World Action Models

The Latent-Only philosophy keeps the video-world-model prior of Render-and-Decode but removes pixel decoding from the inference-time control path. Its motivation is to retain temporal and physical structure learned from video while reducing inference cost. The action signal is decoded directly from a latent, an intermediate denoising feature, a flow field, a semantic mask, a value map, or another video-derived feature future, each of which is cheaper to produce than a fully rendered video. The trade-off is that these methods give up the pixel-space supervision and visual-quality metrics that make Render-and-Decode easy to inspect. The designs differ mainly in where they intercept the predictive path: some consume the model’s own video latent, some tap an intermediate denoising state, some predict a structured substrate such as object flow or masks, and some inherit a joint video-action backbone but disable pixel decoding at inference. In all cases, the key move is to keep a video-shaped dynamic prior while refusing to pay for the final renderer during action selection.

The earliest Latent-Only entries make the interception point explicit. ARDuP [65] focuses latent video diffusion on active regions, VPP [62] conditions inverse dynamics on Stable Video Diffusion representations, and VILP [178] recovers actions from time-aligned multi-view latent video. UVA [93] and Video Policy [100] show the two main training patterns: either share a video-action latent across heads, or freeze the video generator and train the policy on its intermediate features. Genie Envisioner [102] is the clearest bridge between these patterns. GE-Base remains an instruction-conditioned multi-view video DiT, while GE-Act decodes action chunks from a one-step denoised latent cache through flow matching. Act2Goal [208] reuses the same foundation for goal-conditioned transition features, and WAV [91] performs value-guided latent trajectory inference before action decoding. These methods show that the first shortcut is to stop at the representation that action already needs.

As video diffusion backbones scale, Latent-Only designs increasingly intercept the denoising trajectory rather than a finished latent. mimic-video [132] attaches an action decoder to an intermediate ODE checkpoint of Cosmos-Predict2, and DiT4DiT [113] conditions its action expert on intermediate Cosmos-Predict2.5 features in a deterministic step. S-VAM [181] self-distills multi-step denoising into one pass over geometric and semantic features. Fast-WAM [193] and GigaWorld-Policy [186] keep future-video co-training but make full generation optional or unnecessary at inference. LaMP [164] exposes a one-step 3D motion hidden state to a VLA through gated cross-attention, and VAMPO [42] tunes early-step VPM latents against expert visual dynamics. The practical shift is from pixel synthesis to feature extraction.

Another branch predicts structured motion before action decoding. Im2Flow2Act [177] and 3DFlowAction [207] replace appearance with object flow, while 3D-FDP [126] predicts dense 3D flow over query points before a second action diffusion model runs. TraceGen [86] predicts 3D traces that inverse kinematics converts into joint commands. This flow-and-trace group compresses the future toward motion rather than texture.

The structured-substrate branch then broadens beyond flow. 3PoinTr [70] and RoboFlow4D [104] freeze future point-track or 4D-flow predictors and train downstream policies to consume the predicted motion. MWM [109] keeps a video-diffusion architecture but predicts semantic-mask evolution, OmniVTA [206] adds tactile latents for contact-rich control, and EgoExo-WM [152] uses a DINOv3 latent regressor under model-predictive control after exo-to-ego data conversion. This branch makes the future closer to contact and task state than a rendered frame would be.

The clearest Latent-Only inference pattern appears when a joint video-action backbone is trained with visual prediction but used through a latent action path. UWM [213] is the reference case: independent diffusion timesteps let the method collapse the visual branch at inference while retaining the video-trained representation. CosmosPolicy [79], AdaWorldPolicy [191], and MotuBrain [121] follow the pixel-latent version of the same idea, where future RGB or video latents remain useful during training but the action call at inference does not require a fully rendered RGB forecast. LingBot-VA [88] exposes semantic structures before pixel reconstruction, and AIM [32] decodes action through a spatial value map rather than directly through future RGB. The design consequence is clear: visual prediction shapes the representation, but action selection does not have to render the visual prediction.

Recent Latent-Only WAMs extend the same idea across modalities and memory designs. GigaBrain-0.5M* [158] feeds the policy with future-state latents and value embeddings from a Wan2.2 world model. VTAM [192]

treats camera views and GelSight tactile input as decoder-bound latent views, while CLWM [27] forecasts future DINOv3 feature maps and replaces a growing KV cache with Dual-State Test-Time Training memory.

The same inference logic also travels across domains and joint-state designs. DAWN [110] iteratively refines V-JEPA latent world tokens against DiT action hypotheses, JOPAT [49] adds point tracks and visibility to the joint denoising state, and tau0-WM [209] keeps Wan-style video latents in the joint state but decodes action without routine pixel rendering. WALL-WM [92] pushes the same idea into depth-wise cross-attention between matched video blocks and action queries. These examples make Latent-Only a design philosophy rather than a single substrate choice.

A newer branch uses latent imagination as an adaptive control signal. FFDC-WAM [161] caches Motus latent video tokens with the proposed action chunk and uses a causal attention verifier to decide when a full WAM replan is needed. tau0-WM [209] can also reactivate its action-conditioned video simulator only when candidate scoring is useful. This turns imagination into a learned controller decision: the policy invokes the heavier future model only when confidence drops.

Across this family, the central tension is clear. Video pretraining supplies a valuable dynamic prior, but full pixel rendering imposes latency and memory costs that are difficult to justify for every control step. Latent-Only methods preserve the prior while bypassing the renderer, enabling several members of the family to reach real-time inference at the cost of reduced direct visual interpretability.

3.3 Video-Generation-Free World Action Models

The Video-Generation-Free philosophy removes the video-generation backbone altogether. Here *video-generation-free* refers strictly to the absence of a pixel-level video backbone in the predictive path. The LLM, VLM, JEPA, and flow-matching components below remain generative models in the broader sense. Its motivation is to avoid the training and inference cost of video generation while retaining predictive supervision in a more compact representational space. These methods instead leverage large language models, vision-language models, joint-embedding-predictive encoders, deterministic regressors over frozen vision foundation models, or non-video diffusion and hybrid backbones over a compact non-pixel substrate. The dominant pattern attaches a lightweight action expert to a VLM, learns a latent-action vocabulary that bridges actionless video and robot control, or supervises a policy with an auxiliary feature-prediction loss in a vision-foundation embedding space. The family is listed in Table 2, where it complements the Render-and-Decode and Latent-Only entries in Table 1.

The feature-prediction line replaces pixel prediction with teacher-target or foundation-feature prediction. FLARE [204] is the anchor: additional policy tokens are aligned to future-observation embeddings from a frozen teacher, and the action expert consumes those predicted future tokens at inference. FRAPPE [200] extends the idea by aligning a generalist policy to future representations from multiple vision foundation models, while DexWM [44] applies deterministic future-feature regression to dexterous hand-object interaction. LDA-1B [112] moves the same principle into a Qwen3-VL stack by jointly learning dynamics, policy, and forecasting in a structured DINO latent. These works keep the predictive training signal, but place it in an embedding space whose quality is judged through action rather than photorealism.

Another Video-Generation-Free line treats the future as a learned token or latent transition. DUST [171] augments an Eagle-2 VLM with a dual-stream diffusion head that predicts future observation embedding tokens and an action together without a video-generation foundation. ALAM [149] learns latent transitions from action-free video and regularizes them with composition and reversal consistency before a flow-matching action expert consumes the codes. The shared move is to keep a future object in the inference path while choosing a representation that is cheaper than video.

Structured non-pixel futures give the family its second anchor. PointWorld [68] predicts action-conditioned 3D point flow through a deterministic Point Transformer and plans with MPPI over the compact future. PALM [107] structures foresight as affordance maps plus progress, while HiF-VLA [103] uses compact motion-vector foresight tokens and ICLR-VR [124] predicts a gripper-keypoint polyline before each action chunk. These works replace video with geometry, affordance, or trace-like variables that are closer to the control interface.

The remaining generation-free variants broaden the same principle across modalities and feedback. Audio-WM [194] predicts future Mel latents that condition a robot policy when contact sound carries task evidence. DDP [64] switches to latent imagination when a real-imagination gap indicates an out-of-distribution event, and Feedback-WM [2] treats a learned latent transition model as an observer that corrects diffusion-policy denoising online. Together, these methods show what is gained by dropping the video generator entirely, and what must be replaced by tokens, features, audio latents, affordances, or geometry.

3.4 Where the three philosophies meet the formal anatomy

The three philosophies answer the question of *where* action prediction sits along the video-generation inference path: at the rendered pixel output for Render-and-Decode, at an intermediate latent or feature for Latent-Only, or outside the video-generation path for Video-Generation-Free. They do not yet answer *how* the model produces that prediction, where action is injected, or when the model is invoked relative to the control loop. Tables 1 and 2 place every WAM in the same census under a four-axis design space that captures these component-level choices. Section 4 introduces the formal notation behind that four-axis view and uses it to compare WAMs across substrate, backbone family, action coupling, and deployment regime. Together, the philosophy-level taxonomy of this section and the component-level anatomy of the next provide two complementary viewpoints for navigating the field.

4 What Makes a World Action Model

The term *world action model* now spans works built in different communities, with different vocabularies, and on top of different backbones. A paper-by-paper enumeration is insufficient for comparing them. Methods that share the WAM label can differ in nearly every implementation detail, whereas closely related methods may use entirely different terminology. This section takes a different route. It treats WAMs as instances of a common mathematical object: a conditional joint distribution over future predictions and future actions. Under this view, the main design choices specify what future is represented, how actions are coupled to that future, what model family produces the prediction, and how the resulting model is invoked during control. We refer to these choices as the *predictive substrate*, *action coupling*, *architectural backbone*, and *deployment regime*, respectively. Figure 4 previews how the four ingredients combine. Section 4.1 introduces the shared notation. Sections 4.2 through 4.5 then walk through the four ingredients one by one. They sort existing methods along the choices each ingredient offers and note the practical trade-off that each choice carries.

4.1 A Unified Notation for World Action Models

At the right level of generality, WAMs can be described by a common abstract object. The object is a parameterized conditional joint distribution over a window of future predictions and a window of future actions:

$$p_{\Theta}(s_{t+1:t+H}, a_{t:t+H-1} \mid o_{\leq t}, a_{<t}, l). \quad (8)$$

Here $o_{\leq t} = (o_0, \dots, o_t)$ is the observation history, with each $o_i \in \mathcal{O}$ an observation, such as a frame, video chunk, or multimodal sensory packet, $a_{<t} = (a_0, \dots, a_{t-1})$ is the past action history with $a_i \in \mathcal{A}$, and l is a task instruction in language, a goal image, or another auxiliary modality. The model predicts a future trajectory $s_{t+1:t+H}$ of length H in a chosen *substrate space* \mathcal{S} together with a future action chunk $a_{t:t+H-1}$ of the same length. The parameter collection Θ may be one shared backbone or a retained world-model module paired with an action module. To keep later equations compact, we abbreviate the full conditioning context as

$$c \equiv (o_{\leq t}, a_{<t}, l), \quad (9)$$

and write Equation 8 as $p_{\Theta}(s_{t+1:t+H}, a_{t:t+H-1} \mid c)$ whenever no ambiguity arises.

The substrate space \mathcal{S} is the space of the future representation that the WAM exposes to action production or action evaluation. We write it as the image of an interface encoder $\phi : \mathcal{O} \rightarrow \mathcal{S}$, but ϕ should not be taken as the generator’s internal implementation. A pixel-grounded WAM exposes decoded future frames, or the VAE / VQ latent grid that the generator decodes to a clip, even if the action side never invokes the pixel decode at inference. A feature WAM exposes a learned hidden state or token block with no decodable visual

Anatomy of a World Action Model

A WAM is pinned down by four choices that act together at deployment.

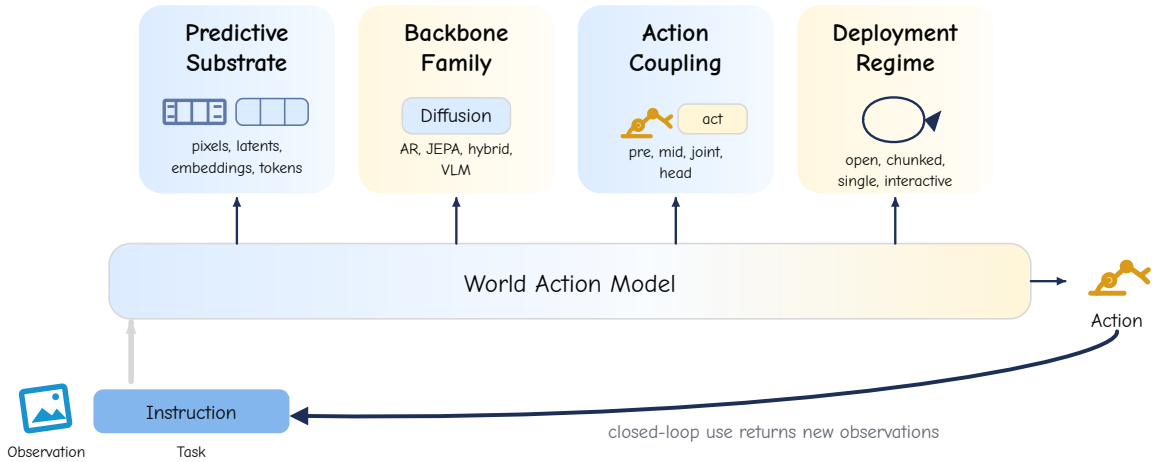


Figure 4 The anatomy of a world action model. Any existing WAM can be specified by four separable but interacting choices: what it predicts (Section 4.2), how action is coupled to that prediction (Section 4.3), the function family that produces the prediction (Section 4.4), and the deployment regime it is intended for (Section 4.5). Two methods with the same labels but different choices on any one axis behave very differently in practice.

interpretation, whether it is a frozen self-supervised teacher’s target, a feature tap of a generative trunk, or a VLM future-token block. A geometric-primitive WAM exposes a structured object whose dimensions are physical coordinates, such as flow, point clouds, depth, pose, or motion vectors. An affordance-map WAM exposes a task-relevant label or score map, such as a value, affordance, or end-effector heatmap. The symbol p_{Θ} in Equation 8 is a placeholder for an abstract joint, not necessarily one monolithic network. Section 4.4 lists the five concrete families that realize its future-prediction factor in practice, each with its own parameterization. Until that point, p_{Θ} should be treated as the action-facing future contract that every WAM must satisfy.

The four ingredients of this section are choices about how Equation 8 is realized in practice.

(i) Predictive substrate. A choice of action-facing future representation, and therefore of the space \mathcal{S} in which $s_{t+1:t+H}$ lives. Section 4.2.

(ii) Action coupling. A choice of factorization of the joint in Equation 8. Three top-level factorization families appear across existing WAMs and return in Section 4.3:

$$\text{action-conditioned rollout} : \begin{cases} q_{\psi}(a_{t:t+H-1} | c) p_{\theta}(s_{t+1:t+H} | c, a_{t:t+H-1}), & \text{chunk-level,} \\ \prod_{k=0}^{H-1} q_{\psi}(a_{t+k} | h_k, c) p_{\theta}(s_{t+k+1} | h_k, a_{t+k}, c), & \text{step-wise,} \end{cases} \quad (10)$$

$$\text{joint generation} : p_{\theta}(s_{t+1:t+H}, a_{t:t+H-1} | c), \quad (11)$$

$$\text{post-prediction head} : p_{\theta}(s_{t+1:t+H} | c) q_{\psi}(a_{t:t+H-1} | s_{t+1:t+H}, c). \quad (12)$$

where $h_k = (s_{t+1:t+k}, a_{t:t+k-1})$ in the step-wise submode. Action-conditioned rollout composes an action source with a world predictor, either by fixing an action chunk before prediction or by updating the action after each predicted step. Joint generation produces the future substrate and the future action chunk from one coupled generative model $p_{\theta}(s_{t+1:t+H}, a_{t:t+H-1} | c)$, where θ may include a shared trunk as well as modality-specific heads or experts. The post-prediction head factorization predicts the substrate first and then decodes action from it with a smaller actor q_{ψ} , which is often a tracker, inverse-dynamics model, optimizer, planner, or policy that can be trained separately from θ . All three families carry the past action history through the abbreviation c , so the conditioning sides remain consistent with Equation 8.

(iii) Architectural backbone. A choice of function family that realizes the future-prediction factor of p_{Θ} , namely an iterative denoising network, an autoregressive decoder, a joint-embedding predictor, a hybrid that combines two of the above, or a large language or vision-language model with an attached action decoder. Each family carries its own parameterization, which we introduce only when the family is treated. Section 4.4.

(iv) Deployment regime. A choice of when the WAM is invoked and over what window. Open-loop rollout invokes the WAM once with $H \approx T$ and specifies the whole window before execution. Chunked closed-loop invokes the WAM every K control steps with $H = K$ and replans after each chunk is executed. Single-step closed-loop invokes the WAM at every control step with $H = 1$. Interactive operation typically reuses a trained model θ and extends H indefinitely while carrying state across calls through a key-value cache or a persistent latent. Section 4.5.

This decomposition is useful not because the four axes are separable, but because they separate different questions about a WAM. The substrate asks what kind of future variable s is represented, the backbone asks how that variable is predicted, the coupling asks how actions enter the prediction or are recovered from it, and the deployment regime asks how the resulting model is invoked inside a control loop. In practice, the substrate and backbone are usually fixed by training, while deployment and some wrapper-level coupling choices can be adjusted at inference. Changing the coupling family itself, however, often requires an additional action head, planner, or fine-tuning stage. We use these four questions as the organizing lens for the subsections below.

Tables 1 and 2 place every WAM in our census as a 4-tuple over these choices, and the subsections that follow walk through the choices one by one and anchor each one in the same equation.

4.2 Predictive Substrate: Where WAMs Dream

The first choice in Equation 8 is where the future variable $s_{t+1:t+H}$ lives. We call this space the predictive substrate. The substrate is not determined by the backbone alone, nor by the last tensor consumed by the action head. For example, a video-diffusion model may denoise a VAE grid and skip pixel decoding during control. As long as a fixed decoder maps that grid back to video, the predicted future remains pixel-grounded. Conversely, a policy may consume an intermediate state from a video-trained trunk, but that state is a feature substrate if it has no fixed observation decoder. In this section, we therefore classify a WAM by the representation in which it forms the future used for action prediction or evaluation.

We use four substrate categories. Pixel-grounded substrates represent future observations, either as decoded RGB, RGB-D, multi-view video, or as VAE / VQ latents with a fixed video decoder. Acoustic observation latents are rare in the current literature, so we treat them as multimodal observation-latent edge cases rather than as a separate category. Feature substrates represent learned states, teacher embeddings, or VLM token blocks without a fixed observation decoder. Geometric substrates represent physical structure or motion, such as optical flow, point tracks, depth, pose, motion vectors, or polylines. Affordance substrates represent task-specific maps or score fields, such as value maps, contact maps, semantic masks, or heatmaps. Some WAMs predict more than one kind of future in the same forward process. The tables mark these joint cases with \wedge .

4.2.1 Pixel-grounded substrates

The pixel-grounded category is the direct descendant of video generation. Its decoded form is an observation tensor,

$$s_{t+1:t+H} \in \mathbb{R}^{H \times C_o \times H_{\text{px}} \times W_{\text{px}}}, \quad (13)$$

where C_o is the observation channel count. A decoder-bound latent variant uses the same category when s is a VAE, VQ, or 3D-VAE code whose fixed decoder maps it back to video. The model may act before that decoder is called, but the future still lives in an observation coordinate frame. This criterion separates decoder-bound observation latents from feature latents below.

Decoded observation futures. Early WAMs take the direct route: generate a future image or video, then decode action from it. UniPi [29] establishes this pattern with text-conditioned video plus inverse dynamics, while AVDC [81] shows that dense correspondences from the decoded video can replace action labels. The same substrate category covers cascaded designs when the action module is separate. DreamGen [73] uses decoded

robot futures to obtain pseudo-actions before policy training, RIGVid [133] retargets tracked generated demonstrations, and VERA [94] lets a Jacobian inverse-dynamics model act from a generated visual lookahead. Later decoded-observation methods keep the same substrate but broaden the content of the forecast. Vidar [37], LVP [17], and EVA [162] use stronger Wan-class planners, while Veo-Act [199] uses a precision VLA executor on top of a generated visual plan. These wrappers still belong to the pixel-grounded category when geometry is recovered after the visual future has been produced.

Pixel-decodable latent futures. The second route keeps the future in the pixel-decodable latent space of a video generator. GR-1 [172] and GR-2 [16] predict VQ image tokens and action tokens in one autoregressive stream. PAD [52], UWM [213], WorldVLA [14], VideoVLA [140], DriveVA [106], and τ_0 -WM [209] instead couple action to a diffusion or flow-matching trajectory over VAE video latents. Hybrid entries such as F1 [111], Motus [10], Pelican-Unify 1.0 [198], NoiseGate [67], and WALL-WM [92] keep the same substrate while changing the routing between video and action streams. CKT-WAM [75] and FFDC-WAM [161] add a teacher context or verifier around that substrate, but the action-facing future remains the student or Motus pixel-decodable visual forecast. These works are not feature-substrate WAMs under the present criterion because the future is still a compressed observation under a fixed observation decoder.

Sparse, multimodal, and joint observation futures. The pixel-grounded category also covers reduced, multi-channel, and rare non-visual observation displays. SWEET [143] shrinks the forecast to task-critical keyframes. Audio-WM [194] predicts future Mel latents that condition a robot policy, and it is grouped here as an acoustic observation-latent edge case because it follows the same decoder-bound criterion without changing the category name. DriveDreamer-Policy [212] predicts decoded video together with depth, while MV-VDP [90] and ActionImages [203] predict RGB together with Gaussian end-effector heatmaps. OmniVTA [206] is also an observation-future case rather than a geometric one, because its future is a pair of visual and tactile decoder-bound latents with fixed decoders, not a coordinate tensor. HarmoWAM [36] is a dual observation case because its Wan backbone produces both rendered frames and pixel-decodable latent features that feed different action experts. The common trade-off is clear: observation futures preserve appearance, contact traces, and broad video priors, but they pay for detail that many action decisions do not need.

4.2.2 Feature substrates

The feature category covers futures that are learned states rather than decoder-bound observation states. A single sample can be written as

$$s_{t+1:t+H} \in \mathbb{R}^{H \times N_{\text{tok}} \times d_{\text{emb}}}, \quad (14)$$

with N_{tok} tokens per step and embedding dimension d_{emb} . The key property is absence of a fixed observation decoder. The state may be trained through future prediction, but it is interpreted only by the action pathway or by an embedding loss.

Encoder-only and feature-tap futures. This variant keeps the future-trained backbone but routes action through an internal state. Fast-WAM [193] is the clean encoder-only case: it keeps video co-training but masks the test-time future branch, so action is predicted from a hidden state shaped by future prediction. VPP [62], Genie Envisioner [102], VidMan [170], mimic-video [132], VAG [85], and WAV [91] keep more of the generative trajectory but decode action from latent or intermediate denoising states. LaMP [164] is the VLM-stack version of the same shortcut, exposing a motion-expert hidden state through gated cross-attention before the action expert fires. The boundary is that an auxiliary hidden state is not enough to make the substrate a feature future when the inference-time action-facing future is still a pixel-decodable visual forecast.

Teacher-target futures. A second feature route predicts the embedding of a frozen vision encoder, a multimodal encoder, or a policy/expert latent state. FLARE [204] aligns policy tokens to future-observation embeddings, and LDA-1B [112] learns dynamics, policy, and forecasting in a structured DINO feature space on a Qwen3-VL backbone. DAWN [110], CLWM [27], EgoExo-WM [152], DexWM [44], DDP [64], Feedback-WM [2], and FRAPPE [200] all keep the future in a learned feature space rather than in a decoder-bound observation latent. The benefit is compact prediction and strong semantic invariance. The cost is that the substrate has no direct visual-fidelity metric.

VLM-token futures. The third feature route represents the future as tokens in a VLM vocabulary. CoT-VLA [201] predicts visual chain-of-thought tokens before action generation. DUST [171] jointly diffuses a

next-state observation token and an action token, and ALAM [149] interleaves view-specific latent-transition tokens with action tokens under a structured attention mask. These methods use tokens as future carriers, not merely as action discretisation. The feature category is therefore broad, but its practical boundary is whether the predicted future lacks a fixed observation decoder and is not a physical coordinate tensor or task map.

4.2.3 Geometric primitive substrates

The third category exposes a structured object whose dimensions are physical coordinates. A common instance is dense optical flow, for which ϕ is an off-the-shelf flow operator and a single sample is

$$s_{t+1:t+H} \in \mathbb{R}^{H \times 2 \times H_{\text{px}} \times W_{\text{px}}}, \quad (15)$$

namely two channels carrying the (u, v) displacement field. Variants replace the flow field with N_{kp} tracked keypoints in $\mathbb{R}^{H \times N_{\text{kp}} \times 3}$, a depth tensor, a 6-DoF pose stream, a sparse waypoint polyline, or a macroblock motion-vector grid. A method belongs here when geometry is the predicted future representation, not when a tracker merely measures geometry after a pixel-grounded future has already been chosen.

Flow and point-motion futures. Im2Flow2Act [177] and 3DFlowAction [207] motivate the category by replacing full appearance with object motion. NovaFlow [87] and Dream2Flow [28] take the cascaded route by extracting 3D object flow from generated videos before planning. TraceGen [86] predicts 3D traces, while 3PoinTr [70] and RoboFlow4D [104] feed separately predicted point tracks or flow into downstream policies. 3D-FDP [126], PointWorld [68], and Dex4D [83] extend this idea from dense image flow to sparse or scene-level 3D point trajectories. The common payoff is a future that is closer to control than RGB.

Structured physical futures. TesserAct [202] reconstructs a 4D point cloud as the control-facing representation. HiF-VLA [103] uses MPEG-4 motion vectors as a compact future-motion field, and ICLR-VR [124] reduces the future to a five-point gripper polyline. JOPAT [49] is one prominent joint case because it denoises a pixel-decodable visual future together with point tracks and visibility. Geometric futures are cheaper and often easier to supervise than pixel-grounded futures, but they lose appearance cues that are irrelevant to motion yet useful for semantic grounding.

The structured tensor is far smaller than the pixel tensor at the same spatial resolution, and it can often be trained stably on less data. The trade-off is that these predictions are most reliable when motion, contact, or geometry carries the task-relevant information. They are less suitable when appearance or semantics is essential.

4.2.4 Affordance map substrates

The fourth category exposes a task-relevant label or score map. A single sample is

$$s_{t+1:t+H} \in \mathbb{R}^{H \times C \times H_{\text{px}} \times W_{\text{px}}}, \quad (16)$$

with C channels carrying value, affordance, contact likelihood, segmentation, or progress information. The substrate is not an observation, a feature, or a physical coordinate field. It is a spatial answer to a task-specific question.

Value, mask, and affordance futures. AIM [32] inserts a spatial value map between future visual dynamics and action, with intent-causal attention forcing action to pass through that value representation. PALM [107] predicts a bundle of future affordance maps, including object masks, contact likelihood, placement candidates, and motion regions. MWM [109] predicts semantic-mask dynamics rather than RGB, so its future is a label map rather than a generic hidden state. ActionImages [203] and MV-VDP [90] are joint pixel-grounded-and-affordance cases because they predict RGB together with Gaussian end-effector heatmaps. Affordance futures are therefore more task-specific than observation or feature futures. They are compact and often directly useful for manipulation, but their labels must be defined for the task at hand.

4.2.5 Substrate and coupling are distinct axes

The four categories classify the space of the future variable s , while Section 4.3 classifies how s and a are tied together. The distinction matters. Fast-WAM [193] is feature-substrate and encoder-only because its future

Predictive Substrate Categories

What future does the action pathway consume?

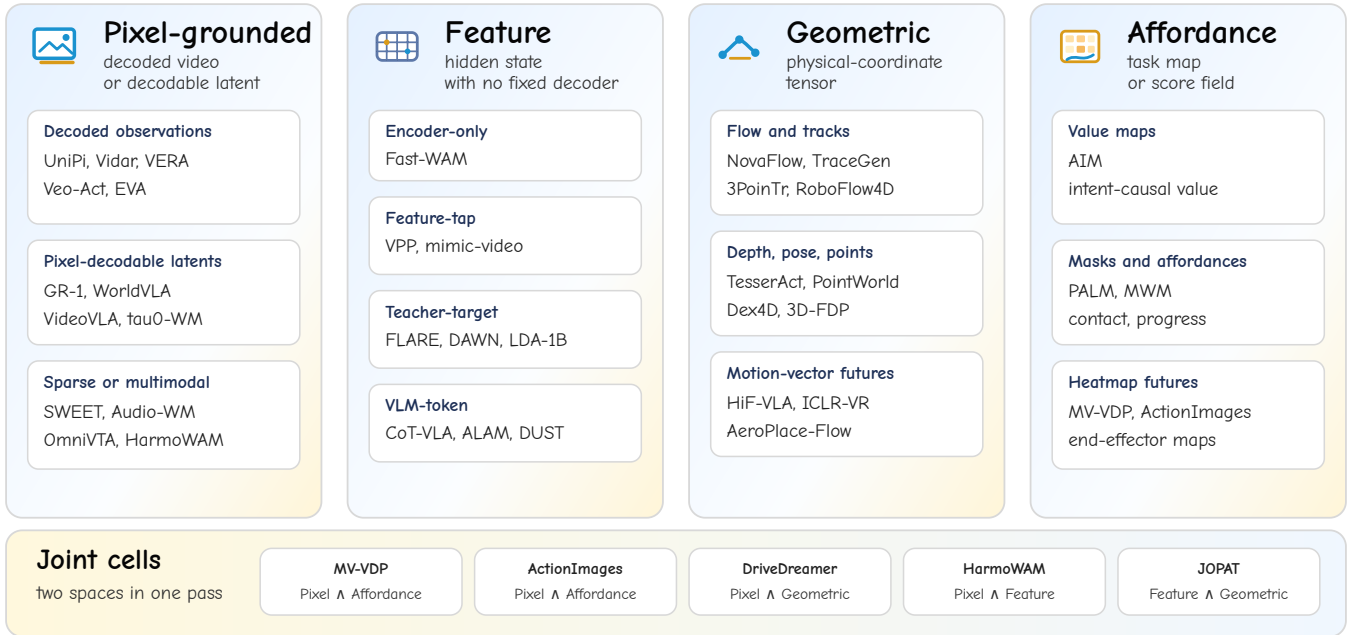


Figure 5 The four predictive-substrate categories of Section 4.2, classified by the representational space in which the WAM forms its future. Pixel-grounded substrates include decoded observations and pixel-decodable latents, with Audio-WM treated as a rare acoustic observation-latent edge case (14, 29, 172, 194, 209). Feature substrates include encoder-only, feature-tap, teacher-target, and VLM-token futures with no fixed observation decoder (62, 110, 149, 193, 204). Geometric substrates expose flow, point clouds, depth, polylines, or motion vectors (68, 103, 202). Affordance substrates expose value maps, masks, affordances, or heatmaps (32, 107, 109). Joint cells use the \wedge separator when one forward pass predicts futures in two categories.

branch is masked at inference and the action expert uses a hidden state. τ_0 -WM [209] is pixel-grounded-substrate and joint-denoising because it predicts a Wan VAE video latent together with action velocity. mimic-video [132] is feature-substrate and post-prediction-head because the action head taps an intermediate ODE state of a frozen video trunk. The substrate names where the WAM forms its future. The coupling axis names how that future reaches action. Figure 5 arranges the categories and representative members on the same panel.

4.3 Action Coupling: How Action Enters and Leaves

Action coupling is the structural decision that turns a predictive world model into a WAM. The factorization families in Equations 10, 11, and 12 pin down the main arrangements. Action-conditioned rollout is a composite family: an outer planner, policy, candidate sampler, or teleoperation stream supplies actions through q_ψ , and an action-conditioned world model predicts their consequences. Joint generation and post-prediction heads are action-producing families, either because the future action chunk is produced by the same joint model or because it is decoded from the produced future by a separate head. Within each arrangement, existing methods further differ on how an action is represented and on how many actions are produced per forward pass. Figure 6 sketches the three top-level arrangements and shows representative candidate forms inside the first. Notely, in this section, when we refer to a method as non-census, we mean that it is cited for background but is not listed as a WAM in Tables 1 and 2.

Three Action-Coupling Families

Action can shape rollout, share generation, or decode from a predicted future.

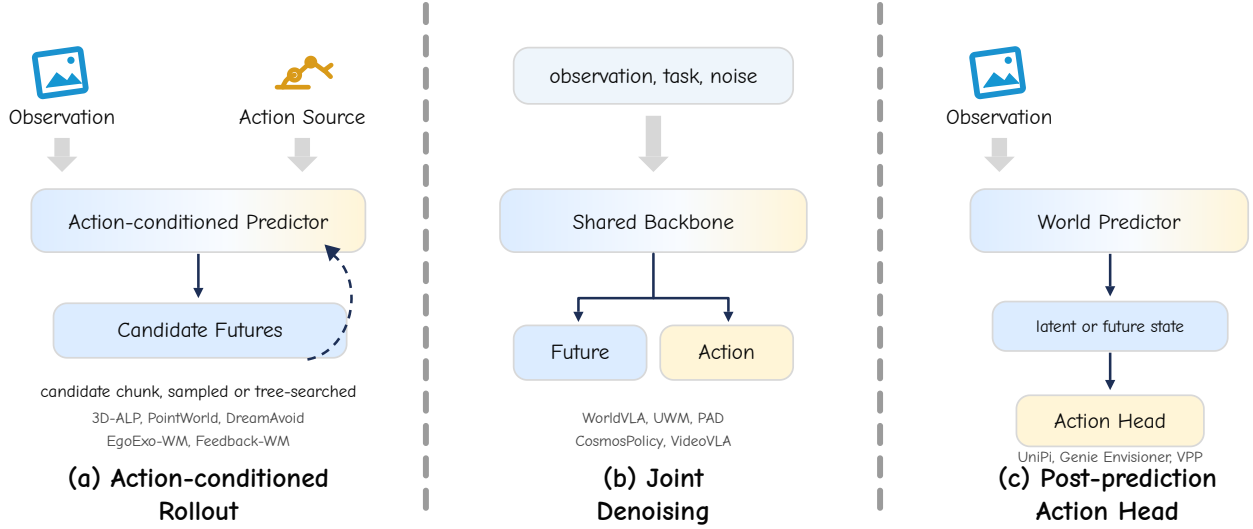


Figure 6 Three common action-coupling families in existing WAMs. The arrangements differ in where the action representation is bound to the substrate prediction, and that binding is what determines latency, controllability, and per-step inference cost.

4.3.1 Action-conditioned rollout

Action-conditioned rollout composes an action source with an action-conditioned world predictor. Repeating Equation 10,

$$p_{\theta, \psi}(s_{t+1:t+H}, a_{t:t+H-1} | c) = \begin{cases} q_{\psi}(a_{t:t+H-1} | c) p_{\theta}(s_{t+1:t+H} | c, a_{t:t+H-1}), & \text{chunk-level,} \\ \prod_{k=0}^{H-1} q_{\psi}(a_{t+k} | h_k, c) p_{\theta}(s_{t+k+1} | h_k, a_{t+k}, c), & \text{step-wise,} \end{cases} \quad (17)$$

$$h_k = (s_{t+1:t+k}, a_{t:t+k-1}).$$

The first line is chunk-level rollout. The future action chunk is output by q_{ψ} and enters the substrate predictor as conditioning. The second line is step-wise rollout. It alternates action selection with action-conditioned prediction, so q_{ψ} can update the next action after each predicted state. Training usually applies the standard substrate-prediction loss of Section 4.4 to p_{θ} , with the proposed action or action chunk added to the conditioning side. The action source may be a planner, a policy, a human command stream, a fixed candidate sampler, or, in the step-wise background lineage, a diffusion policy queried after each predicted state.

In the current census, the WAMs listed in the tables in this family all sit on the chunk-level submode. The step-wise submode is retained to connect the WAM notation to the model-based RL and interactive-prediction lineage discussed below.

The chunk-level submode is the natural form when the method scores candidate futures before execution. Cosmos-Transfer1 [129] illustrates a nearby non-census case: segmentation, depth, edge, and other spatial control maps condition world generation for data generation, but the method does not use generated futures as an inference-time robot-action selector. PointWorld [68] conditions a point-flow dynamics model on a proposed sequence of robot point flows from MPPI and scores the predicted scene flow, so action search happens outside the predictor rather than inside the trunk. DreamAvoid [33] samples flow-policy action chunks only at detected critical phases, renders each chunk with DreamDojo, and executes the candidate with the highest predicted progress value. EgoExo-WM [152] conditions a DINOv3-L latent predictor on horizon-length UniEgoMotion candidate body-motion sequences, rolls each candidate forward under MPC, and selects the sequence whose predicted latent ends closest to a visual goal.

Tree-search and feedback-guided variants keep the same family but change the cadence of candidate evaluation. 3D-ALP [142] uses MCTS to propose joint actions, converts each proposal into a 3D camera pose through kinematics, and uses a 3D-consistent renderer as a rollout oracle before executing the selected action. Feedback-WM [2] uses a lightweight latent transition model to guide a diffusion policy while that policy forms an action chunk, passes the corrected chunk to the retained world-prediction path, and then updates an auxiliary feedback state from the observation after execution. These variants explain why the family should not be split by whether actions arrive as one candidate chunk, a tree node, or a corrected action sequence. In all cases, actions are supplied to a predictor that evaluates their consequences.

The step-wise submode remains useful for background and for future WAMs whose action source reacts inside the imagined trajectory. AdaWorld [40] conditions an SVD-style latent-video diffusion predictor at each rollout step, although it is not part of the census. FlowDreamer [50] cascades a geometric-primitive flow prediction into a diffusion next-frame generator, a two-stage variant of the same arrangement. The latent-imagination Dreamer line of PlaNet [55], DreamerV3 [57], Dreamer 4 [58], and TransDreamer [18] supplies the model-based RL ancestor of this design, where a planner, actor, critic, or compact dynamics model advances imagined steps depending on the member. The non-WAM antecedent iVideoGPT [174] interleaves observation, action, and reward tokens in an autoregressive stream, while InteractiveWorldSimulator [166] and RoboScape [139] extend interactive prediction with streaming inputs or physics-aware branches.

Action-conditioned rollout is therefore a composite WAM factorization rather than an action-producing world model by itself. Its advantage is counterfactual control: candidate actions shape the future before selection. Its limitation depends on the submode. Chunk-level rollout can evaluate many candidates in parallel, but it cannot react to a control signal that changes inside the predicted window. Step-wise rollout can react after each predicted state, but it loses the parallelism of a single generated window.

4.3.2 Joint generation

The joint factorization produces substrate and action from a single generative process on a shared backbone. Repeating Equation 11,

$$p_{\theta}(s_{t+1:t+H}, a_{t:t+H-1} | c). \quad (18)$$

Training is often expressible as, or approximated by, a weighted sum of a substrate-side generative loss and an action-side regression or classification loss,

$$\mathcal{L}_{\text{joint}}(\theta) = \mathcal{L}_{\text{gen}}(s) + \lambda \mathcal{L}_{\text{act}}(a), \quad (19)$$

where \mathcal{L}_{gen} is a diffusion or autoregression loss on s and \mathcal{L}_{act} is a flow-matching, regression, or cross-entropy loss on a . Section 4.4 reuses Equation 19 as the canonical training objective of hybrid backbones.

Joint diffusion of video and action. The joint-diffusion group differs less in its factorization than in the interface through which the two streams meet. UWM [213], CosmosPolicy [79], VideoVLA [140], DriveVA [106], and DriveWAM [141] place video-side latents and action-side variables in the same denoising or flow-matching state, so one sampled trajectory carries both the imagined future and the action chunk. CoVAR [183], UD-VLA [19], GigaWorld-Policy [186], and CKT-WAM [75] keep the same joint target but change the information routing through bridge attention, discrete diffusion, action-centered masking, or teacher-context prefixes. Multisubstrate variants add another future channel without changing the coupling: AIM [32] inserts a value map, DriveDreamer-Policy [212] orders depth, video, and action inside one training pass, MV-VDP [90] and ActionImages [203] attach end-effector heatmaps to the visual stream, VTAM [192] adds a tactile force proxy, HiF-VLA [103] adds motion-vector foresight tokens, and DUST [171] applies the same idea to a VLM-substrate setting. The useful distinction is therefore whether the added channel changes the substrate, not whether the action is still jointly produced.

Joint autoregression of frame and action tokens. GR-1 [172] and GR-2 [16] jointly predict next-frame and next-action tokens autoregressively over a discrete codebook. WorldVLA [14] carries the same recipe to a Chameleon backbone with an action-chunk attention mask between action and frame tokens. PhysGen [144] adapts the autoregressive NOVA backbone with continuous physical tokens and a per-token diffusion detokenizer, so that autoregressive image and action streams share a joint training loss in the sense of Equation 19. ICLR-VR [124] serializes state, reasoning, and action tokens into a single Llama2-style causal stream under a combined next-token loss, with the gripper-keypoint polyline decoded immediately before each action chunk.

Mixture-of-experts hybrids that joint-denoise. The hybrid branch uses separate experts or heads so that video and action can share a trunk without forcing every layer to treat both variables identically. UVA [93] and PAD [52] are the clean baselines: one shared pass supports both visual prediction and action, and the visual branch can be bypassed when only control is needed. F1 [111], Motus [10], MotuBrain [121], RynnVLA-002 [13], BagelVLA [63], Pelican-Unify 1.0 [198], and ALAM [149] instantiate the mixture pattern with different expert allocations for understanding, video, latent transition, and action streams. A second set keeps the same joint target but changes the schedule: DreamZero [188] uses one shared timestep across a Wan-style video-action state, X-WAM [51] lets action finish before multi-view video, NoiseGate [67] learns per-latent denoising increments, τ_0 -WM [209] samples video and action flow times independently, and WALL-WM [92] couples action queries into matched video blocks at every depth. CLWM [27], DAWN [110], JOPAT [49], and DDP [64] show the same coupling on feature, track, or compact latent substrates, while FFDC-WAM [161] keeps the Motus joint rollout and adds a verifier that decides when the cached future is still trustworthy. The family is diverse, but the shared claim is narrow: the action and the future substrate are trained as mutually constrained variables rather than as a planner output followed by a passive decoder.

The benefit of the family is one coupled sample or rollout that produces both predictions and a strong mutual consistency between them. The trade-off is potential training instability because the generation and action losses can pull the shared representation in different directions. Practitioners therefore often pretrain the substrate head on video and introduce the action head later, sometimes with a careful λ schedule in Equation 19.

4.3.3 Post-prediction action head

The post-prediction head factorizes into a substrate generator and a smaller action expert q_ψ that decodes action from the substrate. Repeating Equation 12,

$$p_\theta(s_{t+1:t+H}, a_{t:t+H-1} | c) = p_\theta(s_{t+1:t+H} | c) q_\psi(a_{t:t+H-1} | s_{t+1:t+H}, c), \quad (20)$$

with θ frozen during q_ψ training in many implementations. This arrangement is common because it lets a pretrained predictor stay fixed while the action head changes across embodiments.

Video-plan then action-recovery WAMs. UniPi [29] a canonical early example: a text-conditioned diffusion planner first synthesizes a future image trajectory, and a task-specific inverse-dynamics module then recovers the low-level actions between generated frames. VLP [30] adds a VLM policy and value-guided search over abstract text actions, but the executable controls still come from goal-conditioned policies that consume the synthesized video plan. EVA [162] pairs a Wan2.1-14B planner with a frozen spatial-softmax IDM and uses the IDM-derived smoothness-and-limits reward to align the planner via GRPO without folding the IDM into the joint denoiser. CreFlow [125] post-trains the Vidar text-image-to-video trunk against an LTL-grounded corrective reflow loss and keeps the Vidar inverse-dynamics expert frozen. These methods are often described as video planning or video world modeling, but under the notation here their action coupling is the post-prediction factorization of Equation 20.

Action expert on a latent, token, or multimodal substrate. This group keeps the action expert separate, so the key question is what future object q_ψ consumes. Some members expose a task-structured future before action decoding: $\pi_{0.7}$ [72] uses BAGEL-style subgoal images as visual targets, CoT-VLA [201] uses visual chain-of-thought tokens, PALM [107] uses affordance maps, and OmniVTA [206] uses visual and tactile latents whose fixed decoders tie them back to observation space for slow-fast contact-rich control. Feature-substrate members differ mainly in where the future feature comes from. FLARE [204], LDA-1B [112], FRAPPE [200], and GigaBrain-0.5M* [158] rely on future embeddings or value-bearing latent tokens, while VPP [62], VideoPolicy [100], Genie Envisioner [102], LaMP [164], and VAMPO [42] intercept intermediate states of a video-trained trunk. The action expert can therefore be small, but its reliability depends on whether the chosen substrate preserves the action-relevant part of the forecast.

Action expert decodes structured geometry from a generated future. The geometric post-head line uses motion or structure as the intermediate action-facing representation. Im2Flow2Act [177], 3DFlowAction [207], 3D-FDP [126], NovaFlow [87], Dream2Flow [28], and AeroPlace-Flow [120] make flow the intermediate future, while TesserAct [202], 4DGen [108], MVISTA-4D [160], TC-IDM [119], and Dex4D [83] recover pose or point structure before control. Earlier decoded-video cascades such as AVDC [81], Dreamitate [99], This&That [157],

ARDuP [65], Gen2Act [9], GR-MG [89], and GraspDreamer [147] fit the same factorization even when the extracted geometry is implicit in a tracker or learned head. RoboEnvision [184], Vidar [37], LVP [17], Veo-Act [199], Say-Dream-Act [48], NovaPlan [38], and MoLA [96] keep the same post-head order but add stronger video planners, VLM scoring, or modality-specific inverse dynamics. The shared lesson is that a generated future can be action-facing even when the final action module is trained separately.

Generated future is partly or fully optional at inference. The optional-future pattern trains with a substrate generator, then shortens or bypasses that generator when the inference-time action path can use an internal state. VidMan [170] is the clearest representative: it mounts a policy head on frozen video-diffusion features rather than waiting for a rendered future. mimic-video [132], S-VAM [181], DiT4DiT [113], Fast-WAM [193], and VAG [85] make the same move through ODE checkpoints, distilled foresight states, fixed denoising timesteps, masked future-video tokens, or pooled clean latents. WAV [91], MWM [109], DexWM [44], and SWEET [143] instead keep a compact action-facing substrate, such as valued latent rollouts, semantic-mask dynamics, DINO latents, or sparse keyframes. LingBot-VA [88] and HarmoWAM [36] show the mixture-backbone version, where visual dynamics remain available but a lighter action expert or gate handles the frequent control path. This design keeps q_ψ cheap and embodiment-specific. But the risk is equally clear, since the skipped future can help only if the retained state still preserves the part of the forecast that action needs.

4.3.4 Action representation and chunk size

Across the three coupling families, the action variable $a_t \in \mathcal{A}$ is parameterized differently. Three representations are common. The first is per-step continuous controls, with $\mathcal{A} = \mathbb{R}^{d_a}$ for joint velocities, end-effector deltas, or gripper commands. The second is discrete tokens from a fixed codebook, as in autoregressive joint denoisers such as WorldVLA [14], where \mathcal{A} is a finite vocabulary obtained from per-dimension binning. The third is learned latent actions that act as currency between the world model and the policy. In that case $\mathcal{A} = \mathbb{R}^{d_z}$ for some learned $d_z \ll d_a$, as in Motus [10], ALAM [149], and LDA-1B [112]. The chunk size is the second nested choice. A model that produces one action per forward pass pays a low per-action latency but invokes the full backbone at the control frequency. A model that produces a long chunk per pass spreads the backbone cost over the chunk but cannot react during it. Closed-loop control on real hardware pushes methods toward smaller chunks or toward cheaper backbones, and many of the architectural choices in Section 4.4 can be reinterpreted as a response to that pressure.

4.4 Architectural Backbone: How the Prediction Is Produced

Once a substrate space \mathcal{S} is fixed and a coupling factorization is chosen, the next ingredient is the function family that realizes $p_\theta(\cdot | c)$. We group current backbones into five families, and for each we list the defining parameterization, the canonical training objective, and the WAMs that use it.

4.4.1 Iterative-denoising backbones (video diffusion)

The diffusion family writes p_θ as the marginal of a learned reverse process over a sequence of progressively denoised latents $s^{(N_{\text{diff}})}, s^{(N_{\text{diff}}-1)}, \dots, s^{(0)} = s_{t+1:t+H}$:

$$p_\theta(s_{t+1:t+H} | c) = \int p(s^{(N_{\text{diff}})}) \prod_{n=1}^{N_{\text{diff}}} p_\theta(s^{(n-1)} | s^{(n)}, c) \, ds^{(N_{\text{diff}}:1)}, \quad (21)$$

where $p(s^{(N_{\text{diff}})})$ is a Gaussian noise prior and each reverse step is parameterized by a noise-prediction network $\epsilon_\theta(s^{(n)}, n, c)$. Training minimizes the standard score-matching loss

$$\mathcal{L}_{\text{diff}}(\theta) = \mathbb{E}_{n,s^{(0)},\epsilon} \left[\|\epsilon - \epsilon_\theta(s^{(n)}, n, c)\|^2 \right], \quad s^{(n)} = \alpha_n s^{(0)} + \sigma_n \epsilon. \quad (22)$$

At inference, sampling draws $s^{(N_{\text{diff}})} \sim p(s^{(N_{\text{diff}})})$ and runs N_{diff} reverse steps.

Founding video-diffusion backbones. The dominant topology is a video diffusion transformer that attends jointly over space and time, as in Latte [114], CogVideoX [185], Wan [156], and Cosmos-Predict2 [128]. AnimateDiff [53] represents an earlier generation that adapts a frozen image-diffusion backbone with a learned temporal motion module.

Because standard diffusion denoises a whole window at once rather than advancing one element at a time, the coupling-type landscape is narrower than the autoregressive one. The census mainly uses diffusion in three ways:

$$(\mathcal{Y}, c_{\text{diff}}) = \begin{cases} (s_{t+1:t+H}, (c, a_{t:t+H-1})), & \text{action-conditioned rollout,} \\ ((s_{t+1:t+H}, a_{t:t+H-1}), c), & \text{joint generation,} \\ (s_{t+1:t+H}, c) \text{ followed by } q_\psi(a_{t:t+H-1} \mid s_{t+1:t+H}, c), & \text{post-prediction head.} \end{cases} \quad (23)$$

Pure diffusion supports action-conditioned rollout most naturally when the action source proposes a whole candidate chunk before the reverse chain begins, as in DreamAvoid [33]. Step-wise rollout is less natural for a windowed reverse chain, so it appears mainly in non-census foundation or simulator works such as AdaWorld [40] and FlowDreamer [50], or in hybrid wrappers that alternate an action source and a predictor.

Post-prediction head (Eq. 12). Many diffusion WAMs first denoise the substrate $s^{(0)}$ and then decode action from it with a separate module q_ψ . The important variation is which substrate q_ψ consumes, not the mere fact that the action head is separate.

Decoded-observation members use the generated visual future as the common interface. UniPi [29] is the representative case, since a text-conditioned diffusion rollout is followed by inverse dynamics. VLP [30] changes the conditioning context before sampling, while AVDC [81], Dreamitate [99], This&That [157], Gen2Act [9], and GR-MG [89] recover actions, tool motion, or goals from the produced frames. Later visual planners mostly strengthen the generator or the search wrapper: RoboEnvision [184], Vidar [37], MVISTA-4D [160], LVP [17], Veo-Act [199], and Say-Dream-Act [48] keep the same factorization while changing horizon, view structure, executor, or sampling schedule.

Pixel-latent and geometric variants keep the post-head order but make the substrate cheaper or closer to control. ARDuP [65], VILP [178], DiT4DiT [113], and OmniVTA [206] act from latent video, intermediate denoising states, or visual-tactile decoder-bound latents. Im2Flow2Act [177] is the representative geometric instance because it denoises object flow as the action-facing future. 3DFlowAction [207] and Tesseract [202] extend the same idea to 3D flow or reconstructed point clouds. The trade-off is direct: these substrates reduce appearance burden, but their usefulness depends on whether the chosen geometry captures the task-relevant contact and motion.

Feature-level post heads push the shortcut further by consuming latent states rather than decoded futures. VPP [62], Video Policy [100], Genie Envisioner [102], Act2Goal [208], WAV [91], VAG [85], EnerVerse [66], VidMan [170], mimic-video [132], S-VAM [181], and Fast-WAM [193] use variations of this shortcut, such as tapping intermediate denoising states, distilling foresight features, fixing a small number of denoising steps, or bypassing explicit video generation at inference. Related compact-substrate variants, such as MWM [109], pursue a similar latency goal through task maps rather than through a pure feature tap.

Joint generation (Eq. 11). The second group of diffusion WAMs models the future substrate s and the future action a within a coupled denoising or flow-matching process, so Equation 21 operates on the joint variable (s, a) . PAD [52] is a compact image-action DiT instance, while VideoVLA [140], CosmosPolicy [79], CoVAR [183], UD-VLA [19], GigaWorld-Policy [186], and CKT-WAM [75] vary the shared denoising state through a pretrained video trunk, bridge attention, discrete tokens, action-centered masking, or transfer prefixes. GigaWorld-Policy is the action-centered case in this group: its causal mask prevents future-video tokens from influencing action tokens, so deployment can sample the action chunk without instantiating future-video tokens. AIM [32], X-WAM [51], and AdaWorldPolicy [191] show the same diffusion coupling with value maps, multi-view 4D video, or force-prediction modules attached to the joint state. The benefit is a coupled generative process that encourages consistency between action and future prediction. The trade-off is potential training instability because the generation and action losses can impose competing demands on the shared representation. Practitioners therefore often pretrain the substrate head on video and introduce the action head later, sometimes with a careful λ schedule.

Flow matching as a variant of Equation 21. A large subset of the diffusion family instantiates ϵ_θ as a velocity-field predictor under flow matching rather than as a noise predictor under score matching, but both still satisfy the iterative-reverse-chain parameterization of Equation 21. Conditional flow matching learns $v_\theta(x_t, t, c)$ to regress the linear interpolation velocity $x_1 - x_0$ under $\mathcal{L}_{\text{FM}} = \mathbb{E} \|(x_1 - x_0) - v_\theta(x_t, t, c)\|_2^2$, and at inference

the same Euler reverse integration takes the place of the noise-prediction reverse chain. The flow-matching variant is therefore not a new backbone family but a different choice of the score parameterization inside Equation 21. Among the post-prediction-head WAMs above, 3D-FDP [126], EVA [162], VAMPO [42], and MoLA [96] run flow matching as the score parameterization. Among the joint-denoising WAMs, DriveVA [106], ActionImages [203], MV-VDP [90], and CreFlow [125] take the same route on a Wan-class trunk. The choice between score matching and flow matching remains separable from the coupling family within Equation 21.

4.4.2 Autoregressive backbones (next-frame, next-token)

The autoregressive family is a backbone parameterization, not an action-coupling regime by itself. It realizes whichever generative factor from Section 4.3 is assigned to it by serializing that factor’s output into a causal stream. Let

$$y_{1:M} = \text{Serialize}(\mathcal{Y}), \quad (24)$$

where \mathcal{Y} is the set of variables generated by the autoregressive backbone, and let c_{ar} denote the local context supplied to that backbone. The generic autoregressive parameterization is

$$p_{\theta}(y_{1:M} | c_{\text{ar}}) = \prod_{j=1}^M p_{\theta}(y_j | y_{<j}, c_{\text{ar}}), \quad (25)$$

where y_j may be a frame, a latent-grid block, a discrete visual token, a continuous latent token, an action token, or a compact latent state. Training minimizes the corresponding next-element loss

$$\mathcal{L}_{\text{ar}}(\theta) = - \sum_{j=1}^M \log p_{\theta}(y_j | y_{<j}, c_{\text{ar}}), \quad (26)$$

with cross-entropy for discrete tokens and a continuous likelihood, regression, diffusion, or flow-matching loss for continuous elements. At inference, a key-value cache reuses prefix computation, so each newly appended element requires one forward pass over the new tokens rather than recomputation over the whole history.

The pair $(\mathcal{Y}, c_{\text{ar}})$ is inherited from the coupling choice rather than from autoregression itself:

$$(\mathcal{Y}, c_{\text{ar}}) = \begin{cases} (s_{t+1:t+H}, (c, a_{t:t+H-1})) \text{ or } (s_{t+k+1}, (h_k, a_{t+k}, c)), & \text{action-conditioned rollout,} \\ ((s_{t+1:t+H}, a_{t:t+H-1}), c), & \text{joint autoregression,} \\ (s_{t+1:t+H}, c) \text{ followed by } q_{\psi}(a_{t:t+H-1} | s_{t+1:t+H}, c), & \text{post-prediction head.} \end{cases} \quad (27)$$

Equation 27 is only a backbone-level map. The full joint factorizations remain those in Section 4.3. Its purpose is to prevent the common confusion that autoregression implies joint action prediction. In many WAMs the autoregressive stream predicts only the future substrate. The future action may be supplied to the rollout, generated in the same stream, or decoded later by q_{ψ} .

Autoregressive video priors (not WAMs). The following examples provide the video prior but not the WAM coupling. The list is selective rather than exhaustive, because these works cover the substrate forms that later WAMs reuse. VideoGPT [182] gives the VQ-VAE plus GPT recipe over discrete video latents. VideoPoet [82] is a decoder-only multimodal video language model over tokenized visual and audio streams. NOVA [26] removes the discrete codebook and predicts continuous video latents. These works define reusable autoregressive video priors, but they are not WAMs by themselves. Earlier GAN-based video generators such as MoCoGAN [153] belong to the historical substrate discussion of Section 4.2. They are not instances of the next-element parameterization in Equation 25.

Action-conditioned rollout (Eq. 10). The following methods use autoregression for action-conditioned sequential dynamics. iVideoGPT [174] brings the VideoGPT-style token stream into interactive world modeling by interleaving observations, actions, and rewards. AdaWorld [40] is the boundary case in this group. The paper calls it an autoregressive world model because inference repeats action-conditioned next-frame prediction, but each local transition is an SVD-based denoising problem. It is therefore best understood as an autoregressive rollout with a diffusion transition, not as a homogeneous next-token decoder. PlaNet [55], DreamerV3 [57], Dreamer 4 [58], and TransDreamer [18] are latent-imagination methods based on compact dynamics models.

These methods explain the sequential-dynamics side of the family, but they usually pair the dynamics model with a separate actor, planner, or critic and therefore serve here as background rather than census WAMs.

Joint autoregression (Eq. 11). The following methods use autoregression as the trunk for streams that include both future substrate elements and action-facing elements. WorldVLA [14] is the clearest discrete instance: it serializes language, image, and action tokens in one autoregressive token framework. PhysGen [144] is the continuous-latent counterpart, adapting the NOVA [26] autoregressive video prior to continuous physical tokens. GR-1 [172] establishes the GPT-style video-generative trunk for robot manipulation, and GR-2 [16] scales that trunk with larger video-language pretraining and robot-trajectory fine-tuning. From the backbone perspective, the shared property is causal next-element prediction. The exact action head, detokenizer, or joint-output factorization is the coupling detail handled in Section 4.3.

Post-prediction head (Eq. 12). The post-head line of Equation 27 is still possible for autoregressive backbones, even when it is less common than joint token prediction. LingBot-VA [88] is the closest current member because it models visual dynamics before inverse dynamics, but it uses an autoregressive diffusion mixture-of-transformers rather than the next-element backbone of Equation 25. The distinction matters because the backbone can be causal while the action interface remains a separate decoder over the predicted substrate.

The common pressure across the family is serial cost. Autoregression avoids the many denoising steps of diffusion and benefits from caching, but it still advances the generated stream one element at a time. For WAMs, the relevant latency is therefore governed by what the stream contains: full frames, latent blocks, visual tokens, action tokens, or compact latent states. When an autoregressive WAM misses a real-time control budget, the first places to inspect are the serialized substrate size, the action-chunk length, and whether actions are produced inside the same causal stream or decoded by a cheaper head.

4.4.3 Joint-embedding-predictive backbones

The joint-embedding-predictive family realizes $p_\theta(s_{t+1:t+H} | c)$ as a Dirac at a deterministic predictor rather than as a stochastic density. It learns a context encoder \mathbf{E}^{ctx} and an exponential-moving-average target encoder \mathbf{E}^{tgt} together with a predictor f_θ^{pred} that maps context features to target features. Given a context view x_{ctx} of the past and a target view x_{tgt} of the future, the objective is

$$\mathcal{L}_{\text{jepa}}(\theta) = \mathbb{E}_{x_{\text{ctx}}, x_{\text{tgt}}} \left[\left\| f_\theta^{\text{pred}}(\mathbf{E}^{\text{ctx}}(x_{\text{ctx}})) - \text{sg}[\mathbf{E}^{\text{tgt}}(x_{\text{tgt}})] \right\|^2 \right], \quad (28)$$

with $\text{sg}[\cdot]$ a stop-gradient. Under the unified contract of Equation 8, the substrate $s_{t+1:t+H}$ is $\mathbf{E}^{\text{tgt}}(x_{\text{tgt}})$, and the predictive distribution $p_\theta(s_{t+1:t+H} | c)$ collapses to a Dirac at $f_\theta^{\text{pred}}(\mathbf{E}^{\text{ctx}}(x_{\text{ctx}}))$. The JEPA backbone never inverts \mathbf{E}^{tgt} , which makes the inference loop cheap and is the source of the family’s appeal. The backbone predicts the substrate only. Actions enter through one of two coupling patterns:

$$(\mathcal{Y}, c_{\text{jepa}}) = \begin{cases} (s_{t+k+1}, (h_k, a_{t+k}, c)) \text{ at rollout step } k, & \text{action-conditioned rollout,} \\ (s_{t+1:t+H}, c) \text{ followed by } q_\psi(a_{t:t+H-1} | s_{t+1:t+H}, c), & \text{post-prediction head.} \end{cases} \quad (29)$$

Action-conditioned rollout treats the JEPA predictor as a world model and recovers actions through planning over imagined consequences. A post-prediction head predicts the full substrate window first and then decodes action from it with a separate expert q_ψ .

I-JEPA [3] introduces the architecture on images, predicting the latent representations of large masked target blocks from a spatially distributed context block. A-JEPA [35] carries the same recipe to audio spectrograms under a curriculum time-frequency masking strategy, and MC-JEPA [6] couples the prediction objective with optical-flow learning so that one shared encoder serves both motion and content. V-JEPA [7] and V-JEPA 2 [4] extend the idea to video through pure feature prediction, and the V-JEPA 2 action-conditioned variant shows how such features can support MPC through action-conditioned latent rollout. The tabulated WAMs use this lineage in two ways. FLARE [204] pairs a deterministic JEPA predictor with a downstream action expert in the post-prediction-head factorization by aligning additional DiT tokens to future-observation embeddings from a frozen teacher. DAWN [110] keeps V-JEPA 2-style feature tokens inside a hybrid WAM that alternates compact latent-world prediction with action denoising. The common pressure is that the embedding space has no intrinsic visual quality measure, so downstream-task accuracy becomes the only standard test of the world model.

4.4.4 Hybrid backbones (generative head plus action head)

The hybrid family runs a generative head and an action head on the same backbone. Schematically,

$$p_{\theta}(s_{t+1:t+H}, a_{t:t+H-1} | c) \propto g_{\theta}(c) \longrightarrow \{h_{\theta}^s, h_{\theta}^a\}, \quad (30)$$

where g_{θ} is a shared encoder-decoder trunk and $h_{\theta}^s, h_{\theta}^a$ are two output heads. Training minimizes the joint coupling loss of Equation 19 from Section 4.3, namely $\mathcal{L}_{\text{gen}}(s) + \lambda \mathcal{L}_{\text{act}}(a)$, with the substrate head and action head sharing the trunk g_{θ} .

Because the shared trunk g_{θ} feeds both heads in one place, the hybrid family is the backbone in which every coupling family of Section 4.3 can appear, and which one a given model realizes is fixed by how the substrate head h_{θ}^s and the action head h_{θ}^a are ordered around that trunk:

$$(h_{\theta}^s, h_{\theta}^a) = \begin{cases} (s, a) \text{ jointly from } g_{\theta}(c), & \text{joint generation,} \\ s \text{ from } h_{\theta}^s, \text{ then } p_{\theta}^a(a | s, c), & \text{post-prediction head,} \\ q_{\psi}(a | c) \text{ or external search, then } p_{\theta}^s(s | c, a), & \text{action-conditioned rollout.} \end{cases} \quad (31)$$

Here p_{θ}^s and p_{θ}^a denote the substrate-prediction and action-prediction distributions implemented by h_{θ}^s and h_{θ}^a when a shared trunk feeds multiple heads. The action-producing couplings place the action inside or after the shared pass, whereas action-conditioned rollout wraps the same trunk with a planner, sampler, tree search, or feedback-guided action proposal.

Joint generation (Eq. 11). Joint generation is the native arrangement of the family, because a single pass of $g_{\theta}(c)$ produces the future substrate s and the future action a under the loss of Equation 19. The plain shared-trunk form is easiest to see in UVA [93]: the video head shapes the representation during training, and the action head can be used alone at inference. UWM [213], DUST [171], VTAM [192], DDP [64], DriveDreamer-Policy [212], and DriveWAM [141] keep that contract while changing the substrate from pixel-decodable latents to VLM tokens, tactile force proxies, compact latent states, or driving-specific video and depth.

Mixture-of-Transformer variants keep joint generation but separate expert routes so that video and action do not have to share every layer. F1 [111] is the representative mixture case, with generation and action experts coupled by next-scale prediction. Motus [10], MotuBrain [121], X-WAM [51], and RynnVLA-002 [13] vary the same idea through optical-flow latent actions, Vidu latents, asynchronous denoising, or an autoregressive-image plus continuous-action split. BagelVLA [63], CLWM [27], Pelican-Unify 1.0 [198], DAWN [110], NoiseGate [67], WALL-WM [92], and τ_0 -WM [209] then vary the schedule and exchange pattern through same-timestep attention, feature flow matching, iterative refinement, per-latent timestep gates, depth-wise cross-attention, or propose-rank-rectify control.

FFDC-WAM [161] sits at the boundary between cached chunked rollout and closed-loop re-planning: the Motus joint rollout remains the action-facing substrate, while a lightweight verifier decides when cached imagination should be trusted. Across the group, the shared benefit is consistency between the predicted world and the chosen action. The cost is that the action and future heads must train together without one objective overwhelming the other.

Post-prediction head (Eq. 12). A different arrangement keeps the shared trunk but orders the heads, so h_{θ}^s completes the substrate window before h_{θ}^a decodes the action from it. FRAPPE [200] is the teacher-feature representative, routing multiple predicted future embeddings into one action transformer. Genie Envisioner [102], GigaBrain-0.5M* [158], and LaMP [164] instead expose video-trained latent, value, or motion-expert states to the action expert. LingBot-VA [88] and HarmoWAM [36] keep a visual-dynamics branch active but hand frequent control to inverse-dynamics or gated action experts. Dex4D [83] and DexWM [44] show the same order on geometric or DINO-latent futures. The action expert here conditions only on the predicted substrate, so it stays cheap and can be retrained per embodiment without touching the trunk, provided that substrate is action-informative.

Action-conditioned rollout (Eq. 10). Action-conditioned rollout enters the family when an action source is specified before or during consequence prediction. 3D-ALP [142] is the pixel-grounded tree-search version, where MCTS proposes joint actions, kinematics converts them into 3D camera poses, and a 3D-consistent world model renders candidate futures for scoring. PointWorld [68] is the geometric chunk-level version,

composing an MPPI sampler over end-effector trajectories with an action-conditioned predictor over a robot point-flow substrate. EgoExo-WM [152] uses horizon-length body-motion candidates from UniEgoMotion to condition a DINOv3-L latent predictor under MPC. Feedback-WM [2] uses the retained predictor to guide or correct proposed action chunks after the action source has formed them. The retained world-prediction module therefore scores or corrects proposed actions instead of producing them as an unconstrained action head.

Across the hybrid backbone family, the main cost is maintaining both a world-prediction path and an action path in the same runtime design. The benefit depends on the coupling. Joint generation encourages consistency between predicted futures and actions within a coupled generative process. Post-prediction heads keep the action expert relatively cheap once the substrate has been produced. Action-conditioned rollout supports counterfactual scoring, but repeated candidate evaluation can dominate inference. For joint members, the main optimization challenge is that substrate and action objectives can impose competing demands on the shared trunk. Practitioners therefore often pretrain the substrate path on video before introducing the action path, sometimes with a scheduled weight on the action loss in Equation 19.

4.4.5 LLM- and VLM-backbone WAMs

The fifth family builds a WAM directly on a large language or vision-language backbone without using a video-generation trunk as the main policy backbone. The natural substrate is the VLM-token variant of the feature category in Equation 14, since the future is described as a token block in the VLM’s vocabulary, and several members of the family reach across into the geometric or affordance categories of Section 4.2. The backbone is the VLM itself. The contract of Equation 8 is realized by composing a VLM forward pass with an optional action expert under the post-prediction-head factorization:

$$s_{t+1:t+H} \sim \text{VLM}_\theta(\text{prompt}(c)), \quad a_{t:t+H-1} \sim q_\psi(\cdot \mid s_{t+1:t+H}, c). \quad (32)$$

The factorization in Equation 32 realizes a WAM when a future substrate is part of the control pathway, either produced by the LLM/VLM itself or supplied by an attached world-model module before action decoding. A VLM policy that maps observations directly to action tokens, without predicting or consuming such a future substrate, falls outside this family under our definition. This subsection therefore covers LLM- and VLM-backbone designs in which future tokens, subgoal images, motion representations, affordance maps, or related latent transitions are present in the inference-time control context.

Among the π -series models, $\pi_{0.7}$ [72] satisfies this gate by conditioning the policy on generated subgoal images. At runtime, a lightweight BAGEL-initialized world model produces multi-view near-future visual subgoals, which are encoded together with observations by the Gemma3-based VLA before its 860M-parameter flow-matching action expert predicts the action chunk. CoT-VLA [201] writes an explicit visual chain-of-thought into the substrate before action is decoded. LDA-1B [112] learns dynamics, policy, and forecasting jointly in a structured DINO-based latent on a Qwen3-VL backbone. HiF-VLA [103] adopts a pretrained VLA trunk with DINOv2 and SigLIP encoders that produces both motion-vector substrate tokens and action latents in one inference pass. PALM [107] carries the family down to a GPT-2-class transformer over CLIP text, MAE vision, and MLP-encoded state, with a diffusion transformer attached as the action expert. ICLR-VR [124] adopts a Llama2-style causal transformer with a pretrained ViT image encoder and decodes a future-third-view gripper-keypoint polyline as the structured substrate before the action chunk fires. ALAM [149] pairs a PaliGemma-2B trunk with a Gemma-300M flow-matching expert, with an algebraic consistency regulariser living entirely in the frozen pretraining encoder. These examples keep the VLM family inside the WAM gate because the future substrate remains in the inference forward path.

The whole family inherits the cost of the VLM stack rather than that of video diffusion, which is why many members use action chunking and a frozen backbone at inference.

4.4.6 The five families together

The five families are not mutually exclusive and the dividing lines move every few months. Equations 21, 25, 28, 30, and 32 are the canonical parameterizations and they appear across substrates. A practitioner who fixes the substrate in Section 4.2 has narrowed the backbone choices by construction, since not every substrate

is compatible with every family. The remaining freedom is mostly about where the prediction lands relative to the control loop, which is the subject of Section 4.5.

4.5 Deployment Regime: Interactive, Rollout, Open Loop, Closed Loop

The final ingredient is the regime the WAM is meant to run in. It is the choice of horizon H in Equation 8 and of the cadence with which the WAM is invoked relative to the control loop. Let T be the task length in control steps, let $N_{\text{fwd}}(H)$ be the cost of one forward pass that produces a substrate trajectory of length H , and let f_{ctrl} be the control frequency. The four regimes below differ primarily in how they allocate N_{fwd} over the task of length T and how frequently new observations can revise the current plan.

4.5.1 Open-loop rollout

Open-loop rollout sets $H \approx T$ and invokes the WAM once. A long window of future substrate is generated once before execution and a separate actor consumes it. Total compute spent on the WAM is

$$C_{\text{open}}(T) = N_{\text{fwd}}(T), \quad (33)$$

which is paid once per scenario and does not scale with f_{ctrl} . UniPi [29], AVDC [81], Dreamitate [99], This&That [157], ARDuP [65], Gen2Act [9], RoboEnvision [184], LVP [17], and MVISTA-4D [160] run the WAM open-loop in their original setting, with the actor consuming the offline rollout through tracking or inverse dynamics. PointWorld [68] marks the open-loop limit of planning: MPPI evaluates chunked point-flow predictions over the horizon before execution, but the reported experiments do not replan during execution. Data-generation uses of Cosmos-Transfer1 [129] and Sora-class video models [130, 131] illustrate a related but non-census use of generated futures, where the future model supplies synthetic experience rather than acting as an inference-time WAM. The benefit of the regime is that the cost is paid once per scenario. Its main limitation is that the rollout cannot react to real-world deviations, so the actor that consumes it must absorb every distribution shift on its own.

4.5.2 Chunked closed-loop control

Chunked closed-loop sets $H = K$ for some replan period K and invokes the WAM every K control steps. The WAM produces an action chunk of length K , the actor executes the chunk in K/f_{ctrl} seconds, and the next invocation begins. Total compute over a task of length T is

$$C_{\text{chunk}}(T, K) = \lceil T/K \rceil N_{\text{fwd}}(K), \quad (34)$$

which interpolates between open-loop ($K \rightarrow T$) and single-step ($K \rightarrow 1$). It is feasible as long as $N_{\text{fwd}}(K)/K < 1/f_{\text{ctrl}}$, so that the next chunk is prepared before the previous one is consumed. This is a common choice for real-robot deployment because it amortizes a large backbone over multiple control ticks.

VLM-stack chunked controllers. In VLM-stack WAMs, chunked control appears when a future-substrate token block or structured future is refreshed before an action chunk is decoded. $\pi_{0.7}$ [72] produces a chunked action from its Gemma3 flow-matching expert after a separate BAGEL-class world model refreshes its multi-view subgoal images every four seconds, making it the π -series member that satisfies the inference-time gate. CoT-VLA [201] writes a visual chain-of-thought before each chunk is decoded. HiF-VLA [103] models motion-vector substrate tokens and action latents in parallel on a pretrained VLA trunk. PALM [107] and ICLR-VR [124] place a diffusion or autoregressive action expert on top of an affordance map or future-third-view gripper-keypoint polyline substrate. ALAM [149] uses algebraically consistent latent transitions on a PaliGemma-family trunk. FRAPPE [200], LDA-1B [112], DUST [171], FLARE [204], and GR-2 [16] also use chunked control, with the substrate supplied by a future embedding, a dual-stream diffusion head, or autoregression over frame and action tokens.

Pixel-latent and feature chunked controllers. WAMs that avoid decoded observation output use the same amortization principle, but differ in how much of the future substrate remains active at inference. UVA [93] and UWM [213] are the representative head-skipping cases, because the video side helps training while the inference call can produce only the action chunk. F1 [111], Motus [10], MotuBrain [121], BagelVLA [63], NoiseGate [67], Pelican-Unify 1.0 [198], τ_0 -WM [209], and WALL-WM [92] keep the same chunked regime inside MoT trunks.

Genie Envisioner [102], mimic-video [132], RynnVLA-002 [13], UD-VLA [19], and CKT-WAM [75] show adjacent ways to reduce per-call work, through a slow latent refresh, intermediate features, continuous or discrete joint autoregression, or a frozen student with updated context tokens. VideoVLA [140], CoVAR [183], CosmosPolicy [79], AdaWorldPolicy [191], DriveVA [106], DriveWAM [141], and DriveDreamer-Policy [212] keep chunked joint generation active in manipulation or driving settings.

Other chunked designs place the chunk behind a planning or abstraction stage rather than a bare action head. VLP [30] is an early decoded-video case, using value-guided search over abstract actions and replanning before goal-conditioned policies execute low-level controls. ActionImages [203], MV-VDP [90], CreFlow [125], and EVA [162] decode chunks from pixel-substrate plans. MoLA [96], VAMPO [42], LaMP [164], CLWM [27], DDP [64], and DAWN [110] do the same from latent or JEPa-style substrate plans. Among action-conditioned rollouts, 3D-ALP [142], DreamAvoid [33], and Feedback-WM [2] keep the execution window chunked, but differ in whether candidates come from tree search, critical-phase sampling, or feedback-guided denoising. EgoExo-WM [152] is the horizon-ranking boundary case, since it scores chunk-length body-motion candidates but reports open-loop rollout evaluation. JOPAT [49] and 3D-FDP [126] move the chunk to joint track or 3D-flow substrates, while Vidar [37], GR-MG [89], VILP [178], VideoPolicy [100], Act2Goal [208], Im2Flow2Act [177], TesserAct [202], and 3DFlowAction [207] show that the same chunked-cost equation can sit on pixel-grounded, feature, or geometric futures.

Compact or optional-substrate chunked controllers. Several chunked WAMs reduce the per-call burden by making the future sparse, latent, or optional. Veo-Act [199] and Say-Dream-Act [48] still decode from generated visual plans, while VAG [85], S-VAM [181], DiT4DiT [113], Fast-WAM [193], WAV [91], and DexWM [44] move the action path to feature or teacher-target futures. MWM [109] and AIM [32] replace appearance with task maps, and GigaWorld-Policy [186] and X-WAM [51] keep a joint video-action model but let the inference action path avoid the full visual schedule when possible.

Variants of chunked closed-loop control. Some WAMs keep Equation 34 but make the chunk schedule task-aware. FFDC-WAM [161] is the representative adaptive case, since a lightweight verifier compares predicted and observed latents and triggers an early replan when imagination drifts. NoiseGate [67] adapts denoising increments, WALL-WM [92] ties the chunk window to the next semantic event, and PALM [107] gates subtask progress inside the chunk. HarmoWAM [36] splits reactive and predictive loops, while SWEET [143] plus BageVLA [63] reduce the substrate per call to task-critical keyframes. These variants still pay the chunked cost, but they spend it when the forecast is likely to change the action.

The benefit of the chunked regime is that it spreads the cost of a large backbone over the chunk and tolerates moderate backbone latency. Its main limitation is the staleness of a long chunk, which becomes visible when the world changes during chunk execution.

4.5.3 Single-step closed-loop control

Single-step closed-loop sets $H = 1$ and invokes the WAM at every control step. The model produces one substrate prediction and one action per call, and per-step compute is

$$\mathcal{C}_{\text{single}}(T) = T N_{\text{fwd}}(1), \quad (35)$$

subject to the hard constraint $N_{\text{fwd}}(1) < 1/f_{\text{ctrl}}$. The iVideoGPT [174] antecedent illustrates the autoregressive version of this regime, while current census members such as WorldVLA [14], VidMan [170], GR-1 [172], PAD [52], VPP [62], DreamZero [188], PhysGen [144], and OmniVTA [206] sit closest to the single-step WAM pattern. Dex4D [83] runs one forward pass per control tick on the student backbone to output one action and one next-joint prediction, with the goal pointer advancing only when the point-to-point distance falls below a threshold. The latent-imagination Dreamer line [18, 57, 58] remains background for this regime when the agent acts at every imagination step. The regime gives the highest reactivity but pays the highest per-step inference cost. It is feasible mainly when the substrate has been chosen sparsely, in the feature, geometric, or affordance categories of Section 4.2, or in the pixel-decodable latent variant of the pixel-grounded category. In that case one forward pass can fit inside the control period.

4.5.4 Interactive simulator operation

The fourth regime is the truly interactive simulator, in which a stream of user or actor inputs continuously shapes the ongoing generation and there is no fixed endpoint. The WAM is invoked at every step but reuses the past via a key-value cache or a persistent latent of state size $\mathcal{M}(t)$. Per-step compute is

$$\mathcal{C}_{\text{int}}(t) = N_{\text{fwd}}^{\text{cached}}(\mathcal{M}(t)), \quad (36)$$

where $\mathcal{M}(t)$ grows with the horizon. For attention-based backbones, the cached forward pass is linear in $\mathcal{M}(t)$ at each step and cumulative cost is quadratic in t . InteractiveWorldSimulator [166], EnerVerse [66], and LingBot-VA [88] approach this regime among the cited interactive works, while Dreamer 4 [58] remains a background lineage point for persistent latent imagination. The latent-imagination lineage from PlaNet [55] and DreamerV3 [57] through TransDreamer [18] also fits when the actor and the world model run inside a shared rollout loop. The defining property is that the model can be steered mid-stream and the same cache or persistent state carries the long-horizon trajectory. The cost pressure lands on memory rather than on per-step compute. Section 5 returns to that pressure under the heading of persistence.

4.6 Putting the Four Ingredients Together

Sections 4.1 through 4.5 let us pin down any existing WAM by a compact 4-tuple

$$\text{WAM} \cong (\Phi, \mathcal{F}, \mathcal{B}, \mathcal{D}), \quad (37)$$

whose four coordinates take values matching the table columns below. The substrate Φ from Section 4.2 takes one of {pixel-grounded, feature, geometric, affordance}. The table cells record narrower variants where they matter, with decoded or latent for pixel-grounded entries, audio-latent for the acoustic observation entry, and encoder-only, tap, teacher, or VLM token for feature. Here *Pixel (latent)* abbreviates the pixel-decodable latent route above. The action coupling \mathcal{F} from Section 4.3 takes one of {action-conditioned rollout, joint generation, post-prediction head}. The backbone \mathcal{B} from Section 4.4 takes one of {diffusion, autoregressive, joint-embedding, hybrid, LLM/VLM}. The deployment regime \mathcal{D} from Section 4.5 takes one of {open, chunked, single, interactive}.

Example placements show how the tuple is used. F1 [111] is (pixel-decodable latent grid, joint generation, hybrid, chunked). WorldVLA [14] is (pixel-decodable latent grid, joint generation, autoregressive, single). FLARE [204] is (latent teacher-target, post-prediction head, joint-embedding, chunked). UWM [213] is (pixel-decodable latent grid, joint generation, hybrid, chunked). Frame-level video-diffusion works such as Wan [156] and CogVideoX [185] are *not* WAMs on their own because they fix only Φ and \mathcal{B} . They become a WAM only once \mathcal{F} and \mathcal{D} are added, which is exactly the wrapper-around-a-frozen-backbone pattern that runs through CosmosPolicy [79], VidMan [170], and VideoVLA [140] on top of related backbones.

The 4-tuple also clarifies the current direction of WAM design. The first and third coordinates, Φ and \mathcal{B} , receive substantial attention because they are training-time choices and because choosing a pretrained backbone is often the expensive design decision. The middle coordinate \mathcal{F} and the closing coordinate \mathcal{D} are more flexible at inference time, and this flexibility is where many recent mechanisms appear. The census reinforces this asymmetry. Substrate choices remain inside the four categories and backbone choices cluster around diffusion and hybrid families, while sharper variation appears in coupling mechanisms such as iterative-interactive refinement, layer-wise integration, and asynchronous-adaptive denoising, and in runtime mechanisms such as adaptive- K chunking, event-triggered invocation, and dual-frequency slow-fast loops. The expensive corner of the design space combines a decoded pixel-grounded substrate, joint generation, a diffusion backbone, and single-step closed-loop deployment. Few methods can use that corner directly in production. Designs that appear to use it usually trade one coordinate for simulation, an offline rollout, or a smaller model used at inference. Section 5 examines the same methods through the five properties that embodiment demands of them. Section 6 examines the data and the evaluation practices the field has settled on to make these design choices empirically meaningful.

Tables 1 and 2 place every WAM in our census on the 4-tuple of Equation 37. Table 1 gathers the works in the Render-and-Decode and Latent-Only philosophies of Section 3, while Table 2 gathers the Video-Generation-Free works that drop the video-generation backbone from the predictive path. Within each band, rows follow

arXiv first-submission date so that the family’s evolution unfolds top to bottom. The substrate column reports where the WAM represents its future, following the four categories of Section 4.2. The backbone column refers to the families of Section 4.4. The action coupling column refers to the three families of Equations 10, 11, and 12, with *Cond. rollout* used as the compact table label for action-conditioned rollout. The deployment column refers to the regimes of Section 4.5.

Table 1 Census of World Action Models in the Render-and-Decode and Latent-Only philosophies of Section 3. The upper band lists Render-and-Decode methods and the lower band lists Latent-Only methods. Columns give the four-axis anatomy of Section 4, locating each work on the design tuple of Equation 37, and within each band rows follow arXiv first-submission date.

WAM	Date	Substrate	Backbone	Action coupling	Deployment
Render-and-Decode					
UniPi [29]	2023.02	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
AVDC [81]	2023.10	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
VLP [30]	2023.10	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
GR-1 [172]	2023.12	Pixel (latent)	Autoregressive	Joint generation	Single-step
Dreamitate [99]	2024.06	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
This&That [157]	2024.07	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
GR-MG [89]	2024.08	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
Gen2Act [9]	2024.09	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
GR-2 [16]	2024.10	Pixel (latent)	Autoregressive	Joint generation	Chunked
PAD [52]	2024.11	Pixel (latent)	Diffusion	Joint generation	Single-step
CoT-VLA [201]	2025.03	Feature (VLM token)	LLM/VLM	Post-prediction head	Chunked
Tesseract [202]	2025.04	Geometric	Diffusion	Post-prediction head	Chunked
DreamGen [73]	2025.05	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
WorldVLA [14]	2025.06	Pixel (latent)	Autoregressive	Joint generation	Single-step
RoboEnvision [184]	2025.06	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
4DGen [108]	2025.07	Pixel (decoded) \wedge Geometric	Diffusion	Post-prediction head	Open-loop
RIGVid [133]	2025.07	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
Vidar [37]	2025.07	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
F1 [111]	2025.09	Pixel (latent)	Hybrid	Joint generation	Chunked
NovaFlow [87]	2025.10	Geometric	Diffusion	Post-prediction head	Open-loop
UD-VLA [19]	2025.11	Pixel (latent)	Diffusion	Joint generation	Chunked
RynnVLA-002 [13]	2025.11	Pixel (latent)	Hybrid	Joint generation	Chunked
VideoVLA [140]	2025.12	Pixel (latent)	Diffusion	Joint generation	Chunked
Motus [10]	2025.12	Pixel (latent)	Hybrid	Joint generation	Chunked
LVP [17]	2025.12	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
CoVAR [183]	2025.12	Pixel (latent)	Diffusion	Joint generation	Chunked
Dream2Flow [28]	2025.12	Geometric	Diffusion	Post-prediction head	Open-loop
TC-IDM [119]	2026.01	Geometric	Diffusion	Post-prediction head	Open-loop
BagelVLA [63]	2026.02	Pixel (latent)	Hybrid	Joint generation	Chunked
MVISTA-4D [160]	2026.02	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
Say-Dream-Act [48]	2026.02	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
Dex4D [83]	2026.02	Geometric	Hybrid	Post-prediction head	Single-step
DreamZero [188]	2026.02	Pixel (latent)	AR / Diffusion	Joint generation	Single-step
NovaPlan [38]	2026.02	Geometric	Hybrid	Post-prediction head	Chunked
PhysGen [144]	2026.03	Pixel (latent)	Autoregressive	Joint generation	Single-step
EmboAlign [195]	2026.03	Pixel (decoded)	Diffusion	Post-prediction head	Open-loop
AeroPlace-Flow [120]	2026.03	Geometric	Diffusion	Post-prediction head	Open-loop
EVA [162]	2026.03	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
DriveDreamer-Policy [212]	2026.04	Pixel (decoded) \wedge Geometric	Hybrid	Joint generation	Chunked
MV-VDP [90]	2026.04	Pixel (decoded) \wedge Affordance	Diffusion	Joint generation	Chunked
DriveVA [106]	2026.04	Pixel (latent)	Diffusion	Joint generation	Chunked
Veo-Act [199]	2026.04	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
ActionImages [203]	2026.04	Pixel (decoded) \wedge Affordance	Diffusion	Joint generation	Chunked
GraspDreamer [147]	2026.04	Pixel (decoded) \wedge Geometric	Diffusion	Post-prediction head	Open-loop
3D-ALP [142]	2026.04	Pixel (decoded)	Hybrid	Cond. rollout	Chunked
VAG [85]	2026.04	Feature (tap)	Diffusion	Post-prediction head	Chunked
$\pi_{0.7}$ [72]	2026.04	Pixel (decoded)	Hybrid	Post-prediction head	Chunked
X-WAM [51]	2026.04	Pixel (latent)	Hybrid	Joint generation	Chunked
CKT-WAM [75]	2026.05	Pixel (latent)	Diffusion	Joint generation	Chunked

continued on next page

Table 1 Census of World Action Models in the Render-and-Decode and Latent-Only philosophies (continued).

WAM	Date	Substrate	Backbone	Action coupling	Deployment
NoiseGate [67]	2026.05	Pixel (latent)	Hybrid	Joint generation	Chunked
HarmoWAM [36]	2026.05	Pixel (decoded / latent)	Hybrid	Post-prediction head	Chunked
DreamAvoid [33]	2026.05	Pixel (decoded)	Diffusion	Cond. rollout	Chunked
MoLA [96]	2026.05	Pixel (latent)	Diffusion	Post-prediction head	Chunked
CreFlow [125]	2026.05	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
Pelican-Unify 1.0 [198]	2026.05	Pixel (latent)	Hybrid	Joint generation	Chunked
SWEET [143]	2026.05	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
DriveWAM [141]	2026.05	Pixel (latent)	Hybrid	Joint generation	Chunked
VERA [94]	2026.05	Pixel (decoded)	Diffusion	Post-prediction head	Chunked
Latent-Only					
ARDuP [65]	2024.06	Pixel (latent)	Diffusion	Post-prediction head	Open-loop
Im2Flow2Act [177]	2024.07	Geometric	Diffusion	Post-prediction head	Chunked
VPP [62]	2024.12	Feature (tap)	Diffusion	Post-prediction head	Single-step
VILP [178]	2025.02	Pixel (latent)	Diffusion	Post-prediction head	Chunked
UVA [93]	2025.03	Pixel (latent)	Hybrid	Joint generation	Chunked
UWM [213]	2025.04	Pixel (latent)	Hybrid	Joint generation	Chunked
3DFlowAction [207]	2025.06	Geometric	Diffusion	Post-prediction head	Chunked
Video Policy [100]	2025.08	Feature (tap)	Diffusion	Post-prediction head	Chunked
Genie Envisioner [102]	2025.08	Feature (tap)	Diffusion	Post-prediction head	Chunked
3D-FDP [126]	2025.09	Geometric	Diffusion	Post-prediction head	Chunked
TraceGen [86]	2025.11	Geometric	Diffusion	Post-prediction head	Chunked
mimic-video [132]	2025.12	Feature (tap)	Diffusion	Post-prediction head	Chunked
Act2Goal [208]	2025.12	Feature (tap)	Diffusion	Post-prediction head	Chunked
CosmosPolicy [79]	2026.01	Pixel (latent)	Diffusion	Joint generation	Chunked
LingBot-VA [88]	2026.01	Pixel (latent)	Hybrid	Post-prediction head	Interactive
GigaBrain-0.5M* [158]	2026.02	Pixel (latent)	Hybrid	Post-prediction head	Chunked
AdaWorldPolicy [191]	2026.02	Pixel (latent)	Hybrid	Joint generation	Chunked
3PoinTr [70]	2026.03	Geometric	Hybrid	Post-prediction head	Chunked
DiT4DiT [113]	2026.03	Pixel (latent)	Diffusion	Post-prediction head	Chunked
S-VAM [181]	2026.03	Feature (tap)	Diffusion	Post-prediction head	Chunked
Fast-WAM [193]	2026.03	Feature (encoder-only)	Diffusion	Post-prediction head	Chunked
GigaWorld-Policy [186]	2026.03	Pixel (latent)	Diffusion	Joint generation	Chunked
OmniVTA [206]	2026.03	Pixel (latent)	Diffusion	Post-prediction head	Single-step
VAMPO [42]	2026.03	Feature (tap)	Diffusion	Post-prediction head	Chunked
VTAM [192]	2026.03	Pixel (latent)	Hybrid	Joint generation	Chunked
LaMP [164]	2026.03	Feature (tap)	Hybrid	Post-prediction head	Chunked
AIM [32]	2026.04	Affordance	Diffusion	Joint generation	Chunked
WAV [91]	2026.04	Feature (tap)	Diffusion	Post-prediction head	Chunked
CLWM [27]	2026.04	Feature (teacher)	Hybrid	Joint generation	Chunked
MWM [109]	2026.04	Affordance	Diffusion	Post-prediction head	Chunked
MotuBrain [121]	2026.04	Pixel (latent)	Hybrid	Joint generation	Chunked
FFDC-WAM [161]	2026.05	Pixel (latent)	Hybrid	Joint generation	Chunked
DAWN [110]	2026.05	Feature (teacher)	Hybrid	Joint generation	Chunked
EgoExo-WM [152]	2026.05	Feature (teacher)	Hybrid	Cond. rollout	Open / Chunked
RoboFlow4D [104]	2026.05	Geometric	Diffusion	Post-prediction head	Chunked
JOPAT [49]	2026.05	Pixel (latent) \wedge Geometric	Diffusion	Joint generation	Chunked
τ_0 -WM [209]	2026.06	Pixel (latent)	Hybrid	Joint generation	Chunked
WALL-WM [92]	2026.06	Pixel (latent)	Hybrid	Joint generation	Chunked

Table 2 Census of Video-Generation-Free World Action Models, the third design philosophy of Section 3. No video-generation backbone is in the predictive path. Columns match Table 1, locating each work on the design tuple of Equation 37, and rows follow arXiv first-submission date.

WAM	Date	Substrate	Backbone	Action coupling	Deployment
Video-Generation-Free					
FLARE [204]	2025.05	Feature (teacher)	Joint-embedding	Post-prediction head	Chunked

continued on next page

Table 2 Census of Video-Generation-Free World Action Models (continued).

WAM	Date	Substrate	Backbone	Action coupling	Deployment
DUST [171]	2025.10	Feature (VLM token)	Hybrid	Joint generation	Chunked
Audio-WM [194]	2025.12	Audio (latent)	Diffusion	Post-prediction head	Chunked
HiF-VLA [103]	2025.12	Geometric	LLM/VLM	Joint generation	Chunked
DexWM [44]	2025.12	Feature (teacher)	Hybrid	Post-prediction head	Chunked
PointWorld [68]	2026.01	Geometric	Hybrid	Cond. rollout	Open / Chunked
PALM [107]	2026.01	Affordance	LLM/VLM	Post-prediction head	Chunked
LDA-1B [112]	2026.02	Feature (teacher)	LLM/VLM	Post-prediction head	Chunked
FRAPPE [200]	2026.02	Feature (teacher)	Hybrid	Post-prediction head	Chunked
ICLR-VR [124]	2026.03	Geometric	LLM/VLM	Joint generation	Chunked
DDP [64]	2026.03	Feature (teacher)	Hybrid	Joint generation	Chunked
ALAM [149]	2026.05	Feature (VLM token)	LLM/VLM	Joint generation	Chunked
Feedback-WM [2]	2026.05	Feature (teacher)	Hybrid	Cond. rollout	Chunked

5 Core Properties of World Action Models

The four-axis anatomy of Section 4 explains how a WAM is assembled. This section asks a different question. Once that assembly is placed inside a control loop and made to run, what must remain true of it? We focus on five properties. A WAM must be **interactable**, so that control signals can shape the predicted future rather than only be decoded after generation. It must be **causal**, so that future information cannot leak into the action being executed. It must be **persistent**, so that the predicted state remains coherent as the robot acts, observes, and replans. It must be **physically plausible**, so that the future is realizable by the embodiment rather than only visually convincing. It must be **generalizable**, so that the same predictive-action contract remains useful when tasks, objects, scenes, cameras, or embodiments change. These properties are not independent. A design that improves one property often shifts cost, latency, memory, or error into another, which is why WAM design is best understood as a set of trade-offs among prediction quality, runtime feasibility, and control bandwidth.

5.1 Interactability

A video world model is only weakly interactable by default. Prompted to show a cup being lifted, it will render a plausible lift, but nothing ties that imagined motion to the trajectory the controller actually intends to send. The action is not yet the variable that governs the continuation. A WAM becomes interactable by deciding where the action enters the prediction loop. The factorization families in Equations 10, 11, and 12 lay out this design space. This subsection walks from post-prediction decoding, where action is recovered only after the future has been produced, through action-conditioned rollout, where proposed actions shape candidate futures, to joint prediction, where action and future are produced together.

Post-prediction decoding. In this form, the model synthesizes a whole future first and only then recovers control from it. UniPi, AVDC, and VLP establish this form by generating visual plans before extracting control [29, 30, 81]. Dreamitate, Gen2Act, and RoboEnvision extend the same post-prediction interface to tool use, human video, and long-horizon keyframes [9, 99, 184]. This end of the spectrum is clearest in cascaded WAMs: NovaFlow, Dream2Flow, and TC-IDM keep the future predictor separate, then convert generated motion into object flow, tool trajectories, or inverse-dynamics controls [28, 87, 119]. Audio-WM shows the same pattern outside vision, where generated future sound becomes the policy condition [194]. This interface is easy to attach to a pretrained video generator, but controllability arrives only after the full generation cost has been paid.

In-generation control. A more integrated interface pushes control into the generative process itself, and the works here line up as a near-monotone progression in how early the action binds. This&That starts the progression by conditioning video diffusion on spatial language and paired gesture coordinates, with ablations showing the gesture channel is essential for precise manipulation [157]. ARDuP narrows that conditioning to automatically discovered active regions and decodes action from latent representations rather than rendered pixels [65]. More integrated still, CoVAR and VAG run a dedicated action branch alongside video denoising and couple the two through bridge attention or step-synchronized denoising, passing only compact pooled

video context to the action side [85, 183]. UWM, UVA, F1, Motus, and Cosmos Policy drop the separate branch and share substrate prediction and action heads on a single backbone [10, 79, 93, 111, 213]. AdaWorld binds earliest of all, injecting learned latent actions at every denoising step [40, 191]. Taken in order, the progression is a dial. The later the action binds, the cheaper and more modular the model, and the earlier it binds, the more the predicted future is genuinely shaped by the act it must inform.

Bottlenecked action pathways. The most bottlenecked designs restrict the information that the action pathway can consume. AIM exposes the future through a predicted spatial value map rather than raw future-pixel tokens [32]. GigaWorld-Policy makes action the primary prediction target and treats video as an optional action-conditioned by-product [186]. From the subgoal side, $\pi_{0.7}$ supplies sparse multi-view goal images to the policy context rather than a dense rollout [72]. Narrowing the channel this way is exactly what buys cheap, controllable inference, and it is also what caps the approach. The action can react only to whatever future variables survive the bottleneck.

5.2 Causality

Causality governs both correctness and latency. A standard video diffusion model denoises an entire clip jointly, so future frames can influence the representation that later drives the current action. This bidirectional pattern improves temporal smoothness, but it prevents the model from acting before imagination has finished. Worse, a model that quietly attends to the next frame can grasp flawlessly inside its own rollout and then close on empty air at inference, because the cue it leaned on was never available in real time. A WAM must instead produce predictions that depend only on past observations and the chosen action, and it must expose the relevant part of that prediction quickly enough for control.

Causal token streams. Causal streaming makes the ordering explicit: future and action are decoded as a token stream, and past computation is reused through a key-value cache [14, 16, 144, 172, 174]. PhysGen makes the mechanism explicit by combining causal masking, lookahead multi-token prediction, and key-value caching. It predicts several future action tokens at each step, while execution consumes only the leading token [144]. The gain comes from avoiding repeated computation over context that has already been fixed.

Leakage-free denoising. When prediction remains chunked, the key constraint is leakage control inside the chunk. Action tokens are forbidden from attending to future-pixel tokens or later actions, so the action head cannot benefit from information that would be unavailable at inference time [14, 19, 201]. WorldVLA introduced this mask to control error propagation in autoregressive action chunks, and the same mask also removes expensive attention between action tokens and future-pixel tokens [14]. CoT-VLA uses causal attention to generate a discrete visual chain-of-thought, then switches to full attention only for the short action sequence that reaches the goal [201].

Action-prioritized inference. Once leakage is controlled, latency is what is left to fight. Intermediate ODE checkpoints, single forward passes, and early denoising exits work because the controller usually needs the next action more than it needs a fully refined future video [19, 113, 132, 181, 193]. UD-VLA formalizes this point with synchronous joint discrete denoising and reports roughly fourfold faster inference than autoregressive counterparts [19]. The saved computation is mainly the tail of the denoising trajectory, where distant pixels are refined but the next action has already become stable.

The same prioritization also overlaps prediction with execution. LingBot-VA, MotuBrain, and DreamZero hide part of world-model latency behind the execution of the previous action chunk [88, 121, 188]. LingBot-VA pairs a Mixture-of-Transformers backbone, which gives shared and modality-specific tokens their own expert streams, with an asynchronous pipeline that runs action prediction and motor execution at the same time [88]. DreamZero pushes this overlap furthest, sustaining closed-loop control from a large autoregressive video-diffusion model through the cache-level re-grounding we examine under persistence in Section 5.3 [188]. Diffusion forcing supplies a related middle point: it assigns low noise to the near-certain past and high noise to the uncertain future under causal masking, so diffusion quality is retained where uncertainty remains [17]. Causal structure thus does two jobs at once. It blocks the future from leaking into the current action, and it decides where computation is worth spending, on the part of the rollout that can still move the next decision.

5.3 Persistence

Persistence is coherence under repeated action, observation, and replanning. It fails through drift, cost growth, and forgetting. Drift appears because each predicted step conditions the next, so small errors compound until the rollout leaves the data manifold. Cost grows when a model attends to full history, since memory and attention increase with the episode horizon. Forgetting appears when a finite context discards scene identity, so an occluded object can reappear in the wrong place. These failures pull against each other. Full-history attention fights drift and forgetting but raises cost. Sliding context bounds cost but discards old evidence.

Observation replacement. The cheapest persistence mechanism is re-grounding. After each executed chunk, the method replaces imagined observations with measured observations and updates the cache before forecasting again. DreamZero applies this idea through key-value-cache observation replacement and reaches 7 Hz closed-loop control with a 14B autoregressive video-diffusion model [188]. LingBot-VA follows the same principle with a closed-loop rollout that re-grounds on feedback before predicting the next chunk [88]. Re-grounding is cheap because it uses the robot’s new observation rather than generating a longer imagined future.

Bounded memory. Re-grounding limits drift, but memory still needs a fixed budget. DexWM predicts future latent states from finger-keypoint actions and maintains bounded test-time memory instead of an unbounded cache, which lets it train and operate over 900 hours of human and robot interaction [44]. Act2Goal uses multi-scale temporal hashing to keep near-future frames dense while sparsifying distal frames, preserving long-horizon goal consistency at bounded cost [208]. LingBot-VA and MotuBrain add per-step cost control through Mixture-of-Transformers streams, which cap the cost of processing history even when a longer context is available [88, 121].

5.4 Physical Plausibility

The video-world-model instinct equates a good future with a photorealistic one. A WAM needs something stricter. Its future exists only to constrain action, so what matters is whether the predicted dynamics are realizable by the body that will execute them, not whether they look convincing. Contact, force, and kinematics are what decide that, and three kinds of physical constraint push prediction toward it.

Visual-geometric abstraction ladder. Prediction targets form a ladder, and a WAM can choose how far down it to predict. At the top sits full RGB. One rung down, RGB-D, surface normals, and 4D point maps add explicit geometry to generated video, as in WAMs that reconstruct structure from RGB-D-Normal rollouts [51, 160, 202]. Below that, optical flow and 3D object flow keep only motion, the variable that transfers most cleanly across humans, robots, simulation, and reality. Im2Flow2Act and 3DFlowAction use flow to bridge domains and embodiments [177, 207], and FlowDreamer carries the same abstraction into flow-conditioned generation [50]. Lower still, semantic masks discard photometric detail and keep only what the policy acts on. MWM makes the case for this rung directly, arguing that RGB prediction is misaligned with control because masks preserve the geometry while discarding the photometric nuisance [109].

At the bottom, the future stops being rendered at all. VPP and S-VAM decode action from predicted or distilled foresight features, S-VAM foreseeing geometric and semantic representations in a single forward pass rather than a video [62, 181]. GigaWorld-Policy, Fast-WAM, WAV, and DexWM keep the future latent or action-centered at inference [44, 91, 186, 193], and FLARE and LDA-1B push the target into frozen or structured embedding spaces [112, 204]. The descent obeys a single logic. Each rung down trades away appearance the controller never needed and forces capacity toward geometry, contact, and dynamics, at the price of a future that is harder to inspect.

Richer physical modalities. Geometry narrows the target, but vision alone does not expose the variables that decide contact-rich manipulation, such as contact state, friction, slip, and normal force. A WAM that predicts only pixels is therefore fragile in the regime where manipulation is hardest. OmniVTA predicts future tactile contact states and pairs a slow policy with a high-rate reflexive controller that corrects discrepancies between predicted and observed tactile signals [206]. AdaWorldPolicy adds force prediction as a dedicated branch alongside the world model and the action expert, then uses force-torque mismatch for online adaptation [191]. DexWM predicts finger-level interaction rather than coarse end-effector motion [44], the same fine-grained bet that buys it bounded memory under persistence (Section 5.3). These modalities are low-dimensional, but

they carry dense information exactly where contact determines success. They often offer a cheaper and more action-relevant target than the pixels of the same event.

Proprioceptive and kinematic-chain coherence. Contact signals still need to be grounded in the body that executes the action. A predicted future can look plausible and still be impossible for a particular robot. The arm may pass through itself, the end-effector may be unreachable, or the hand configuration may violate joint limits. Physical plausibility therefore also requires embodiment coherence. PAD, GR-1, Cosmos Policy, and X-WAM condition prediction on proprioception so that the forecast is tied to joint state [51, 52, 79, 172]. Cosmos Policy goes further by treating future proprioception and image observations as jointly modeled latent frames [79]. DexWM adds a hand-consistency loss that penalizes kinematically inconsistent predictions [44]. MVISTA-4D uses trajectory-level inference with a residual inverse model to project an imagined future onto an action the embodiment can execute [160]. Each design shrinks the effective search space to the robot’s feasible manifold.

These constraints all draw on physical priors learned from internet-scale video pretraining [114, 128, 130, 131, 156, 185]. The main lesson is that how useful the imagined future is for the action predicts task success better than how good it looks [91, 140, 186, 193]. A WAM should therefore predict the coarsest embodiment-grounded representation that still constrains the action, and flow, masks, tactile and force signals, and proprioceptively consistent latents often fill that role better than photorealism.

5.5 Generalization

Generalization in robotics is a bundle of shifts rather than a single axis: new tasks, objects, appearances, rooms, cameras, embodiments, and action spaces. Under the WAM contract, generalization requires both sides of the model to survive the shift. The predictive substrate must remain useful, and the action pathway must still execute the prediction. This is stricter than VLA-style semantic generalization and stricter than video generalization, because an action-relevant future is required in addition to a plausible instruction-to-action mapping or a plausible video.

Video-prior transfer. The broadest transfer source is the video prior itself. GR-1 and GR-2 show that large-scale language-conditioned video pretraining improves manipulation in unseen scenes and novel scenarios once action prediction is fine-tuned [16, 172]. VideoVLA ties the prior directly to action by jointly denoising future frame latents and action tokens, improving novel-object and cross-embodiment performance while retaining the cost of multi-step diffusion sampling [140]. DreamZero makes a similar argument with an autoregressive video-action model and reports unseen-motion gains over VLA baselines, together with robot-to-robot and human-to-robot transfer from short video-only adaptation sets [188]. LVP keeps the connection cascaded: it trains a large video planner on human and robot clips, then extracts executable motion through pose estimation and retargeting for zero-shot real-robot tasks [17]. RoboEnvision targets task transfer by turning instructions into future keyframes that condition the downstream policy, so unseen tasks can still be routed through a visual subgoal interface [184]. $\pi_{0.7}$ uses generated futures as subgoal context rather than as the online action generator [72]. The common thread is that scale helps only once the future is connected to the action, and making that connection usually costs an inverse model, a subgoal generator, or a large video backbone.

Substrate transfer. When pixels carry too much appearance-specific detail, transfer shifts to the substrate. Flow-based methods treat object motion as the transferable variable. Im2Flow2Act uses object flow to bridge human-to-robot and simulation-to-real gaps before a flow-conditioned policy acts [177]. 3DFlowAction lifts the same idea into 3D object flow, separating object dynamics from the robot that later executes the trajectory [207]. What flow buys is suppression of texture, camera style, and part of the embodiment gap, but it leaves the hard part to the executor, which still has to solve grasping, collision avoidance, and contact.

Mask, latent, and feature substrates make the same trade-off at a deeper level. MWM replaces RGB prediction with semantic-mask dynamics, improving robustness to background, lighting, and object-color shifts [109]. S-VAM distills a diffusion backbone into geometric and semantic foresight features, so inference no longer depends on full video denoising [181]. Fast-WAM and GigaWorld-Policy keep video prediction as a training signal but remove future-video generation from part or all of inference [186, 193]. The deepest version of the move forecasts in feature space rather than pixels. FRAPPE and LDA-1B predict in self-supervised vision-

foundation features, where appearance variation washes out and what survives the shift is the action-relevant structure [112, 200]. These substrates make generalization less dependent on photorealistic rendering, but they also make the predicted future harder to inspect and harder to evaluate with video metrics.

Action-abstraction transfer. Substrate transfer still needs an executor, so another line of work transfers the action abstraction and the data recipe. LDA-1B scales this separation with universal embodied data ingestion, a hand-centric action space, and DINO latent forecasting across high-quality demonstrations, low-quality trajectories, and actionless videos [112]. DUST shows the diffusion-side version of the point: action-free BridgeV2 video pretraining improves a dual-stream policy even when future observations are predicted only in an embedding space [171]. ALAM adds an algebraic consistency constraint over latent action transitions, so the abstract action space keeps a predictable composition law while the action expert remains tied to the embodiment [149]. The shared move is to separate what should transfer from what should stay local. The substrate carries task and environment dynamics. The action decoder absorbs embodiment-specific kinematics. The data recipe supplies enough variation that neither side memorizes the training platform.

What falls out in practice is a single principle. Predict at the invariant level that still constrains control. Pixels transfer semantic and physical priors but carry appearance cost. Flow and masks transfer object motion and geometry but require a reliable executor. Latent actions and feature forecasts scale to action-free video and heterogeneous data, but they need careful grounding before they become precise control. No design dominates across all shifts. A WAM that claims generalization should therefore state which shift it targets, which design axis is meant to transfer, and which adapter remains specific to the robot that executes the action. This framing turns generalization from a broad claim into a checkable property.

Across the five properties, one lesson outweighs the rest. They are not a checklist but a set of competing pressures. The stronger action interface that makes a model interactable also narrows the channel that causality must keep leak-free. The bounded memory that buys persistence is the same budget that starves long-horizon plausibility. The abstraction that lets a substrate transfer is what makes its predicted future hard to evaluate. A WAM is therefore never tuned for one property alone. Every design sits at a point where a gain in one property is paid for in another, and turning that balance into something measurable, rather than asserted, is the task we take up in Section 6.

6 Data and Evaluation

A WAM can only preserve the properties of Section 5 when its training data and evaluation protocol expose those properties. Data determines which future variables the model can learn, which action labels it can trust, and which embodiments it can transfer across. Evaluation determines whether those learned variables are useful in a control loop rather than only plausible as video. The two questions are therefore coupled. This section first organizes WAM data by the trade-off between scale, action-label quality, and embodiment match. It then organizes evaluation by the trade-off between visual fidelity, closed-loop success, physical plausibility, and evaluation cost.

6.1 Data sources

WAM training data falls into five groups. Each group provides a different compromise between scale, action-label fidelity, physical grounding, and access.

Robot teleoperation. Robot teleoperation gives the cleanest action-conditioned trajectories because a human operator drives the embodiment and the command stream is recorded with the visual observation. Open X-Embodiment [21] aggregates this recipe across laboratories and robot platforms, making it the canonical cross-embodiment resource for robot policy training. RoboMIND [175], RoboSet [8], RoboTurk [115], and Robo360 [101] extend the same supervision pattern with different platform mixes, task suites, and camera configurations. Human2Robot [176] pairs each teleoperated trajectory with a synchronized video of a bare human hand performing the same task, adding a human-to-robot correspondence on top of the exact action labels. RoboNet [24] departs from human teleoperation by collecting real-robot trajectories through scripted and random policies, trading demonstration precision for throughput while still logging exact action commands.

The benefit of this group is label fidelity. The cost is that every additional hour consumes robot time, operator time, or both.

Portable human demonstrations. Portable demonstrations increase throughput by moving data collection from a robot platform to a human-worn or handheld capture setup. EgoMimic [77] records egocentric video and three-dimensional hand tracks with a lightweight wearable rig. EgoVerse [135] scales the same idea with wearable and phone-based capture, then retargets the demonstrations across robot embodiments. EgoDex [60] carries the head-worn route to large scale, recording dexterous bimanual hand-pose tracks on Apple Vision Pro with no robot in the loop. This group buys scale because human demonstrations are faster to collect than robot teleoperation. It pays for that scale with an embodiment gap that must be closed before the data can train a robot policy.

Internet-scale egocentric and instructional video. Internet video gives WAMs their scale, but it usually lacks robot actions. Ego4D [45], Ego-Exo4D [46], and EgoPAT3D [98] emphasize first-person activity and hand motion. EPIC-KITCHENS [23], HD-EPIC [134], EgoVid-5M [163], OpenEgo [74], IndEgo [15], EgoScale [205], and Egocentric-10k [1] broaden the pool with instructional, household, and large-scale egocentric video. This group is valuable because it exposes broad visual variation, object dynamics, and human manipulation patterns. Its limitation is the missing action channel. One route uses such video, together with broader internet video, to pretrain a video backbone, as in V-JEPA 2 [4] and Wan [156]. Another route attaches an inverse-dynamics or latent-action model that recovers a control proxy from video. AVDC [81] recovers flow-conditioned action from unlabeled video, while LDA-1B [112] and DUST [171] use action-free video to strengthen latent future prediction before action decoding. Internet video is therefore a scale source only after the action-label problem has been made explicit.

Simulation. Simulation provides exact action labels, controlled curricula, and low marginal cost after the environment has been authored. Classical physics platforms for manipulation include Robosuite [215], ManiSkill 2 [47], ManiSkill 3 [150], MetaWorld [190], and LIBERO [105]. General simulation backends such as IsaacSim [127], CoppeliaSim [137], and SimplerEnv [95] provide the execution substrate for broader embodied tasks. Household and digital-twin datasets build on these engines to generate demonstrations programmatically, including RoboCasa [123], RoboTwin [122], and RoboTwin 2 [20]. RoboCerebra [59] instead collects human-teleoperated trajectories inside the simulator for long-horizon evaluation. Aggregators sit above both lines: RoboVerse [43] unifies several simulator backends, and RoboData [180] merges simulation benchmarks with a real-robot set. The benefit is scalable label quality. The limitation is the sim-to-real gap, which a WAM often tries to close with an internet-pretrained visual prior or a more abstract substrate.

Synthetic data from WAMs themselves. A WAM can also become a data engine. IRASim [214] follows this pattern from the simulator side by replacing a physics engine with a learned action-conditioned diffusion model whose rollouts stand in for simulated trajectories. Cosmos-Transfer1 [129] and InteractiveWorldSimulator [166] support similar data-generation use, and Sora-class video models are sometimes used to seed downstream policies [130, 131]. DreamGen [73] makes the cascaded WAM version concrete by generating robot videos, recovering pseudo-actions or latent actions, and training the final policy from those neural trajectories. RIGVid [133] and GraspDreamer [147] use generated demonstrations in a more direct action-recovery path, where tracking and retargeting replace manual action labels. Synthetic neural trajectories sit between simulation and real teleoperation. They inherit the realism of internet pretraining, but they also inherit the generator’s failure modes.

In practice, WAM training mixes these sources. Internet video supplies visual and physical priors, teleoperation supplies trusted action labels, portable human data adds scale, simulation supplies controllable coverage, and synthetic trajectories fill gaps that are expensive to collect physically. The unresolved question is how to choose the mixture for a target deployment regime. Access is a second constraint. Sora-class generators [130, 131] and the EgoScale release [205] are closed or not yet public, which limits independent verification of their contribution.

6.2 Evaluation

Evaluation must judge both prediction and action. Video-generation metrics ask whether the future looks plausible. Robot-learning benchmarks ask whether the policy succeeds in the loop. A WAM needs both perspectives, but neither is sufficient on its own. A useful protocol must also ask whether the prediction is physically realizable, whether coherence survives long horizons, and what compute, memory, and latency are required to obtain the result.

Visual fidelity metrics. The video-generation lineage contributes FVD [154], FID, LPIPS [196], PSNR, SSIM [168], and DreamSim [39]. These metrics are cheap and familiar, but they reward visual realism rather than action utility. A crisp rollout can still be useless when contact, geometry, or proprioception is wrong. A visually sparse representation can still drive a successful policy when it captures the variables needed for control. Fast-WAM [193] and GigaWorld-Policy [186] bear this out by dropping the rendered forecast at inference without losing control, while WorldScore [31] and WorldSimBench [136] fold visual fidelity into broader assessments that pair perceptual scores against closed-loop action success.

Closed-loop benchmarks. Closed-loop evaluation measures whether the policy actually completes the task. These benchmarks differ mainly in where the rollout comes from, which sets a throughput-versus-validity trade-off. Simulation sits at the high-throughput end, where benchmarks probe different competences. The LIBERO family [34, 105, 159, 211] and RoboMME [22] focus on language-conditioned manipulation and multimodal evaluation. robomimic [116], MetaWorld [190], ManiSkill 2 [47], and ManiSkill 3 [150] provide manipulation suites with different levels of control diversity and physics coverage. HomeRobot [189], VLABench [197], RoboEval [165], and RoboVerse [43] broaden the benchmark space toward navigation, embodied reasoning, and multi-backend evaluation. Real-robot arenas such as RoboArena [5] and RoboChallenge [179] trade throughput for physical validity by running the policy on hardware. DWorldEval [97] closes the loop inside a learned world model instead, scoring a policy from imagined rollouts whose success correlates with real execution. Cascaded WAMs add another evaluation layer because the generated future and the action translator can each fail separately. Video-to-flow and video-to-trajectory methods therefore need both substrate checks and execution checks, as illustrated by NovaFlow, Dream2Flow, TraceGen, and RoboFlow4D [28, 86, 87, 104]. Closed-loop evaluation measures the right outcome, but every trial consumes a robot, a simulator, or a learned generator for the duration of the task.

Physical plausibility and long-horizon coherence. Some WAM properties are visible in a rollout but poorly captured by a single success number. Physical plausibility can be evaluated through tactile and force prediction errors, as in OmniVTA [206] and AdaWorldPolicy [191], or through kinematic and geometry-consistency checks, as in DexWM [44] and MVISTA-4D [160]. A standard metric has not yet emerged. Long-horizon coherence is similarly under-specified. Reported numbers usually reduce to task success after a fixed number of replans, and current benchmarks do not test hour-scale rollout at control frequency. RoboScape [139] brings physics signals into evaluation, while PhysWorld [117] points to the broader problem of judging physical world models.

What current evaluation misses. Current evaluation has a predictable ordering. Visual-fidelity metrics are cheap but weakly tied to downstream success. Closed-loop simulation is more predictive but consumes task-time rollouts. Real-robot arenas are physically valid but expensive to repeat. Physical-plausibility and long-horizon metrics remain underdeveloped. The same ordering appears in data collection. A teleoperation hour costs a robot and an operator, a simulator hour costs compute, an internet-video hour costs curation, and a learned-model rollout costs a forward pass. This is why proxy evaluators are attractive despite depending on the generator’s own fidelity. The WAMs that matter for real use are also the hardest to evaluate, because chunked and interactive closed-loop control require longer horizons, stricter latency constraints, and real-environment variation. The field is therefore moving toward a two-stage protocol: inexpensive visual and representation screens, followed by selective closed-loop tests under explicit compute, memory, and latency budgets. The conversion factor between the cheap screen and the expensive test remains the central open question in WAM evaluation, and it motivates the frontiers in Section 7.

7 Open Challenges

Sections 4 through 6 define the design axes, core properties, data sources, and evaluation practice of current World Action Models. The open challenges below use those axes to separate problems that are often discussed together. The boundaries are analytical rather than absolute: data affects generalization, action abstraction affects physics, and inference cost affects almost every design choice. Still, each subsection asks a different question. We first ask how much future computation a WAM should build into its design or spend at runtime, then ask which data should train each stage. We then move through memory, generalization, action grounding, physical plausibility, and evaluation.

7.1 Dream More or Act More?

Prediction helps action only when its latency fits the decision loop. Richer futures, larger backbones, and deeper denoising can improve the substrate available to the action side, but they also lengthen the path from observation to control. Current WAMs therefore split between computing a detailed future and extracting just enough future-trained structure for action. S-VAM distills video diffusion into geometric and semantic foresight features, Fast-WAM removes the future-video branch at inference, GigaWorld-Policy masks attention so future video can be skipped, and X-WAM lets action denoising finish before video denoising completes [51, 181, 186, 193]. These methods suggest that the video objective can be more important than a fully rendered test-time forecast.

The other side keeps stronger imagination in the control path and pays for it in latency. DreamZero brings a large video backbone into closed-loop control through a custom optimization stack, CosmosPolicy uses repeated model and value queries for planning, and NovaPlan places video generation before flow extraction and geometric grounding [38, 79, 188]. The open problem is not simply to make every model smaller. It is to expose a controllable fidelity-latency curve, so a WAM can choose how much future information the action actually needs.

That choice should also become a runtime decision. A video policy may run the same denoising schedule at every step, and a latent planner may fix its horizon and search depth before it knows whether the next state is routine transit or contact-sensitive manipulation [55, 140]. Coarse switches help but remain limited. A forecast-confidence method can choose when to replan, yet a binary execute-or-replan decision cannot decide how much of the remaining horizon should be regenerated [161]. DreamAvoid, SANTS, NoiseGate, and HarmoWAM show that the trigger can be learned from task phase, denoising state, or process cues [33, 36, 67, 146]. The open problem is to spend generative compute by expected action value: more when uncertainty, contact, or irreversible error is near, and less when the next action is routine.

This budget question also clarifies the relation between video world models and WAMs. Video generation has pushed toward longer and more detailed futures, while WAMs increasingly try to compute only the future information needed for action. The two lines now meet at the substrate level of Section 4.2: video generators add action conditioning, and WAMs borrow latent or joint-embedding substrates from video backbones [128, 129, 166, 167, 198]. The practical obstacle is that a video-scale model reaches the control rate only after substantial optimization, while skipping video tokens recovers speed by giving up the rendered forecast [186, 188, 193]. A stronger design would train one backbone with video objectives, then let the action path choose how much of that backbone to run at inference.

7.2 What Data Should Each Stage Learn From?

Section 6 organizes WAM data by source, but the harder question is where each source belongs in training. A WAM usually needs visual dynamics, embodiment alignment, executable action, and sometimes a post-training controller or scheduler. These variables are not learned from the same evidence. Internet video can teach broad visual regularities, egocentric or third-person human demonstrations can align motion with intent, robot teleoperation can ground the final action decoder, and rollout data can train gates or budget controllers. Pooling these sources hides the variable each source supervises. The open challenge is to assign each data source to a stage, an objective, and a model component.

The pretraining stage is the clearest starting point and also the easiest stage to overgeneralize. GR-2 and

DreamDojo show the appeal of broad video pretraining before robot-specific adaptation, while VPP shows that a video prior can be decisive for downstream action when the robot data alone is too narrow [16, 41, 62]. Yet robot-domain video and general human video are not interchangeable. VidMan reports that robot-domain pretraining helps while a mismatched egocentric source can hurt the same downstream setting, and VLA-JEPA finds that human-video gains depend on the evaluation split rather than appearing uniformly [145, 170]. The question is therefore not whether video pretraining helps. The question is which kind of video teaches the dynamics that the later action path will need.

The alignment stage is less settled. Egocentric human data, paired human-robot demonstrations, and wearable capture can bridge the gap between passive video and robot control, but each format carries a different bias. EgoScale separates broad human pretraining from an aligned mid-training stage, EgoMimic uses first-person human demonstrations to improve robot imitation, and EgoVerse finds that scene diversity can matter more than raw volume [77, 135, 205]. FRAPPE makes the same point from a different angle: web-scale egocentric data, task-specific egocentric data, and robot data play distinct roles in transfer [200]. These results argue for a staged curriculum, but they do not yet tell us the minimum capture hardware, viewpoint match, or data quality needed for each embodiment.

The final action stage is where the data bottleneck becomes sharpest. Cross-platform corpora and simulation benches provide a practical base for action supervision, but their action spaces remain uneven across arms, hands, mobile bases, and deformable-object tasks [20, 21, 43, 59, 123, 175]. Action-free recipes try to reduce this bottleneck by discovering latent actions, flow actions, or structured transition codes from video, as in LAPA, DUST, ALAM, and LDA-1B [112, 149, 171, 187]. These recipes reduce the amount of labeled robot control needed, but they do not remove it. Every abstract action still needs a decoder, an inverse-dynamics model, or a robot-specific calibration step before it can become executable control.

Synthetic trajectories add a further stage rather than a complete replacement for action labels. DreamGen uses generated futures to produce neural trajectories, and physical-world-model pipelines can turn video generation into policy training data, but both still inherit generator errors and still need grounding in robot dynamics [73, 117]. Human2Robot shows the other side of the same problem: paired human-robot data can be valuable, but contact-rich skills remain hard to capture even with specialized teleoperation hardware [176]. The missing result is a joint scaling law over source, stage, quality, model size, and action grounding. Without such a law, adding more data remains an engineering choice rather than a predictable design decision.

7.3 Can Memory Keep Up?

Section 5 catalogs mechanisms for persistence, including observation replacement, bounded memory, and spatially indexed state. What remains open is not simply a longer prediction horizon. A useful WAM must maintain task state while actions are being produced, and its memory cost should grow with scene complexity rather than with episode length. Autoregressive prediction exposes the failure mode directly: IRASim accumulates temporal drift despite overlapping context, PhysGen is bounded by a fixed context of physical packages, and sparse-memory chunking still degrades once the scene becomes dynamic [66, 144, 214]. These are memory failures, not only prediction-quality failures.

Promising ingredients exist in separate places. A test-time memory can keep per-step recall bounded, but its update remains sensitive enough to risk forgetting earlier state [27]. Dreamer 4 shows that compact latent state can support continual imagination under closed-loop control, but the abstraction leaves open how to preserve object-level spatial detail across long manipulation episodes [58]. The challenge is to combine bounded memory, spatial indexing, and observation replacement into one method that remains reactive over long tasks.

7.4 How Can WAMs Generalize?

Section 5 treats generalization as a family of shifts, not as one scalar property. A WAM therefore cannot be made generalizable by adding scale alone. The target shift has to be named first, because visual appearance, object dynamics, morphology, action space, camera placement, and contact regime stress different parts of the model. Domain randomization remains a sharp test of this separation. Act2Goal drops steeply under harder randomized conditions, and FRAPPE remains far below a practical rate even when it leads the compared

baselines [200, 208]. Transfer to new motions and objects is also bounded: a video-conditioned policy improves over baselines, but the hardest motion and object splits still leave a large gap [9]. The lesson is that future prediction helps only when the predicted substrate matches the shift that the action path must absorb.

The path toward generalizable WAMs is to match data, substrate, and adaptation to the shift being targeted. Cross-embodiment transfer is the clearest case, because the action space itself changes. DreamZero reports transfer from short video-only adaptation sets, but the need for adaptation shows that morphology shift is not solved by scale alone [188]. Removing video pretraining from a strong policy sharply reduces task completion, which shows that the borrowed video prior matters, but it does not explain which embodiment changes the prior can absorb [62]. EgoScale offers one positive sign through a clean scaling law for dexterous pretraining volume [205]. The open challenge is to design for a declared shift before training, then test whether the chosen substrate and action decoder generalize across that shift rather than reporting transfer after the fact.

7.5 What Grounds Abstract Action?

The action-coupling axis of Section 4 often replaces raw control with an abstract target such as a latent code, a flow field, a point track, or a structured transition. This substitution is useful because it lets a method learn from action-free or weakly labeled video. It is also risky because the action can lose a physical handle. One failure mode is estimator dependence. Two-dimensional object flow cannot represent depth translation, out-of-plane rotation, or contact force, and a rigid-transform decoder breaks when the object is articulated or deformable [81, 177]. Point-track and flow-based adapters make the action interface inspectable, but they inherit errors from trackers, depth estimators, and local linearity assumptions [49, 70, 94].

The second failure mode is opacity. LAPA learns discrete latent actions from video and improves action-free pretraining, but fine grasping remains difficult when the codebook is too coarse [187]. ALAM adds algebraic consistency to latent transitions, yet the resulting structure is still a composition law rather than a force, torque, or contact explanation [149]. LDA-1B learns a broad latent action space, but execution still depends on the decoder that maps that space back to an embodiment-specific command [112]. The open challenge is to keep the data advantage of abstract actions while adding a physical grounding signal and an inspection handle for failure analysis.

7.6 When Is a Future Physical?

Section 5 frames physical plausibility through visual-geometric abstraction, tactile and force prediction, and proprioceptive or kinematic coherence. The difficulty is that a future can look plausible while failing the embodiment. Momentum is a simple example. A flow-based world model conditioned on one frame cannot encode object velocity, so push tasks fail when the object keeps moving after contact ends [50]. Contact and force are harder still. A tactile world model derives a force proxy from tactile-flow divergence, but the proxy is uncalibrated, and a point-cloud substrate omits force magnitude, friction, and surface compliance [68, 192]. Even an executability reward can miss contact if it only scores smoothness and joint limits [162].

Some methods point toward a better standard by making physics part of the substrate or loss. RoboScape and RoboScape-R add geometric and consistency constraints, OmniVTA and DexWM bring tactile or dexterous state into the prediction problem, and AdaWorldPolicy includes action-conditioned future structure in its training objective [44, 139, 148, 191, 206]. These examples remain partial. The open challenge is to train and evaluate futures for being executable by the embodiment, not only visually convincing.

7.7 What Should Evaluation Report?

Section 6 describes visual screens, simulator tests, and closed-loop robot evaluation. The open challenge is to report the variables that make these tiers comparable. World-model quality is not yet a reliable proxy for policy utility, as MotuBrain finds only a weak correlation between its world-model score and downstream success [121]. A single success number can also hide opposite regimes. HarmoWAM shows that a schedule useful for transit can fail at precise interaction, which means that averaging the phases removes the failure mode the evaluation should expose [36].

The missing protocol is an accuracy-at-budget report. A WAM that succeeds slowly is not equivalent to one that acts at the control rate, and a WAM that succeeds over a short horizon is not equivalent to one that keeps state over a long task. Some papers report latency, but often apart from accuracy and on different hardware or task settings [30, 69, 90, 113, 183]. Better examples are emerging: Fast-WAM isolates the value of test-time video generation once video co-training is preserved, while SANTS and NoiseGate report success together with learned reductions in denoising cost [67, 146, 193]. Evaluation should place success, latency, sustained horizon, peak memory, and contact-sensitive failure tags on the same axis.

The challenges are coupled

These challenges are coupled even when the boundaries above separate their main questions. A memory fix that deepens search raises the compute budget. A richer substrate can improve contact realism while making evaluation and inference more expensive. Latent actions reduce the need for labeled robot control, but the same abstraction weakens physical grounding unless the decoder is calibrated. Contact-aware prediction requires force, tactile, or failure data that cheap visual screens do not provide. Upstream estimator errors can appear as both grounding failures and transfer failures. The shared lesson is that a WAM should be advanced as one coupled design space, where progress on one challenge changes the practical trade-offs of the others.

8 Conclusion

This survey has treated World Action Models as a contract between prediction and control: given past observations, past actions, and a task context, a WAM must anticipate the actionable future while producing the action that should follow. This framing separates WAMs from broad world models, video generation models, and plain VLAs, and it explains why works with different names can still belong to the same family. We organized existing works through two complementary views. The philosophy-level view asks what a method is required to generate before action is decoded, while the component-level anatomy locates the same method by predictive substrate, backbone, action coupling, and deployment regime. These views then support the core properties that matter after a model enters a control loop: interactability, causality, persistence, physical plausibility, and generalization. They also show why data and evaluation are not secondary details, since a benchmark can only validate the properties it exposes and a dataset can only teach the action variables it records or infers. Across the works considered here, the common trajectory is not simply toward larger video prediction, but toward more selective imagination that keeps the parts of the future needed for action and discards the rest. The remaining challenges therefore lie in choosing the right substrate, coupling action without leakage, preserving long-horizon state, grounding predictions in physical constraints, and reporting capability together with compute, memory, and latency. By giving these choices a shared vocabulary, this survey aims to make future WAMs easier to compare and to help the field generate less of the future while preserving more of what embodied action requires.

References

- [1] Build AI. Egocentric-10k, 2025. <https://huggingface.co/datasets/builddotai/Egocentric-10K>.
- [2] Tuo An, Jindou Jia, Gen Li, Jingliang Li, Chuhao Zhou, Pengfei Liu, Bofan Lyu, Jiaqi Bai, Xinying Guo, Geng Li, and Jianfei Yang. Feedback world model enables precise guidance of diffusion policy, 2026. <https://arxiv.org/abs/2605.15705>.
- [3] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture, 2023. <https://arxiv.org/abs/2301.08243>.
- [4] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Mojtaba Komeili, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, Sergio Arnaud, Abha Gejji, Ada Martin, Francois Robert Hogan, Daniel Dugas, Piotr Bojanowski, Vasil Khalidov, Patrick Labatut, Francisco Massa, Marc Szafraniec, Kapil Krishnakumar, Yong Li, Xiaodong Ma, Sarath Chandar, Franziska Meier, Yann LeCun, Michael Rabbat, and Nicolas Ballas. V-jepa 2: Self-supervised video models enable understanding, prediction and planning, 2025. <https://arxiv.org/abs/2506.09985>.

- [5] Pranav Atreya, Karl Pertsch, Tony Lee, Moo Jin Kim, Arhan Jain, Artur Kuramshin, Clemens Eppner, Cyrus Neary, Edward Hu, Fabio Ramos, Jonathan Tremblay, Kanav Arora, Kirsty Ellis, Luca Macesanu, Marcel Torne Villasevil, Matthew Leonard, Meedeum Cho, Ozgur Aslan, Shivin Dass, Jie Wang, William Reger, Xingfang Yuan, Xuning Yang, Abhishek Gupta, Dinesh Jayaraman, Glen Berseth, Kostas Daniilidis, Roberto Martin-Martin, Youngwoon Lee, Percy Liang, Chelsea Finn, and Sergey Levine. RoboArena: Distributed real-world evaluation of generalist robot policies. In *Conference on Robot Learning*, 2025. <https://arxiv.org/abs/2506.18123>.
- [6] Adrien Bardes, Jean Ponce, and Yann LeCun. Mc-jepa: A joint-embedding predictive architecture for self-supervised learning of motion and content features, 2023. <https://arxiv.org/abs/2307.12698>.
- [7] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video, 2024. <https://arxiv.org/abs/2404.08471>.
- [8] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking, 2023. <https://arxiv.org/abs/2309.01918>.
- [9] Homanga Bharadhwaj, Debidatta Dwibedi, Abhinav Gupta, Shubham Tulsiani, Carl Doersch, Ted Xiao, Dhruv Shah, Fei Xia, Dorsa Sadigh, and Sean Kirmani. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. In *CoRL Workshop on Cross-Embodiment*, 2024. <https://arxiv.org/abs/2409.16283>.
- [10] Hongzhe Bi, Hengkai Tan, Shenghao Xie, Zeyuan Wang, Shuhe Huang, Haitian Liu, Ruowen Zhao, Yao Feng, Chendong Xiang, Yinze Rong, Hongyan Zhao, Hanyu Liu, Zhizhong Su, Lei Ma, Hang Su, and Jun Zhu. Motus: A unified latent action world model, 2025. <https://arxiv.org/abs/2512.13030>.
- [11] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Motukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0 : A vision-language-action flow model for general robot control, 2024. <https://arxiv.org/abs/2410.24164>.
- [12] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control, 2023. <https://arxiv.org/abs/2307.15818>.
- [13] Jun Cen, Siteng Huang, Yuqian Yuan, Kehan Li, Hangjie Yuan, Chaohui Yu, Bohan Hou, Yuming Jiang, Jiayan Guo, Xin Li, Hao Luo, Fan Wang, Deli Zhao, and Hao Chen. Rynvla-002: A unified vision-language-action and world model. *CoRR*, abs/2511.17502, 2025. doi: 10.48550/ARXIV.2511.17502. <https://doi.org/10.48550/arXiv.2511.17502>.
- [14] Jun Cen, Chaohui Yu, Hangjie Yuan, Yuming Jiang, Siteng Huang, Jiayan Guo, Xin Li, Yibing Song, Hao Luo, Fan Wang, Deli Zhao, and Hao Chen. Worldvla: Towards autoregressive action world model, 2025. <https://arxiv.org/abs/2506.21539>.
- [15] Vivek Chavan, Yasmina Imgrund, Tung Dao, Sanwantri Bai, Bosong Wang, Ze Lu, Oliver Heimann, and Jörg Krüger. Indego: A dataset of industrial scenarios and collaborative work for egocentric assistants, 2025. <https://arxiv.org/abs/2511.19684>.
- [16] Chi-Lam Cheang, Guangzeng Chen, Ya Jing, Tao Kong, Hang Li, Yifeng Li, Yuxiao Liu, Hongtao Wu, Jiafeng Xu, Yichu Yang, Hanbo Zhang, and Minzhao Zhu. Gr-2: A generative video-language-action model with web-scale knowledge for robot manipulation, 2024. <https://arxiv.org/abs/2410.06158>.
- [17] Boyuan Chen, Tianyuan Zhang, Haoran Geng, Caiyi Zhang, Peihao Li, Kiwhan Song, William T. Freeman, Jitendra Malik, Pieter Abbeel, Russ Tedrake, Vincent Sitzmann, and Yilun Du. Large video planner enables generalizable robot control. *arXiv preprint arXiv:2512.15840*, 2025. <https://arxiv.org/abs/2512.15840>.

- [18] Chang Chen, Yi-Fu Wu, Jaesik Yoon, and Sungjin Ahn. Transdreamer: Reinforcement learning with transformer world models, 2022. <https://arxiv.org/abs/2202.09481>.
- [19] Jiayi Chen, Wenxuan Song, Pengxiang Ding, Ziyang Zhou, Han Zhao, Feilong Tang, Donglin Wang, and Haoang Li. Unified diffusion v1a: Vision-language-action model via joint discrete denoising diffusion process, 2026. <https://arxiv.org/abs/2511.01718>.
- [20] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation, 2025. <https://arxiv.org/abs/2506.18088>.
- [21] Open X-Embodiment Collaboration, Abby O’Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tung, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anchit Gupta, Andrew Wang, Andrey Kolobov, Anikait Singh, Animesh Garg, Aniruddha Kembhavi, Annie Xie, Anthony Brohan, Antonin Raffin, Archit Sharma, Arefeh Yavary, Arhan Jain, Ashwin Balakrishna, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Blake Wulfe, Brian Ichter, Cewu Lu, Charles Xu, Charlotte Le, Chelsea Finn, Chen Wang, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Christopher Agia, Chuer Pan, Chuyuan Fu, Coline Devin, Danfei Xu, Daniel Morton, Danny Driess, Daphne Chen, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Felipe Vieira Frujeri, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guangwen Yang, Guanzhi Wang, Hao Su, Hao-Shu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homanga Bharadhwaj, Homer Walke, Hongjie Fang, Huy Ha, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jad Abou-Chakra, Jaehyung Kim, Jaimyn Drake, Jan Peters, Jan Schneider, Jasmine Hsu, Jay Vakil, Jeannette Bohg, Jeffrey Bingham, Jeffrey Wu, Jensen Gao, Jiaheng Hu, Jiajun Wu, Jialin Wu, Jiankai Sun, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jimmy Wu, Jingpei Lu, Jingyun Yang, Jitendra Malik, João Silvério, Joey Hejna, Jonathan Booher, Jonathan Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Black, Kevin Lin, Kevin Zhang, Kiana Ehsani, Kiran Lekkala, Kirsty Ellis, Krishan Rana, Krishnan Srinivasan, Kuan Fang, Kunal Pratap Singh, Kuo-Hao Zeng, Kyle Hatch, Kyle Hsu, Laurent Itti, Lawrence Yunliang Chen, Lerrel Pinto, Li Fei-Fei, Liam Tan, Linxi "Jim" Fan, Lionel Ott, Lisa Lee, Luca Weihs, Magnum Chen, Marion Lepert, Marius Memmel, Masayoshi Tomizuka, Masha Itkina, Mateo Guaman Castro, Max Spero, Maximilian Du, Michael Ahn, Michael C. Yip, Mingtong Zhang, Mingyu Ding, Minh Heo, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Muhammad Zubair Irshad, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Ning Liu, Norman Di Palo, Nur Muhammad Mahi Shafullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Pannag R Sanketi, Patrick "Tree" Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundaesan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Rohan Bajjal, Rosario Scalise, Rose Hendrix, Roy Lin, Runjia Qian, Ruohan Zhang, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Shan Lin, Sherry Moore, Shikhar Bahl, Shivin Dass, Shubham Sonawani, Shubham Tulsiani, Shuran Song, Sichun Xu, Siddhant Haldar, Siddharth Karamcheti, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Subramanian Ramamoorthy, Sudeep Dasari, Suneel Belkhale, Sungjae Park, Suraj Nair, Suvir Mirchandani, Takayuki Osa, Tanmay Gupta, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Thomas Kollar, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Trinity Chung, Vidhi Jain, Vikash Kumar, Vincent Vanhoucke, Vitor Guizilini, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiangyu Chen, Xiaolong Wang, Xinghao Zhu, Xinyang Geng, Xiyuan Liu, Xu Liangwei, Xuanlin Li, Yansong Pang, Yao Lu, Yecheng Jason Ma, Yejin Kim, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Yilin Wu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yongqiang Dou, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yue Cao, Yueh-Hua Wu, Yujin Tang, Yuke Zhu, Yunchu Zhang, Yunfan Jiang, Yunshuang Li, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zehan Ma, Zhuo Xu, Zichen Jeff Cui, Zichen Zhang, Zipeng Fu, and Zipeng Lin. Open x-embodiment: Robotic learning datasets and rt-x models, 2023. <https://arxiv.org/abs/2310.08864>.
- [22] Yinpei Dai, Hongze Fu, Jayjun Lee, Yuejiang Liu, Haoran Zhang, Jianing Yang, Chelsea Finn, Nima Fazeli, and Joyce Chai. RoboMME: Benchmarking and understanding memory for robotic generalist policies. In *International Conference on Machine Learning*, 2026. <https://arxiv.org/abs/2603.04639>.
- [23] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos,

- Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset, 2018. <https://arxiv.org/abs/1804.02748>.
- [24] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning, 2020. <https://arxiv.org/abs/1910.11215>.
- [25] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining, 2025. <https://arxiv.org/abs/2505.14683>.
- [26] Haoge Deng, Ting Pan, Haiwen Diao, Zhengxiong Luo, Yufeng Cui, Huchuan Lu, Shiguang Shan, Yonggang Qi, and Xinlong Wang. Autoregressive video generation without vector quantization, 2024. <https://arxiv.org/abs/2412.14169>.
- [27] Yueci Deng, Guiliang Liu, and Kui Jia. Dexworldmodel: Causal latent world modeling towards automated learning of embodied tasks, 2026. <https://arxiv.org/abs/2604.16484>.
- [28] Karthik Dharmarajan, Wenlong Huang, Jiajun Wu, Li Fei-Fei, and Ruohan Zhang. Dream2flow: Bridging video generation and open-world manipulation with 3d object flow. *arXiv preprint arXiv:2512.24766*, 2025. <https://arxiv.org/abs/2512.24766>.
- [29] Yilun Du, Mengjiao Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Joshua B. Tenenbaum, Dale Schuurmans, and Pieter Abbeel. Learning universal policies via text-guided video generation. In *Advances in Neural Information Processing Systems*, 2023. <https://arxiv.org/abs/2302.00111>.
- [30] Yilun Du, Mengjiao Yang, Pete Florence, Fei Xia, Ayzaan Wahid, Brian Ichter, Pierre Sermanet, Tianhe Yu, Pieter Abbeel, Joshua B. Tenenbaum, Leslie Kaelbling, Andy Zeng, and Jonathan Tompson. Video language planning. In *International Conference on Learning Representations*, 2024. <https://arxiv.org/abs/2310.10625>.
- [31] Haoyi Duan, Hong-Xing Yu, Sirui Chen, Li Fei-Fei, and Jiajun Wu. Worldscore: A unified evaluation benchmark for world generation. In *IEEE/CVF International Conference on Computer Vision*, 2025. <https://arxiv.org/abs/2504.00983>.
- [32] Liaoyuan Fan, Zetian Xu, Chen Cao, Wenyao Zhang, Mingqi Yuan, and Jiayu Chen. Aim: Intent-aware unified world action modeling with spatial value maps, 2026. <https://arxiv.org/abs/2604.11135>.
- [33] Xianzhe Fan, Yuxiang Lu, Shenyuan Gao, Xiaoyang Wu, Ruihua Han, Manling Li, and Hengshuang Zhao. Dreamavoid: Critical-phase test-time dreaming to avoid failures in vla policies, 2026. <https://arxiv.org/abs/2605.11750>.
- [34] Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, Jinlan Fu, Jingjing Gong, and Xipeng Qiu. LIBERO-Plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025. <https://arxiv.org/abs/2510.13626>.
- [35] Zhengcong Fei, Mingyuan Fan, and Junshi Huang. A-jepa: Joint-embedding predictive architecture can listen, 2024. <https://arxiv.org/abs/2311.15830>.
- [36] Qiuxuan Feng, Jiale Yu, Jiaming Liu, Yueru Jia, Zhuangzhe Wu, Hao Chen, Zezhong Qian, Shuo Gu, Peng Jia, Siwei Ma, and Shanghang Zhang. Harmowam: Harmonizing generalizable and precise manipulation via adaptive world action models, 2026. <https://arxiv.org/abs/2605.10942>.
- [37] Yao Feng, Hengkai Tan, Xinyi Mao, Chendong Xiang, Guodong Liu, Shuhe Huang, Hang Su, and Jun Zhu. Vidar: Embodied video diffusion model for generalist manipulation, 2025. <https://arxiv.org/abs/2507.12898>.
- [38] Jiahui Fu, Junyu Nan, Lingfeng Sun, Hongyu Li, Jianing Qian, Jennifer L. Barry, Kris Kitani, and George Konidaris. Novaplan: Zero-shot long-horizon manipulation via closed-loop video language planning, 2026. <https://arxiv.org/abs/2602.20119>.
- [39] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data. In *Advances in Neural Information Processing Systems*, 2023.
- [40] Shenyuan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, and Chuang Gan. Adaworld: Learning adaptable world models with latent actions, 2025. <https://arxiv.org/abs/2503.18938>.
- [41] Shenyuan Gao, William Liang, Kaiyuan Zheng, Ayaan Malik, Seonghyeon Ye, Sihyun Yu, Wei-Cheng Tseng, Yuzhu Dong, Kaichun Mo, Chen-Hsuan Lin, Qianli Ma, Seungjun Nah, Loic Magne, Jiannan Xiang, Yuqi Xie, Ruijie

- Zheng, Dantong Niu, You Liang Tan, K. R. Zentner, George Kurian, Suneel Indupuru, Pooya Jannaty, Jinwei Gu, Jun Zhang, Jitendra Malik, Pieter Abbeel, Ming-Yu Liu, Yuke Zhu, Joel Jang, and Linxi Jim Fan. Dreamdojo: A generalist robot world model from large-scale human videos, 2026. <https://arxiv.org/abs/2602.06949>.
- [42] Zirui Ge, Pengxiang Ding, Baohua Yin, Qishen Wang, Zhiyong Xie, Yemin Wang, Jinbo Wang, Hengtao Li, Runze Suo, Wenxuan Song, Han Zhao, Shangke Lyu, Zhaoxin Fan, Haoang Li, Ran Cheng, Cheng Chi, Huibin Ge, Yaozhi Luo, and Donglin Wang. Vampo: Policy optimization for improving visual dynamics in video action models, 2026. <https://arxiv.org/abs/2603.19370>.
- [43] Haoran Geng, Feishi Wang, Songlin Wei, Yuyang Li, Bangjun Wang, Boshi An, Charlie Tianyue Cheng, Haozhe Lou, Peihao Li, Yen-Jen Wang, Yutong Liang, Dylan Goetting, Chaoyi Xu, Haozhe Chen, Yuxi Qian, Yiran Geng, Jiageng Mao, Weikang Wan, Mingtong Zhang, Jiangran Lyu, Siheng Zhao, Jiazhao Zhang, Jialiang Zhang, Chengyang Zhao, Haoran Lu, Yufei Ding, Ran Gong, Yuran Wang, Yuxuan Kuang, Ruihai Wu, Baoxiong Jia, Carlo Sferrazza, Hao Dong, Siyuan Huang, Yue Wang, Jitendra Malik, and Pieter Abbeel. RoboVerse: Towards a unified platform, dataset and benchmark for scalable and generalizable robot learning, 2025. <https://arxiv.org/abs/2504.18904>.
- [44] Raktim Gautam Goswami, Amir Bar, David Fan, Tsung-Yen Yang, Gaoyue Zhou, Prashanth Krishnamurthy, Michael Rabbat, Farshad Khorrani, and Yann LeCun. World models for learning dexterous hand-object interactions from human videos, 2025. <https://arxiv.org/abs/2512.13644>.
- [45] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abrahm Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video, 2022. <https://arxiv.org/abs/2110.07058>.
- [46] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, Eugene Byrne, Zach Chavis, Joya Chen, Feng Cheng, Fu-Jen Chu, Sean Crane, Avijit Dasgupta, Jing Dong, Maria Escobar, Cristhian Forigua, Abrahm Gebreselasie, Sanjay Haresh, Jing Huang, Md Mohaiminul Islam, Suyog Jain, Rawal Khirodkar, Devansh Kukreja, Kevin J Liang, Jia-Wei Liu, Sagnik Majumder, Yongsen Mao, Miguel Martin, Effrosyni Mavroudi, Tushar Nagarajan, Francesco Ragusa, Santhosh Kumar Ramakrishnan, Luigi Seminara, Arjun Somayazulu, Yale Song, Shan Su, Zihui Xue, Edward Zhang, Jinxu Zhang, Angela Castillo, Changan Chen, Xinzhu Fu, Ryosuke Furuta, Cristina Gonzalez, Prince Gupta, Jiabo Hu, Yifei Huang, Yiming Huang, Weslie Khoo, Anush Kumar, Robert Kuo, Sach Lakhavani, Miao Liu, Mi Luo, Zhengyi Luo, Brighid Meredith, Austin Miller, Oluwatuminu Oguntola, Xiaqing Pan, Penny Peng, Shraman Pramanick, Merey Ramazanova, Fiona Ryan, Wei Shan, Kiran Somasundaram, Chenan Song, Audrey Southerland, Masatoshi Tateno, Huiyu Wang, Yuchen Wang, Takuma Yagi, Mingfei Yan, Xitong Yang, Zecheng Yu, Shengxin Cindy Zha, Chen Zhao, Ziwei Zhao, Zhifan Zhu, Jeff Zhuo, Pablo Arbelaez, Gedas Bertasius, David Crandall, Dima Damen, Jakob Engel, Giovanni Maria Farinella, Antonino Furnari, Bernard Ghanem, Judy Hoffman, C. V. Jawahar, Richard Newcombe, Hyun Soo Park, James M. Rehg, Yoichi Sato, Manolis Savva, Jianbo Shi, Mike Zheng Shou, and Michael Wray. Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives, 2024. <https://arxiv.org/abs/2311.18259>.
- [47] Jiayuan Gu, Fanbo Xiang, Xuanlin Li, Zhan Ling, Xiqiang Liu, Tongzhou Mu, Yihe Tang, Stone Tao, Xinyue Wei, Yunchao Yao, Xiaodi Yuan, Pengwei Xie, Zhiao Huang, Rui Chen, and Hao Su. Maniskill2: A unified benchmark for generalizable manipulation skills, 2023. <https://arxiv.org/abs/2302.04659>.
- [48] Songen Gu, Yunuo Cai, Tianyu Wang, Simo Wu, and Yanwei Fu. Say, dream, and act: Learning video world models for instruction-driven robot manipulation. *arXiv preprint arXiv:2602.10717*, 2026. <https://arxiv.org/abs/2602.10717>.
- [49] Jiarui Guan, Wenshuai Zhao, Yue Pei, Ziliang Chen, Arno Solin, and Juho Kannala. Point tracking improves world action models, 2026. <https://arxiv.org/abs/2605.23856>.

- [50] Jun Guo, Xiaojuan Ma, Yikai Wang, Min Yang, Huaping Liu, and Qing Li. Flowdreamer: A rgb-d world model with flow-based motion representations for robot manipulation, 2025. <https://arxiv.org/abs/2505.10075>.
- [51] Jun Guo, Qiwei Li, Peiyan Li, Zilong Chen, Nan Sun, Yifei Su, Heyun Wang, Yuan Zhang, Xinghang Li, and Huaping Liu. Unified 4d world action modeling from video priors with asynchronous denoising, 2026. <https://arxiv.org/abs/2604.26694>.
- [52] Yanjiang Guo, Yucheng Hu, Jianke Zhang, Yen-Jen Wang, Xiaoyu Chen, Chaochao Lu, and Jianyu Chen. Prediction with action: Visual policy learning via joint denoising process, 2024. <https://arxiv.org/abs/2411.18179>.
- [53] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning, 2024. <https://arxiv.org/abs/2307.04725>.
- [54] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning, 2024.
- [55] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels, 2019. <https://arxiv.org/abs/1811.04551>.
- [56] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models, 2022. <https://arxiv.org/abs/2010.02193>.
- [57] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models, 2023. <https://arxiv.org/abs/2301.04104>.
- [58] Danijar Hafner, Wilson Yan, and Timothy Lillicrap. Training agents inside of scalable world models, 2025. <https://arxiv.org/abs/2509.24527>.
- [59] Songhao Han, Boxiang Qiu, Yue Liao, Siyuan Huang, Chen Gao, Shuicheng Yan, and Si Liu. Robocerebra: A large-scale benchmark for long-horizon robotic manipulation evaluation, 2025. <https://arxiv.org/abs/2506.06677>.
- [60] Ryan Hoque, Peide Huang, David J. Yoon, Mouli Sivapurapu, and Jian Zhang. Egodex: Learning dexterous manipulation from large-scale egocentric video, 2026. <https://arxiv.org/abs/2505.11709>.
- [61] Bohan Hou, Gen Li, Jindou Jia, Tuo An, Xinying Guo, Sicong Leng, Haoran Geng, Yanjie Ze, Tatsuya Harada, Philip Torr, Oier Mees, Marc Pollefeys, Zhuang Liu, Jiajun Wu, Pieter Abbeel, Jitendra Malik, Yilun Du, and Jianfei Yang. World model for robot learning: A comprehensive survey, 2026. <https://arxiv.org/abs/2605.00080>.
- [62] Yucheng Hu, Yanjiang Guo, Pengchao Wang, Xiaoyu Chen, Yen-Jen Wang, Jianke Zhang, Koushil Sreenath, Chaochao Lu, and Jianyu Chen. Video prediction policy: A generalist robot policy with predictive visual representations. *arXiv preprint arXiv:2412.14803*, 2024. <https://arxiv.org/abs/2412.14803>.
- [63] Yucheng Hu, Jianke Zhang, Yuanfei Luo, Yanjiang Guo, Xiaoyu Chen, Xinshu Sun, Kun Feng, Qingzhou Lu, Sheng Chen, Yangang Zhang, Wei Li, and Jianyu Chen. BagelVLA: Enhancing long-horizon manipulation via interleaved vision-language-action generation, 2026. <https://arxiv.org/abs/2602.09849>. arXiv:2602.09849.
- [64] Ziou Hu, Xiangtong Yao, Yuan Meng, Zhenshan Bing, and Alois Knoll. Dreaming the unseen: World model-regularized diffusion policy for out-of-distribution robustness, 2026. <https://arxiv.org/abs/2603.21017>.
- [65] Shuaiyi Huang, Mara Levy, Zhenyu Jiang, Anima Anandkumar, Yuke Zhu, Linxi Fan, De-An Huang, and Abhinav Shrivastava. ARDuP: Active region video diffusion for universal policies. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 8465–8472, 2024. <https://arxiv.org/abs/2406.13301>.
- [66] Siyuan Huang, Liliang Chen, Pengfei Zhou, Shengcong Chen, Zhengkai Jiang, Yue Hu, Yue Liao, Peng Gao, Hongsheng Li, Maoqing Yao, and Guanghui Ren. Enerverse: Envisioning embodied future space for robotics manipulation, 2025. <https://arxiv.org/abs/2501.01895>.
- [67] Wen Huang, Haoran Sun, Yongjian Guo, Yunxuan Ma, Haoran Li, Jing Long, Zhouying Mo, Zhong Guan, Yucheng Guo, Shuai Di, and Junwu Xiong. Noisegate: Learning per-latent timestep schedules as information gating in world action models, 2026. <https://arxiv.org/abs/2605.07794>.
- [68] Wenlong Huang, Yu-Wei Chao, Arsalan Mousavian, Ming-Yu Liu, Dieter Fox, Kaichun Mo, and Li Fei-Fei. Pointworld: Scaling 3d world models for in-the-wild robotic manipulation, 2026. <https://arxiv.org/abs/2601.03782>.

- [69] Xijie Huang, Weiqi Gai, Tianyue Wu, Congyu Wang, Zhiyang Liu, Xin Zhou, Yuze Wu, and Fei Gao. Navdreamer: Video models as zero-shot 3d navigators, 2026. <https://arxiv.org/abs/2602.09765>.
- [70] Adam Hung, Bardienus Pieter Duisterhof, and Jeffrey Ichnowski. 3pointr: 3d point tracks for learning manipulation from unconstrained human videos, 2026. <https://arxiv.org/abs/2603.08485>.
- [71] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization, 2025. <https://arxiv.org/abs/2504.16054>.
- [72] Physical Intelligence, Bo Ai, Ali Amin, Raichelle Aniceto, Ashwin Balakrishna, Greg Balke, Kevin Black, George Bokinsky, Shihao Cao, Thomas Charbonnier, Vedant Choudhary, Foster Collins, Ken Conley, Grace Connors, James Darpinian, Karan Dhabalia, Maitrayee Dhaka, Jared DiCarlo, Danny Driess, Michael Equi, Adnan Esmail, Yunhao Fang, Chelsea Finn, Catherine Glossop, Thomas Godden, Ivan Goryachev, Lachlan Groom, Haroun Habeeb, Hunter Hancock, Karol Hausman, Gashon Hussein, Victor Hwang, Brian Ichter, Connor Jacobsen, Szymon Jakubczak, Rowan Jen, Tim Jones, Gregg Kammerer, Ben Katz, Liyiming Ke, Mairbek Khadikov, Chandra Kuchi, Marinda Lamb, Devin LeBlanc, Brendon LeCount, Sergey Levine, Xinyu Li, Adrian Li-Bell, Vladislav Lialin, Zhonglin Liang, Wallace Lim, Yao Lu, Enyu Luo, Vishnu Mano, Nandan Marwaha, Aikys Mongush, Liam Murphy, Suraj Nair, Tyler Patterson, Karl Pertsch, Allen Z. Ren, Gavin Schelske, Charvi Sharma, Baifeng Shi, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, Will Stoeckle, Jiaming Tang, Jimmy Tanner, Shalom Tekeste, Marcel Torne, Kyle Vedder, Quan Vuong, Anna Walling, Haohuan Wang, Jason Wang, XuDong Wang, Chris Whalen, Samuel Whitmore, Blake Williams, Charles Xu, Sukwon Yoo, Lili Yu, Wuming Zhang, Zhuoyang Zhang, and Ury Zhilinsky. $\pi_{0.7}$: a steerable generalist robotic foundation model with emergent capabilities, 2026. <https://arxiv.org/abs/2604.15483>.
- [73] Joel Jang, Seonghyeon Ye, Zongyu Lin, Jiannan Xiang, Johan Bjorck, Yu Fang, Fengyuan Hu, Spencer Huang, Kaushil Kundalia, Yen-Chen Lin, Loic Magne, Ajay Mandlekar, Avnish Narayan, You Liang Tan, Guanzhi Wang, Jing Wang, Qi Wang, Yinzheng Xu, Xiaohui Zeng, Kaiyuan Zheng, Ruijie Zheng, Ming-Yu Liu, Luke Zettlemoyer, Dieter Fox, Jan Kautz, Scott Reed, Yuke Zhu, and Linxi Fan. Dreamgen: Unlocking generalization in robot learning through video world models, 2025. <https://arxiv.org/abs/2505.12705>.
- [74] Ahad Jawaaid and Yu Xiang. Openego: A large-scale multimodal egocentric dataset for dexterous manipulation, 2025. <https://arxiv.org/abs/2509.05513>.
- [75] Yuhua Jiang, Yijun Guo, Hongbing Yang, Guojun Lei, Nuo Chen, Yinuo Zhang, Shaoqiang Yan, Bo Lin, Feifei Gao, and Biqing Qi. Ckt-wam: Parameter-efficient context knowledge transfer between world action models, 2026. <https://arxiv.org/abs/2605.06247>.
- [76] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. CoTracker: It is better to track together. In *European Conference on Computer Vision*, 2024. <https://arxiv.org/abs/2307.07635>.
- [77] Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranay Mathur, Shuo Cheng, Chen Wang, Judy Hoffman, and Danfei Xu. Egomimic: Scaling imitation learning via egocentric video, 2024. <https://arxiv.org/abs/2410.24221>.
- [78] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model, 2024. <https://arxiv.org/abs/2406.09246>.
- [79] Moo Jin Kim, Yihuai Gao, Tsung-Yi Lin, Yen-Chen Lin, Yunhao Ge, Grace Lam, Percy Liang, Shuran Song, Ming-Yu Liu, Chelsea Finn, and Jinwei Gu. Cosmos policy: Fine-tuning video models for visuomotor control and planning, 2026. <https://arxiv.org/abs/2601.16163>.
- [80] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. <https://arxiv.org/abs/2304.02643>.
- [81] Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, and Joshua B. Tenenbaum. Learning to act from

- actionless videos through dense correspondences. In *International Conference on Learning Representations*, 2024. <https://arxiv.org/abs/2310.08576>.
- [82] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, Krishna Somandepalli, Hassan Akbari, Yair Alon, Yong Cheng, Josh Dillon, Agrim Gupta, Meera Hahn, Anja Hauth, David Hendon, Alonso Martinez, David Minnen, Mikhail Sirotenko, Kihyuk Sohn, Xuan Yang, Hartwig Adam, Ming-Hsuan Yang, Irfan Essa, Huisheng Wang, David A. Ross, Bryan Seybold, and Lu Jiang. Videopoet: A large language model for zero-shot video generation, 2024. <https://arxiv.org/abs/2312.14125>.
- [83] Yuxuan Kuang, Sungjae Park, Katerina Fragkiadaki, and Shubham Tulsiani. Dex4d: Task-agnostic point track policy for sim-to-real dexterous manipulation, 2026. <https://arxiv.org/abs/2602.15828>.
- [84] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. MegaPose: 6d pose estimation of novel objects via render & compare. In *Conference on Robot Learning*, 2022. <https://arxiv.org/abs/2212.06870>.
- [85] Xiaolei Lang, Yang Wang, Yukun Zhou, Chaojun Ni, Kerui Li, Jiagang Zhu, Tianze Liu, Jiajun Lv, Xingxing Zuo, Yun Ye, Guan Huang, Xiaofeng Wang, and Zheng Zhu. Vag: Dual-stream video-action generation for embodied data synthesis, 2026. <https://arxiv.org/abs/2604.09330>.
- [86] Seungjae Lee, Yoonkyo Jung, Inkook Chun, Yao-Chih Lee, Zikui Cai, Hongjia Huang, Aayush Talreja, Tan Dat Dao, Yongyuan Liang, Jia-Bin Huang, and Furong Huang. Tracegen: World modeling in 3d trace space enables learning from cross-embodiment videos, 2025. <https://arxiv.org/abs/2511.21690>.
- [87] Hongyu Li, Lingfeng Sun, Yafei Hu, Duy Ta, Jennifer Barry, George Konidakis, and Jiahui Fu. NovaFlow: Zero-shot manipulation via actionable flow from generated videos. *arXiv preprint arXiv:2510.08568*, 2025. <https://arxiv.org/abs/2510.08568>.
- [88] Lin Li, Qihang Zhang, Yiming Luo, Shuai Yang, Ruilin Wang, Fei Han, Mingrui Yu, Zelin Gao, Nan Xue, Xing Zhu, Yujun Shen, and Yinghao Xu. Causal world modeling for robot control, 2026. <https://arxiv.org/abs/2601.21998>.
- [89] Peiyan Li, Hongtao Wu, Yan Huang, Chilam Cheang, Liang Wang, and Tao Kong. GR-MG: leveraging partially-annotated data via multi-modal goal-conditioned policy. *IEEE Robotics Autom. Lett.*, 10(2):1912–1919, 2025. doi: 10.1109/LRA.2025.3526436. <https://doi.org/10.1109/LRA.2025.3526436>.
- [90] Peiyan Li, Yixiang Chen, Yuan Xu, Jiabing Yang, Xiangnan Wu, Jun Guo, Nan Sun, Long Qian, Xinghang Li, Xin Xiao, Jing Liu, Nianfeng Liu, Tao Kong, Yan Huang, Liang Wang, and Tieniu Tan. Multi-view video diffusion policy: A 3d spatio-temporal-aware video action model, 2026. <https://arxiv.org/abs/2604.03181>.
- [91] Runze Li, Hongyin Zhang, Junxi Jin, Qixin Zeng, Zifeng Zhuang, Yiqi Tang, Shangke Lyu, and Donglin Wang. World-value-action model: Implicit planning for vision-language-action systems, 2026. <https://arxiv.org/abs/2604.14732>.
- [92] Shalfun Li, Victor Yao, Charles Yang, Truth Qu, Regis Cheng, Ryan Yu, Howard Lu, Newton Von, Vincent Chen, Yohann Tang, Maeve Zhang, Ellie Ma, Gody Li, Sage Yang, Lorien Shu, J. W. Gao, Ethan Chen, Colin Ye, Yu Sun, Elise Mon, PS Zhang, Neo Li, Lily Li, James Wang, Ping Yang, Chris Pan, Lucy Liang, Hang Su, Roy Gan, Hao Wang, and Qian Wang. Wall-wm: Carving world action modeling at the event joints, 2026. <https://arxiv.org/abs/2606.01955>.
- [93] Shuang Li, Yihuai Gao, Dorsa Sadigh, and Shuran Song. Unified video action model, 2025. <https://arxiv.org/abs/2503.00200>.
- [94] Sizhe Lester Li, Evan Kim, Xingjian Bai, Tong Zhao, Tao Pang, Max Simchowitz, and Vincent Sitzmann. Turning video models into generalist robot policies, 2026. <https://arxiv.org/abs/2605.27817>.
- [95] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. In *Conference on Robot Learning*, 2024. <https://arxiv.org/abs/2405.05941>.
- [96] Yajie Li, Bozhou Zhang, Chun Gu, Zipei Ma, Jiahui Zhang, Jiankang Deng, Xiatian Zhu, and Li Zhang. From imagined futures to executable actions: Mixture of latent actions for robot manipulation, 2026. <https://arxiv.org/abs/2605.12167>.
- [97] Yaxuan Li, Zhongyi Zhou, Yefei Chen, Yaokai Xue, and Yichen Zhu. dworldval: Scalable robotic policy evaluation via discrete diffusion world model, 2026. <https://arxiv.org/abs/2604.22152>.

- [98] Yiming Li, Ziang Cao, Andrew Liang, Benjamin Liang, Luoyao Chen, Hang Zhao, and Chen Feng. Egocentric prediction of action target in 3d, 2022. <https://arxiv.org/abs/2203.13116>.
- [99] Junbang Liang, Ruoshi Liu, Ege Ozguroglu, Sruthi Sudhakar, Achal Dave, Pavel Tokmakov, Shuran Song, and Carl Vondrick. Dreamitate: Real-world visuomotor policy learning via video generation. In *Conference on Robot Learning*, 2024. <https://arxiv.org/abs/2406.16862>.
- [100] Junbang Liang, Pavel Tokmakov, Ruoshi Liu, Sruthi Sudhakar, Paarth Shah, Rares Ambrus, and Carl Vondrick. Video generators are robot policies. *arXiv preprint arXiv:2508.00795*, 2025. <https://arxiv.org/abs/2508.00795>.
- [101] Litian Liang, Liuyu Bian, Caiwei Xiao, Jialin Zhang, Linghao Chen, Isabella Liu, Fanbo Xiang, Zhiao Huang, and Hao Su. Robo360: A 3d omniscpective multi-material robotic manipulation dataset, 2023. <https://arxiv.org/abs/2312.06686>.
- [102] Yue Liao, Pengfei Zhou, Siyuan Huang, Donglin Yang, Shengcong Chen, Yuxin Jiang, Yue Hu, Jingbin Cai, Si Liu, Jianlan Luo, Liliang Chen, Shuicheng Yan, Maoqing Yao, and Guanghui Ren. Genie envisioner: A unified world foundation platform for robotic manipulation. *CoRR*, abs/2508.05635, 2025. doi: 10.48550/arXiv.2508.05635. <https://arxiv.org/abs/2508.05635>.
- [103] Minghui Lin, Pengxiang Ding, Shu Wang, Zifeng Zhuang, Yang Liu, Xinyang Tong, Wenxuan Song, Shangke Lyu, Siteng Huang, and Donglin Wang. Hif-vla: Hindsight, insight and foresight through motion representation for vision-language-action models, 2025. <https://arxiv.org/abs/2512.09928>.
- [104] Sixu Lin, Junliang Chen, Huaiyuan Xu, Zhuohao Li, Guangming Wang, Yixiong Jing, Sheng Xu, Runyi Zhao, Brian Sheil, Lap-Pui Chau, and Guiliang Liu. Roboflow4d: A lightweight flow world model toward real-time flow-guided robotic manipulation, 2026. <https://arxiv.org/abs/2605.17522>.
- [105] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning, 2023. <https://arxiv.org/abs/2306.03310>.
- [106] Mengmeng Liu, Diankun Zhang, Jiuming Liu, Jianfeng Cui, Hongwei Xie, Guang Chen, Hangjun Ye, Michael Ying Yang, Francesco Nex, and Hao Cheng. Driveva: Video action models are zero-shot drivers, 2026. <https://arxiv.org/abs/2604.04198>.
- [107] Yuanzhe Liu, Jingyuan Zhu, Yuchen Mo, Gen Li, Xu Cao, Jin Jin, Yifan Shen, Zhengyuan Li, Tianjiao Yu, Wenzhen Yuan, Fangqiang Ding, and Ismini Lourentzou. Palm: Progress-aware policy learning via affordance reasoning for long-horizon robotic manipulation, 2026. <https://arxiv.org/abs/2601.07060>.
- [108] Zeyi Liu, Shuang Li, Eric Cousineau, Siyuan Feng, Benjamin Burchfiel, and Shuran Song. Geometry-aware 4d video generation for robot manipulation, 2025.
- [109] Yunfan Lou, Xiaowei Chi, Xiaojie Zhang, Zezhong Qian, Chengxuan Li, Rongyu Zhang, Yaoxu Lyu, Guoyu Song, Chuyao Fu, Haoxuan Xu, Pengwei Wang, and Shanghang Zhang. Mask world model: Predicting what matters for robust robot policy learning, 2026. <https://arxiv.org/abs/2604.19683>.
- [110] Hongbo Lu, Liang Yao, Chenghao He, Haoyu Wang, Xiang Gu, Xianfei Li, Wenlong Liao, Tao He, and Pai Peng. The DAWN of world-action interactive models, 2026. <https://arxiv.org/abs/2605.11550>. arXiv:2605.11550.
- [111] Qi Lv, Weijie Kong, Hao Li, Jia Zeng, Zherui Qiu, Delin Qu, Haoming Song, Qizhi Chen, Xiang Deng, and Jiangmiao Pang. F1: A vision-language-action model bridging understanding and generation to actions. *CoRR*, abs/2509.06951, 2025. doi: 10.48550/ARXIV.2509.06951. <https://doi.org/10.48550/arXiv.2509.06951>.
- [112] Jiangran Lyu, Kai Liu, Xuheng Zhang, Haoran Liao, Yusen Feng, Wenxuan Zhu, Tingrui Shen, Jiayi Chen, Jiazhao Zhang, Yifei Dong, Wenbo Cui, Senmao Qi, Shuo Wang, Yixin Zheng, Mi Yan, Xuesong Shi, Haoran Li, Dongbin Zhao, Ming-Yu Liu, Zhizheng Zhang, Li Yi, Yizhou Wang, and He Wang. Lda-1b: Scaling latent dynamics action model via universal embodied data ingestion, 2026. <https://arxiv.org/abs/2602.12215>.
- [113] Teli Ma, Jia Zheng, Zifan Wang, Chunli Jiang, Andy Cui, Junwei Liang, and Shuo Yang. Dit4dit: Jointly modeling video dynamics and actions for generalizable robot control, 2026. <https://arxiv.org/abs/2603.10448>.
- [114] Xin Ma, Yaohui Wang, Xinyuan Chen, Gengyun Jia, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent diffusion transformer for video generation, 2024. <https://arxiv.org/abs/2401.03048>.
- [115] Ajay Mandlekar, Jonathan Booher, Max Spero, Albert Tung, Anchit Gupta, Yuke Zhu, Animesh Garg, Silvio Savarese, and Li Fei-Fei. Scaling robot supervision to hundreds of hours with roboturk: Robotic manipulation dataset through human reasoning and dexterity, 2019. <https://arxiv.org/abs/1911.04052>.

- [116] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning*, 2021. <https://arxiv.org/abs/2108.03298>.
- [117] Jiageng Mao, Sicheng He, Hao-Ning Wu, Yang You, Shuyang Sun, Zhicheng Wang, Yanan Bao, Huizhong Chen, Leonidas Guibas, Vitor Guizilini, Howard Zhou, and Yue Wang. Robot learning from a physical world model, 2025. <https://arxiv.org/abs/2511.07416>.
- [118] Russell Mendonca, Shikhar Bahl, and Deepak Pathak. Structured world models from human videos, 2023. <https://arxiv.org/abs/2308.10901>.
- [119] Weishi Mi, Yong Bao, Xiaowei Chi, Xiaozhu Ju, Zhiyuan Qin, Kuangzhi Ge, Kai Tang, Peidong Jia, Shanghang Zhang, and Jian Tang. Te-idm: Grounding video generation for executable zero-shot robot motion, 2026. <https://arxiv.org/abs/2601.18323>.
- [120] Sarthak Mishra, Rishabh Dev Yadav, Naveen Nair, Wei Pan, and Spandan Roy. Aeroplace-flow: Language-grounded object placement for aerial manipulators via visual foresight and object flow, 2026. <https://arxiv.org/abs/2603.07744>.
- [121] MotuBrain Team, Chendong Xiang, Fan Bao, Haitian Liu, Hengkai Tan, Hongzhe Bi, James Li, Jiabao Liu, Jingrui Pang, Kiro Jing, Louis Liu, Mengchen Cai, Rongxu Cui, Ruowen Zhao, Runqing Wang, Shuhe Huang, Yao Feng, Yinze Rong, Zeyuan Wang, and Jun Zhu. Motubrain: An advanced world action model for robot control, 2026. <https://arxiv.org/abs/2604.27792>.
- [122] Yao Mu, Tianxing Chen, Zanzin Chen, Shijia Peng, Zhiqian Lan, Zeyu Gao, Zhixuan Liang, Qiaojun Yu, Yude Zou, Mingkun Xu, Lunkai Lin, Zhiqiang Xie, Mingyu Ding, and Ping Luo. Robotwin: Dual-arm robot benchmark with generative digital twins, 2025. <https://arxiv.org/abs/2504.13059>.
- [123] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots, 2024. <https://arxiv.org/abs/2406.02523>.
- [124] Toan Nguyen, Weiduo Yuan, Songlin Wei, Hui Li, Daniel Seita, and Yue Wang. Iclr: In-context imitation learning with visual reasoning, 2026. <https://arxiv.org/abs/2603.07530>.
- [125] Zhenyang Ni, Yijiang Li, Ruochen Jiao, Simon Sinong Zhan, Sipeng Chen, Zhenfei Yin, Minshuo Chen, Philip Torr, Zhaoran Wang, and Qi Zhu. Creflow: Corrective reflow for sparse-reward embodied video diffusion rl, 2026. <https://arxiv.org/abs/2605.14274>.
- [126] Sangjun Noh, Dongwoo Nam, Kangmin Kim, Geonhyup Lee, Yeonguk Yu, Raeyoung Kang, and Kyoobin Lee. 3d flow diffusion policy: Visuomotor policy learning via generating flow in 3d space, 2025. <https://arxiv.org/abs/2509.18676>.
- [127] NVIDIA. NVIDIA Isaac Sim. <https://developer.nvidia.com/isaac/sim>. NVIDIA Developer. Accessed: 2026-04-14.
- [128] NVIDIA. Cosmos-predict2: World simulation model for physical ai, 2025. <https://github.com/nvidia-cosmos/cosmos-predict2>.
- [129] NVIDIA, Hassan Abu Alhaja, Jose Alvarez, Maciej Bala, Tiffany Cai, Tianshi Cao, Liz Cha, Joshua Chen, Mike Chen, Francesco Ferroni, Sanja Fidler, Dieter Fox, Yunhao Ge, Jinwei Gu, Ali Hassani, Michael Isaev, Pooya Jannaty, Shiyi Lan, Tobias Lasser, Huan Ling, Ming-Yu Liu, Xian Liu, Yifan Lu, Alice Luo, Qianli Ma, Hanzi Mao, Fabio Ramos, Xuanchi Ren, Tianchang Shen, Xinglong Sun, Shitao Tang, Ting-Chun Wang, Jay Wu, Jiashu Xu, Stella Xu, Kevin Xie, Yuchong Ye, Xiaodong Yang, Xiaohui Zeng, and Yu Zeng. Cosmos-transfer1: Conditional world generation with adaptive multimodal control, 2025. <https://arxiv.org/abs/2503.14492>.
- [130] OpenAI. Video generation models as world simulators. <https://openai.com/index/video-generation-models-as-world-simulators/>, 2024. Accessed: 2026-03-31.
- [131] OpenAI. Sora 2. <https://openai.com/sora>, 09 2025. Accessed: 2026-04-08.
- [132] Jonas Pai, Liam Achenbach, Victoriano Montesinos, Benedek Forrai, Oier Mees, and Elvis Nava. mimic-video: Video-action models for generalizable robot control beyond vlas. *arXiv preprint arXiv:2512.15692*, 2025. <https://arxiv.org/abs/2512.15692>.

- [133] Shivansh Patel, Shraddhaa Mohan, Hanlin Mai, Unnat Jain, Svetlana Lazebnik, and Yunzhu Li. Robotic manipulation by imitating generated videos without physical demonstrations. *arXiv preprint arXiv:2507.00990*, 2025. <https://arxiv.org/abs/2507.00990>.
- [134] Toby Perrett, Ahmad Darkhalil, Saptarshi Sinha, Omar Emara, Sam Pollard, Kranti Parida, Kaiting Liu, Prajwal Gatti, Siddhant Bansal, Kevin Flanagan, Jacob Chalk, Zhifan Zhu, Rhodri Guerrier, Fahd Abdelazim, Bin Zhu, Davide Moltisanti, Michael Wray, Hazel Doughty, and Dima Damen. Hd-epic: A highly-detailed egocentric video dataset, 2025. <https://arxiv.org/abs/2502.04144>.
- [135] Ryan Punamiya, Simar Kareer, Zeyi Liu, Josh Citron, Ri-Zhao Qiu, Xiongyi Cai, Alexey Gavryushin, Jiaqi Chen, Davide Liconti, Lawrence Y. Zhu, Patcharapong Aphiwetsa, Baoyu Li, Aniketh Cheluva, Pranav Kuppili, Yangcen Liu, Dhruv Patel, Aidan Gao, Hye-Young Chung, Ryan Co, Renee Zbizika, Jeff Liu, Xiaomeng Xu, Haoyu Xiong, Geng Chen, Sebastiano Oliani, Chenyu Yang, Xi Wang, James Fort, Richard Newcombe, Josh Gao, Jason Chong, Garrett Matsuda, Aseem Doriwala, Marc Pollefeys, Robert Katzschmann, Xiaolong Wang, Shuran Song, Judy Hoffman, and Danfei Xu. Egoverse: An egocentric human dataset for robot learning from around the world, 2026. <https://arxiv.org/abs/2604.07607>.
- [136] Yiran Qin, Zhelun Shi, Jiwen Yu, Xijun Wang, Enshen Zhou, Lijun Li, Zhenfei Yin, Xihui Liu, Lu Sheng, Jing Shao, Lei Bai, Wanli Ouyang, and Ruimao Zhang. Worldsimbench: Towards video generation models as world simulators, 2024. <https://arxiv.org/abs/2410.18072>.
- [137] Eric Rohmer, Surya Singh, and Marc Freese. V-rep: A versatile and scalable robot simulation framework. pages 1321–1326, 11 2013. doi: 10.1109/IROS.2013.6696520.
- [138] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. <https://arxiv.org/abs/1505.04597>.
- [139] Yu Shang, Xin Zhang, Yinzhou Tang, Lei Jin, Chen Gao, Wei Wu, and Yong Li. Roboscape: Physics-informed embodied world model, 2025. <https://arxiv.org/abs/2506.23135>.
- [140] Yichao Shen, Fangyun Wei, Zhiying Du, Yaobo Liang, Yan Lu, Jiaolong Yang, Nanning Zheng, and Baining Guo. Videovla: Video generators can be generalizable robot manipulators, 2025. <https://arxiv.org/abs/2512.06963>.
- [141] Chen Shi, Jinrui Xu, Shaoshuai Shi, Kehua Sheng, Bo Zhang, and Li Jiang. Drivewam: Video generative priors enable scalable world-action modeling for autonomous driving, 2026. <https://arxiv.org/abs/2605.28544>.
- [142] Bronislav Sidik and Dror Mizrahi. 3d-anchored lookahead planning for persistent robotic scene memory via world-model-based mcts, 2026. <https://arxiv.org/abs/2604.11302>.
- [143] Yiren Song, Yihan Wang, Xiyao Deng, Zhuoran Yan, and Mike Zheng Shou. Sweet: Sparse world modeling with image editing for embodied task execution, 2026. <https://arxiv.org/abs/2605.19319>.
- [144] Zijian Song, Qichang Li, Sihan Qin, Yuhao Chen, Tianshui Chen, Liang Lin, and Guangrun Wang. Learning physics from pretrained video models: A multimodal continuous and sequential world interaction models for robotic manipulation, 2026. <https://arxiv.org/abs/2603.00110>.
- [145] Jingwen Sun, Wenyao Zhang, Zekun Qi, Shaojie Ren, Zezhi Liu, Hanxin Zhu, Guangzhong Sun, Xin Jin, and Zhibo Chen. Vla-jepa: Enhancing vision-language-action model with latent world model, 2026. <https://arxiv.org/abs/2602.10098>.
- [146] Yirui Sun, Guangyu Zhuge, Keliang Liu, Jie Gu, Xinyu Bing, Zhongxue Gan, and Chunxu Tian. Sants: A state-adaptive scheduler for world action models, 2026. <https://arxiv.org/abs/2605.27947>.
- [147] Chao Tang, Jiacheng Xu, Haofei Lu, Bolin Zou, Wenlong Dong, Hong Zhang, and Danica Kragic. Grasp as you dream: Imitating functional grasping from generated human demonstrations, 2026. <https://arxiv.org/abs/2604.07517>.
- [148] Yinzhou Tang, Yu Shang, Yinuo Chen, Bingwen Wei, Xin Zhang, Shu’ang Yu, Liangzhi Shi, Chao Yu, Chen Gao, Wei Wu, and Yong Li. Roboscape-r: Unified reward-observation world models for generalizable robotics training via rl, 2025. <https://arxiv.org/abs/2512.03556>.
- [149] Zuojin Tang, Haoyun Liu, Xinyuan Chang, Changjie Wu, Dongjie Huo, Yandan Yang, Bin Liu, Zhejia Cai, Feng Xiong, Mu Xu, Jiachen Luo, De Ma, Zhiheng Ma, and Gang Pan. ALAM: Algebraically consistent latent action model for vision-language-action models, 2026. <https://arxiv.org/abs/2605.10819>. arXiv:2605.10819.
- [150] Stone Tao, Fanbo Xiang, Arth Shukla, Yuzhe Qin, Xander Hinrichsen, Xiaodi Yuan, Chen Bao, Xinsong Lin, Yulin Liu, Tse-kai Chan, Yuan Gao, Xuanlin Li, Tongzhou Mu, Nan Xiao, Arnav Gurha, Viswesh Nagaswamy

- Rajesh, Yong Woo Choi, Yen-Ru Chen, Zhiao Huang, Roberto Calandra, Rui Chen, Shan Luo, and Hao Su. ManiSkill3: GPU parallelized robotics simulation and rendering for generalizable embodied AI. *arXiv preprint arXiv:2410.00425*, 2024. <https://arxiv.org/abs/2410.00425>.
- [151] Gemini Robotics Team, Krzysztof Choromanski, Coline Devin, Yilun Du, Debidatta Dwivedi, Ruiqi Gao, Abhishek Jindal, Thomas Kipf, Sean Kirmani, Isabel Leal, Fangchen Liu, Anirudha Majumdar, Andrew Marmon, Carolina Parada, Yulia Rubanova, Dhruv Shah, Vikas Sindhwani, Jie Tan, Fei Xia, Ted Xiao, Sherry Yang, Wenhao Yu, and Allan Zhou. Evaluating gemini robotics policies in a veo world simulator, 2025. <https://arxiv.org/abs/2512.10675>.
- [152] Danny Tran, Roberto Martín-Martín, and Kristen Grauman. Egoexo-wm: Unlocking exo video for ego world models, 2026. <https://arxiv.org/abs/2605.15477>.
- [153] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation, 2017. <https://arxiv.org/abs/1707.04993>.
- [154] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.
- [155] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017. <https://arxiv.org/abs/1711.00937>.
- [156] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenten Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025. <https://arxiv.org/abs/2503.20314>.
- [157] Boyang Wang, Nikhil Sridhar, Chao Feng, Mark Van der Merwe, Adam Fishman, Nima Fazeli, and Jeong Joon Park. This&that: Language-gesture controlled video generation for robot planning. *arXiv preprint arXiv:2407.05530*, 2024. <https://arxiv.org/abs/2407.05530>.
- [158] Boyuan Wang, Bohan Li, Chaojun Ni, Guan Huang, Guosheng Zhao, Hao Li, Jie Li, Jindi Lv, Jingyu Liu, Lv Feng, Mingming Yu, Peng Li, Qiuping Deng, Tianze Liu, Xinyu Zhou, Xinze Chen, Xiaofeng Wang, Yang Wang, Yifan Li, Yifei Nie, Yilong Li, Yukun Zhou, Yun Ye, Zhichao Liu, and Zheng Zhu. Gigabrain-0.5m*: a vla that learns from world model-based reinforcement learning, 2026. <https://arxiv.org/abs/2602.12099>.
- [159] Guodong Wang, Chenkai Zhang, Qingjie Liu, Jinjin Zhang, Jiancheng Cai, Junjie Liu, and Xinmin Liu. LIBERO-X: Robustness litmus for vision-language-action models. *arXiv preprint arXiv:2602.06556*, 2026. <https://arxiv.org/abs/2602.06556>.
- [160] Jiaxu Wang, Yicheng Jiang, Tianlun He, Jingkai Sun, Qiang Zhang, Junhao He, Jiahang Cao, Zesen Gan, Mingyuan Sun, Qiming Shao, and Xiangyu Yue. MVISTA-4D: View-consistent 4d world model with test-time action inference for robotic manipulation. *arXiv preprint arXiv:2602.09878*, 2026. <https://arxiv.org/abs/2602.09878>.
- [161] Rui Wang, Yue Zhang, Jiehong Lin, Kuncheng Luo, Jianan Wang, Zhongrui Wang, and Xiaojuan Qi. When to trust imagination: Adaptive action execution for world action models, 2026. <https://arxiv.org/abs/2605.06222>.
- [162] Ruixiang Wang, Qingming Liu, Yueci Deng, Guiliang Liu, Zhen Liu, and Kui Jia. Eva: Aligning video world models with executable robot actions via inverse dynamics rewards, 2026. <https://arxiv.org/abs/2603.17808>.
- [163] Xiaofeng Wang, Kang Zhao, Feng Liu, Jiayu Wang, Guosheng Zhao, Xiaoyi Bao, Zheng Zhu, Yingya Zhang, and Xingang Wang. Egovid-5m: A large-scale video-action dataset for egocentric video generation, 2024. <https://arxiv.org/abs/2411.08380>.
- [164] Xinkai Wang, Chenyi Wang, Yifu Xu, Mingzhe Ye, Fu-Cheng Zhang, Jialin Tian, Xinyu Zhan, Lifeng Zhu, Cewu Lu, and Lixin Yang. Lamp: Learning vision-language-action policies with 3d scene flow as latent motion prior, 2026. <https://arxiv.org/abs/2603.25399>.

- [165] Yi Ru Wang, Carter Ung, Christopher Tan, Grant Tannert, Jiafei Duan, Josephine Li, Anh Le, Rishabh Oswal, Markus Grotz, Wilbert Pumacay, Yuquan Deng, Ranjay Krishna, Dieter Fox, and Siddhartha Srinivasa. RoboEval: Where robotic manipulation meets structured and scalable evaluation. *arXiv preprint arXiv:2507.00435*, 2025. <https://arxiv.org/abs/2507.00435>.
- [166] Yixuan Wang, Rhythm Syed, Fangyu Wu, Mengchao Zhang, Aykut Onol, Jose Barreiros, Hooshang Nayyeri, Tony Dear, Huan Zhang, and Yunzhu Li. Interactive world simulator for robot policy training and evaluation, 2026. <https://arxiv.org/abs/2603.08546>.
- [167] Yuqi Wang, Xinghang Li, Wenxuan Wang, Junbo Zhang, Yingyan Li, Yuntao Chen, Xinlong Wang, and Zhaoxiang Zhang. Unified vision-language-action model, 2025.
- [168] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [169] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. <https://arxiv.org/abs/2312.08344>.
- [170] Youpeng Wen, Junfan Lin, Yi Zhu, Jianhua Han, Hang Xu, Shen Zhao, and Xiaodan Liang. Vidman: Exploiting implicit dynamics from video diffusion model for effective robot manipulation, 2024.
- [171] John Won, Kyungmin Lee, Huiwon Jang, Dongyoung Kim, and Jinwoo Shin. Dual-stream diffusion for world-model augmented vision-language-action model, 2025. <https://arxiv.org/abs/2510.27607>.
- [172] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation, 2023.
- [173] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. <https://openreview.net/forum?id=NxoFmGgWC9>.
- [174] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. ivideoGPT: Interactive videoGPTs are scalable world models. In *Advances in Neural Information Processing Systems*, 2024. <https://arxiv.org/abs/2405.15223>.
- [175] Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, Shichao Fan, Xinhua Wang, Fei Liao, Zhen Zhao, Guangyu Li, Zhao Jin, Lecheng Wang, Jilei Mao, Ning Liu, Pei Ren, Qiang Zhang, Yaoxu Lyu, Mengzhen Liu, Jingyang He, Yulin Luo, Zeyu Gao, Chenxuan Li, Chenyang Gu, Yankai Fu, Di Wu, Xingyu Wang, Sixiang Chen, Zhenyu Wang, Pengju An, Siyuan Qian, Shanghang Zhang, and Jian Tang. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. In *Robotics: Science and Systems XXI*, RSS2025. Robotics: Science and Systems Foundation, June 2025. doi: 10.15607/rss.2025.xxi.152. <http://dx.doi.org/10.15607/RSS.2025.XXI.152>.
- [176] Sicheng Xie, Haidong Cao, Zejia Weng, Zhen Xing, Haoran Chen, Shiwei Shen, Jiaqi Leng, Zuxuan Wu, and Yu-Gang Jiang. Human2robot: Learning robot actions from paired human-robot videos, 2025. <https://arxiv.org/abs/2502.16587>.
- [177] Mengda Xu, Zhenjia Xu, Yinghao Xu, Cheng Chi, Gordon Wetzstein, Manuela Veloso, and Shuran Song. Flow as the cross-domain manipulation interface. In *8th Annual Conference on Robot Learning*, 2024. <https://arxiv.org/abs/2407.15208>.
- [178] Zhengtong Xu, Qiang Qiu, and Yu She. VILP: Imitation learning with latent video planning. *IEEE Robotics and Automation Letters*, 10(4):3350–3357, 2025. <https://arxiv.org/abs/2502.01784>.
- [179] Adina Yakefu, Bin Xie, Chongyang Xu, Enwen Zhang, Erjin Zhou, Fan Jia, Haitao Yang, Haoqiang Fan, Haowei Zhang, Hongyang Peng, Jing Tan, Junwen Huang, Kai Liu, Kaixin Liu, Kefan Gu, Qinglun Zhang, Ruitao Zhang, Saike Huang, Shen Cheng, Shuaicheng Liu, Tiancai Wang, Tiezhen Wang, Wei Sun, Wenbin Tang, Yajun Wei, Yang Chen, Youqiang Gui, Yucheng Zhao, Yunchao Ma, Yunfei Wei, Yunhuan Yang, Yutong Guo, Ze Chen, Zhengyuan Du, Ziheng Zhang, Ziming Liu, and Ziwei Yan. RoboChallenge: Large-scale real-robot evaluation of embodied policies. *arXiv preprint arXiv:2510.17950*, 2025. <https://arxiv.org/abs/2510.17950>.
- [180] Feng Yan, Fanfan Liu, Liming Zheng, Yufeng Zhong, Yiyang Huang, Zechao Guan, Chengjian Feng, and Lin Ma. Robotron-mani: All-in-one multimodal large model for robotic manipulation, 2025. <https://arxiv.org/abs/2412.07215>.

- [181] Haodong Yan, Zhide Zhong, Jiaguan Zhu, Junjie He, Weilin Yuan, Wenxuan Song, Xin Gong, Yingjie Cai, Guanyi Zhao, Xu Yan, Bingbing Liu, Ying-Cong Chen, and Haoang Li. S-vam: Shortcut video-action model by self-distilling geometric and semantic foresight, 2026. <https://arxiv.org/abs/2603.16195>.
- [182] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers, 2021. <https://arxiv.org/abs/2104.10157>.
- [183] Liudi Yang, Yang Bai, George Eskandar, Fengyi Shen, Mohammad Altillawi, Dong Chen, Ziyuan Liu, and Abhinav Valada. Covar: Co-generation of video and action for robotic manipulation via multi-modal diffusion, 2025. <https://arxiv.org/abs/2512.16023>.
- [184] Liudi Yang, Yang Bai, George Eskandar, Fengyi Shen, Mohammad Altillawi, Dong Chen, Soumajit Majumder, Ziyuan Liu, Gitta Kutyniok, and Abhinav Valada. Roboenvision: A long-horizon video generation model for multi-task robot manipulation. *arXiv preprint arXiv:2506.22007*, 2025. <https://arxiv.org/abs/2506.22007>.
- [185] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan Zhang, Weihao Wang, Yean Cheng, Bin Xu, Xiaotao Gu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer, 2025. <https://arxiv.org/abs/2408.06072>.
- [186] Angen Ye, Boyuan Wang, Chaojun Ni, Guan Huang, Guosheng Zhao, Hao Li, Hengtao Li, Jie Li, Jindi Lv, Jingyu Liu, Min Cao, Peng Li, Qiuping Deng, Wenjun Mei, Xiaofeng Wang, Xinze Chen, Xinyu Zhou, Yang Wang, Yifan Chang, Yifan Li, Yukun Zhou, Yun Ye, Zhichao Liu, and Zheng Zhu. Gigaworld-policy: An efficient action-centered world-action model, 2026. <https://arxiv.org/abs/2603.17240>.
- [187] Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Se June Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Liden, Kimin Lee, Jianfeng Gao, Luke Zettlemoyer, Dieter Fox, and Minjoon Seo. Latent action pretraining from videos. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. <https://openreview.net/forum?id=VYOe2eBQeh>.
- [188] Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, Ayaan Malik, Kyungmin Lee, William Liang, Nadun Ranawaka, Jiasheng Gu, Yinzhen Xu, Guanzhi Wang, Fengyuan Hu, Avnish Narayan, Johan Bjorck, Jing Wang, Gwanghyun Kim, Dantong Niu, Ruijie Zheng, Yuqi Xie, Jimmy Wu, Qi Wang, Ryan Julian, Danfei Xu, Yilun Du, Yevgen Chebotar, Scott Reed, Jan Kautz, Yuke Zhu, Linxi “Jim” Fan, and Joel Jang. World action models are zero-shot policies, 2026. <https://arxiv.org/abs/2602.15922>.
- [189] Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, Zsolt Kira, Manolis Savva, Angel Chang, Devendra Singh Chaplot, Dhruv Batra, Roozbeh Mottaghi, Yonatan Bisk, and Chris Paxton. HomeRobot: Open-vocabulary mobile manipulation. In *Conference on Robot Learning*, 2023. <https://arxiv.org/abs/2306.11565>.
- [190] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Avnish Narayan, Hayden Shively, Adithya Bellathur, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-World: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, 2019. <https://arxiv.org/abs/1910.10897>.
- [191] Ge Yuan, Qiyuan Qiao, Jing Zhang, and Dong Xu. Adaworldpolicy: World-model-driven diffusion policy with online adaptive learning for robotic manipulation, 2026. <https://arxiv.org/abs/2602.20057>.
- [192] Haoran Yuan, Weigang Yi, Zhenyu Zhang, Wendi Chen, Yuchen Mo, Jiashi Yin, Xinzhuo Li, Xiangyu Zeng, Chuan Wen, Cewu Lu, Katherine Driggs-Campbell, and Ismini Lourentzou. Vtam: Video-tactile-action models for complex physical interaction beyond vlas, 2026. <https://arxiv.org/abs/2603.23481>.
- [193] Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-wam: Do world action models need test-time future imagination?, 2026. <https://arxiv.org/abs/2603.16666>.
- [194] Fan Zhang and Michael Gienger. Learning robot manipulation from audio world models, 2025. <https://arxiv.org/abs/2512.08405>.
- [195] Gehao Zhang, Zhenyang Ni, Payal Mohapatra, Han Liu, Ruohan Zhang, and Qi Zhu. Emboalign: Aligning video generation with compositional constraints for zero-shot manipulation, 2026. <https://arxiv.org/abs/2603.05757>.
- [196] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.

- [197] Shiduo Zhang, Zhe Xu, Peiju Liu, Xiaopeng Yu, Yuan Li, Qinghui Gao, Zhaoye Fei, Zhangyue Yin, Zuxuan Wu, Yu-Gang Jiang, and Xipeng Qiu. VLABench: A large-scale benchmark for language-conditioned robotics manipulation with long-horizon reasoning tasks, 2024. <https://arxiv.org/abs/2412.18194>.
- [198] Yi Zhang, Yinda Chen, Che Liu, Zeyuan Ding, Jin Xu, Shilong Zou, Junwei Liao, Jiayu Hu, Xiancong Ren, Xiaopeng Zhang, et al. Pelican-unify 1.0: A unified embodied intelligence model for understanding, reasoning, imagination and action, 2026. <https://arxiv.org/abs/2605.15153>.
- [199] Zhongru Zhang, Chenghan Yang, Qingzhou Lu, Yanjiang Guo, Jianke Zhang, Yucheng Hu, and Jianyu Chen. Veo-act: How far can frontier video models advance generalizable robot manipulation?, 2026. <https://arxiv.org/abs/2604.04502>.
- [200] Han Zhao, Jingbo Wang, Wenxuan Song, Shuai Chen, Yang Liu, Yan Wang, Haoang Li, and Donglin Wang. Frappe: Infusing world modeling into generalist policies via multiple future representation alignment, 2026. <https://arxiv.org/abs/2602.17259>.
- [201] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, Ankur Handa, Tsung-Yi Lin, Gordon Wetzstein, Ming-Yu Liu, and Donglai Xiang. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 1702–1713. Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.00166. https://openaccess.thecvf.com/content/CVPR2025/html/Zhao_CoT-VLA_Visual_Chain-of-Thought_Reasoning_for_Vision-Language-Action_Models_CVPR_2025_paper.html.
- [202] Haoyu Zhen, Qiao Sun, Hongxin Zhang, Junyan Li, Siyuan Zhou, Yilun Du, and Chuang Gan. Tesseract: Learning 4d embodied world models, 2025. <https://arxiv.org/abs/2504.20995>.
- [203] Haoyu Zhen, Zixian Gao, Qiao Sun, Yilin Zhao, Yuncong Yang, Yilun Du, Pengsheng Guo, Tsun-Hsuan Wang, Yi-Ling Qiao, and Chuang Gan. Action images: End-to-end policy learning via multiview video generation, 2026. <https://arxiv.org/abs/2604.06168>. arXiv:2604.06168.
- [204] Ruijie Zheng, Jing Wang, Scott Reed, Johan Bjorck, Yu Fang, Fengyuan Hu, Joel Jang, Kaushil Kundalia, Zongyu Lin, Loic Magne, Avnish Narayan, You Liang Tan, Guanzhi Wang, Qi Wang, Jiannan Xiang, Yinzhen Xu, Seonghyeon Ye, Jan Kautz, Furong Huang, Yuke Zhu, and Linxi Fan. Flare: Robot learning with implicit world modeling, 2025. <https://arxiv.org/abs/2505.15659>.
- [205] Ruijie Zheng, Dantong Niu, Yuqi Xie, Jing Wang, Mengda Xu, Yunfan Jiang, Fernando Castañeda, Fengyuan Hu, You Liang Tan, Letian Fu, Trevor Darrell, Furong Huang, Yuke Zhu, Danfei Xu, and Linxi Fan. Egoscale: Scaling dexterous manipulation with diverse egocentric human data. *arXiv preprint arXiv:2602.16710*, 2026.
- [206] Yuhang Zheng, Songen Gu, Weize Li, Yupeng Zheng, Yujie Zang, Shuai Tian, Xiang Li, Ce Hao, Chen Gao, Si Liu, Haoran Li, Yilun Chen, Shuicheng Yan, and Wenchao Ding. Omnivta: Visuo-tactile world modeling for contact-rich robotic manipulation, 2026. <https://arxiv.org/abs/2603.19201>.
- [207] Hongyan Zhi, Peihao Chen, Siyuan Zhou, Yubo Dong, Quanxi Wu, Lei Han, and Mingkui Tan. 3DFlowAction: Learning cross-embodiment manipulation from 3d flow world model. *arXiv preprint arXiv:2506.06199*, 2025. <https://arxiv.org/abs/2506.06199>.
- [208] Pengfei Zhou, Liliang Chen, Shengcong Chen, Di Chen, Wenzhi Zhao, Rongjun Jin, Guanghui Ren, and Jianlan Luo. Act2goal: From world model to general goal-conditioned policy, 2025. <https://arxiv.org/abs/2512.23541>.
- [209] Pengfei Zhou, Shengcong Chen, Di Chen, Jiaxu Wang, Rongjun Jin, Bingwen Zhu, Yike Pan, Songen Gu, Kuanning Wang, Shufeng Nan, Xingyu Qiu, Chenhao Qiu, Pu Yang, Yunuo Cai, Jianxiong Gao, Yifan Li, Yanwei Fu, Xiangyu Yue, Zhi Chen, and Jianlan Luo. τ_0 -WM: A unified video-action world model for robotic manipulation, 2026. <https://arxiv.org/abs/2606.01027>.
- [210] Siyuan Zhou, Yilun Du, Jiaben Chen, Yandong Li, Dit-Yan Yeung, and Chuang Gan. Robodreamer: Learning compositional world models for robot imagination, 2024. <https://arxiv.org/abs/2404.12377>.
- [211] Xueyang Zhou, Yangming Xu, Guiyao Tie, Yongchao Chen, Guowen Zhang, Duanfeng Chu, Pan Zhou, and Lichao Sun. LIBERO-PRO: Towards robust and fair evaluation of vision-language-action models beyond memorization. *arXiv preprint arXiv:2510.03827*, 2025. <https://arxiv.org/abs/2510.03827>.
- [212] Yang Zhou, Xiaofeng Wang, Hao Shao, Letian Wang, Guosheng Zhao, Jiangnan Shao, Jiagang Zhu, Tingdong Yu, Zheng Zhu, Guan Huang, and Steven L. Waslander. Drivedreamer-policy: A geometry-grounded world-action model for unified generation and planning, 2026. <https://arxiv.org/abs/2604.01765>.

- [213] Chuning Zhu, Raymond Yu, Siyuan Feng, Benjamin Burchfiel, Paarth Shah, and Abhishek Gupta. Unified world models: Coupling video and action diffusion for pretraining on large robotic datasets, 2025. <https://arxiv.org/abs/2504.02792>.
- [214] Fangqi Zhu, Hongtao Wu, Song Guo, Yuxiao Liu, Chilam Cheang, and Tao Kong. Irasim: A fine-grained world model for robot manipulation, 2025. <https://arxiv.org/abs/2406.14540>.
- [215] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Kevin Lin, Abhiram Maddukuri, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning, 2020. <https://arxiv.org/abs/2009.12293>.