

MotuBrain: An Advanced World Action Model for Robot Control

MotuBrain Team

1 Introduction

Recent progress in embodied foundation models has been driven by Vision-Language-Action (VLA) policies [6, 8, 10, 20, 22, 26, 33, 38], which map visual observations and language instructions in pretrained Vision-Language Models (VLMs) [1, 4] directly to robot actions. Inheriting rich semantic priors, VLAs generalize well across diverse objects and language instructions, achieving strong performance on a wide range of robotic tasks. However, since they are primarily pretrained on static image-text data, they often neglect the perception and prediction of fine-grained world dynamics that are essential for precise robotic control, leading to a superficial imitation of behaviors rather than temporal understanding of the physics of the world.

With the rise of video generation models [3, 23, 30, 34, 39], a growing body of research has begun exploring how to adapt these models for world modeling [9, 15–18, 28]. A world model aims to predict how the environment evolves in response to actions—a capability that aligns directly with what a video generation model does when conditioned on past observations and actions, namely forecasting future visual states. Conceptually, video generation models are naturally suited to this task, as video generation models are pretrained on vast and diverse web video data, which equips them with rich spatiotemporal priors regarding object permanence, physical dynamics, and human-object interactions. This enables them to generalize more effectively and reason about novel scenarios where in-domain data are scarce. As a result, leveraging video generation models as a foundation for world modeling has emerged as a promising paradigm for scalable robot policy learning.

Following this insight, early attempts largely followed a two-stage paradigm: Video Generation Model (VGM) plus Inverse Dynamics Model (IDM) [13, 14, 19, 29, 40]. In this framework, a video diffusion model pretrained on large-scale web data first predicts future visual trajectories from current observations and language instructions. Subsequently, an inverse dynamics model infers actions from the generated future frames. While this paradigm successfully leverages rich spatiotemporal priors from video data to achieve broad generalization, it suffers from a critical drawback in that errors in video prediction accumulate over time, which leads to compromised action accuracy and downstream policy performance.

To mitigate this issue, a subsequent line of research has explored World Action Models (WAMs) that unify visual dynamics and action prediction within a single, jointly optimized objective [5, 21, 24, 25, 27, 36, 37, 41]. Unlike the VGM+IDM pipeline, which decouples forecasting from action inference, WAMs simultaneously predict future visual states and actions in an aligned manner. This integration offers two key advantages: (1) it avoids the cascading errors inherent in sequential video prediction, and (2) it overcomes the fragmented functionality of conventional embodied systems, where semantic understanding, dynamics modeling, and action generation are typically learned from disparate supervision sources.

In our view, the core source of intelligence in such a unified model is its ability to absorb large-scale heterogeneous multimodal data under one unified training recipe. In principle, this includes pure video data without action annotations, robot data with aligned video-language-action trajectories across different embodiments, and even task-agnostic interaction data with partially missing modalities [32]. By contrast, VLA learning primarily relies on robot task trajectories with aligned observation-language-action supervision, and adaptation to the target robot is primarily coupled to embodiment-specific action data [6, 12, 22].

Motus [5] was an early step in this direction. It established a unified world-action formulation in which video and action are modeled in a shared generative framework, so that policy modeling, world modeling,

video generation, inverse dynamics, and joint video-action prediction become different inference modes of the same model. It also showed that, with UniDiffuser-style continuous multimodal modeling and a Mixture-of-Transformers design, a world action model can absorb heterogeneous multimodal data rather than being restricted to embodiment-specific task trajectories.

Building on this foundation, we present MotuBrain. Like Motus, MotuBrain adopts UniDiffuser [2] to jointly model and schedule the two continuous modalities, namely video and action, and uses a three-stream Mixture-of-Transformers architecture to integrate video generation, action modeling, and language conditioning within one system. This unified formulation again supports inference over five distributions with the same model: vision-language-action policy modeling, world modeling, video generation, inverse dynamics, and joint video-action prediction. More importantly, it preserves the key advantage of unified world-action modeling: the model can learn from a much broader family of multimodal data, including video-only data without action labels, interaction data without explicit task language, and heterogeneous robot trajectories collected from different embodiments.

MotuBrain further extends this paradigm in several practically important directions. It introduces a unified multiview representation that supports an arbitrary number of camera views under different camera layouts, rather than depending on a fixed visual input format. It uses an independent text stream to more tightly couple high-level semantics with low-level control, making instruction following an explicit part of action generation and improving semantic understanding. It adopts a unified action representation across embodiments, enabling the model to capture transferable control regularities rather than overfitting to the action format of a single robot. Beyond architecture and pre-training, MotuBrain also involves a post-training and deployment recipe tailored to long-horizon real-world control: autoregressive diffusion rollout enables temporally extended execution, V2A-style asymmetric dependency enables action-only inference without explicitly generating future video, and real-time chunked closed-loop execution reduces boundary discontinuities during asynchronous control. Finally, we develop a systems-oriented inference stack including denoising-step reduction, compilation with CUDA-graph-friendly execution, FP8 quantization, and DiT caching, which together deliver more than $50\times$ end-to-end speedup over the naive baseline and make large world action models practical for real-time robotic deployment.

2 Method

In this section, we present the overall methodology of MotuBrain from four aspects: model architecture, pre-training, post-training, and inference. We first introduce the core architectural design and the main acceleration techniques integrated into MotuBrain in Section 2.1. We then describe the pretrained foundation model used in our system in Section 2.2. After that, we present the post-training pipeline, including sparse adaptation and step distillation, in Section 2.3. Finally, we describe the inference-time acceleration strategies deployed in MotuBrain in Section 2.4.

2.1 Model Architecture

MotuBrain first adopts UniDiffuser [2] to jointly model and schedule the two continuous modalities, i.e., video and action, so that all interaction patterns between them are captured in a unified generative framework. As a result, a single training procedure supports inference over five distributions: vision-language-action policy modeling, world modeling, video generation, inverse dynamics, and joint video-action prediction, which are formulated in Table 1.

On top of this unified video-action backbone, MotuBrain introduces a three-stream Mixture-of-Transformers (MoT) architecture with a text stream, a video stream, and an action stream. The text stream serves as a conditioning branch: its hidden states participate in Transformer attention, but no output head is applied to text tokens. Maintaining a dedicated text stream also improves semantic understanding and instruction-following capability. The video and action streams are both trained with flow matching, predicting the velocity fields of video latents and action tokens, respectively.

The inputs consist of text tokens, condition-image latents encoded by a Vidu VAE, noisy future video latents, and noisy action tokens. The condition image is represented as the first video latent frame and is teacher-forced




Figure 1 Overview of MotuBrain’s architecture. MotuBrain builds on a unified video-action backbone and adopts a three-stream Mixture-of-Transformers architecture with text, video, and action streams. It further uses an H-bridge attention design to balance cross-modal interaction and efficiency, while supporting flexible multiview inputs through view-dependent 3D RoPE offsets.

in the video stream, while the remaining video latents and action tokens are denoised by the video and action streams. Cross-modal interaction is implemented through joint attention over video, action, and text tokens.

Instead of using full video-action joint attention in all layers, we adopt an H-bridge design following HBridge [35]. Concretely, full V-A joint attention is applied only in the middle 50% of Transformer layers, while the bottom and top 25% layers use decoupled attention. In the decoupled layers, video tokens and action tokens are processed independently and do not perform joint-attention. This design reduces the cost of dense cross-modal attention, improves efficiency, preserves modality-specific representations in shallow and deep layers, and avoids injecting excessive modality-irrelevant information into every layer, while still allowing semantic alignment and policy grounding in intermediate layers.

For multiview inputs, each view is encoded independently by the Vidu VAE and concatenated at the token level. Leveraging 3D RoPE, we introduce view-dependent offsets only along the spatial dimensions while keeping the temporal dimension unchanged. This effectively maps different views to distinct regions in a shared spatial positional space, enabling seamless support for an arbitrary number of camera views without modifying the backbone architecture.

2.2 Pre-training

Our pre-training data follows a four-level data pyramid [5], organized from broad visual diversity to embodiment-specific control signals: Internet videos, ego-centric videos, heterogeneous-embodiment data, and specific-embodiment data. The main motivation is to maximize data efficiency. The bottom level uses large-scale Internet videos to train the video generation model Vidu [3], which serves as the foundation model of MotuBrain. The second level introduces ego-centric videos, which provide first-person interaction patterns

Table 1 Different prediction modes (non-autoregressive mode).

Model	Prediction Objective
VLA	$p(\mathbf{a}_{t+1:t+k} \mid \mathbf{o}_t, \ell)$
WM	$p(\mathbf{o}_{t+1:t+k} \mid \mathbf{o}_t, \mathbf{a}_{t+1:t+k})$
IDM	$p(\mathbf{a}_{t+1:t+k} \mid \mathbf{o}_{t:t+k})$
VGM	$p(\mathbf{o}_{t+1:t+k} \mid \mathbf{o}_t, \ell)$
Joint Video-Action Prediction	$p(\mathbf{o}_{t+1:t+k}, \mathbf{a}_{t+1:t+k} \mid \mathbf{o}_t, \ell)$

and hand-object dynamics that are closer to embodied manipulation. The third level uses heterogeneous-embodiment data collected from different robot platforms, tasks, and scenes. In our setting, we only use dual-arm robot data at this level. The top level consists of specific-embodiment data collected on the target robot configuration, which further aligns the model with the final action space, kinematics, camera setup, and deployment distribution. This hierarchy reflects a central design principle of MotuBrain: the model should learn from heterogeneous multimodal data wherever available, rather than restricting supervision to a single narrow data format.

Starting from the pretrained Vidu weights [3], we perform two-stage pre-training corresponding to the second and third levels of the data pyramid. In the first stage, we train only the video branch on ego-centric and heterogeneous-embodiment data, while keeping the randomly initialized action branch unchanged. Accordingly, the optimization objective in this stage contains only the video loss. This stage is designed to adapt the Internet-scale video prior to embodied manipulation and to obtain a video world model that can understand and predict bimanual interaction dynamics. To improve robustness to imperfect visual conditioning, we follow the noisy-conditioning strategy of LingBot-VA [24] throughout training, including stage 1, stage 2, and non-autoregressive post-training, but not the autoregressive policy setting. In addition, for multiview data, we randomly drop auxiliary views with probability 0.1 during pre-training, so that the model can better adapt to varying numbers of camera views and imperfect visual observations. Concretely, with probability $p = 0.5$, we perturb the conditioned-frame latent as

$$\tilde{z}_0 = s_{\text{aug}} z_0 + (1 - s_{\text{aug}}) \epsilon, \quad s_{\text{aug}} \sim \mathcal{U}[0.3, 0.7], \quad \epsilon \sim \mathcal{N}(0, I), \quad (1)$$

and leave it unchanged otherwise. In the second stage, we initialize from the first-stage checkpoint and train only the action branch on heterogeneous-embodiment data, while freezing the video branch. At this stage, we use a unified action representation across embodiments.

Concretely, let the absolute end-effector chunk be $E^{\text{abs}} = \{e_1^{\text{abs}}, \dots, e_n^{\text{abs}}\}$ and let s denote the end-effector state of the conditioned frame. We define the corresponding relative chunk as $E^{\text{rel}} = \{e_1^{\text{rel}}, \dots, e_n^{\text{rel}}\}$, where each action is represented as

$$e_i^{\text{rel}} = e_i^{\text{abs}} \ominus s. \quad (2)$$

Here $\ominus$ denotes a component-wise pose difference: the position is computed by direct subtraction, the rotation is computed by composition with the inverse reference rotation, and the gripper state is kept unchanged. Writing $e = (p, R, g)$ with position p , rotation R , and gripper state g , we have

$$e_i^{\text{abs}} = (p_i, R_i, g_i), \quad s = (p_s, R_s, g_s), \quad (3)$$

and

$$e_i^{\text{rel}} = (p_i - p_s, R_s^{-1} R_i, g_i). \quad (4)$$

The raw input pose is provided in quaternion format, while the training target uses a 6D rotation representation. Each end-effector action therefore has dimension 10, consisting of position, rotation, and gripper state. We only normalize the gripper dimension to $[-1, 1]$ and keep the remaining dimensions in their original physical scales. Using relative end-effector coordinates with respect to the conditioned frame improves compatibility across different robot embodiments and initial poses, making the action space more consistent under heterogeneous pre-training and allowing the target embodiment to be adapted with less embodiment-specific data. Although only the action branch is updated, stage 2 still optimizes both video and action objectives under the unified

formulation. We use separate SNR-based timestep sampling for the two modalities, with $\text{timeshift} = 6$ for video and $\text{timeshift} = 1$ for action. The overall stage-2 objective is the weighted sum of the video and action losses,

$$\mathcal{L} = \lambda_v \mathcal{L}_v + \lambda_a \mathcal{L}_a, \quad (5)$$

where

$$\mathcal{L}_v = \text{MSE}(v_{\text{out}}, v_{\text{target}}), \quad \mathcal{L}_a = \text{MSE}(a_{\text{out}}, a_{\text{target}}). \quad (6)$$

This design encourages alignment between the action and video modalities within the unified model: action learning benefits from the visual dynamics encoded by the video branch, while the action-conditioned interaction also improves video prediction by injecting control-relevant information into the shared modeling process.




Figure 2 Attention masks used in all training stages and post-training modes. (a) Stage-1 pre-training updates only the video branch. (b) Stage-2 pre-training uses full joint attention over language, conditioned video, video, and action tokens. (c) Non-autoregressive post-training disables video-to-action attention. (d) Autoregressive post-training applies a causal mask over temporally ordered video and action tokens.

2.3 Post-training

While pre-training aims to build a unified world action model from broad and heterogeneous data, post-training focuses on adapting the model to the target embodiment. We start from the checkpoint obtained after the second stage of pre-training and finetune it on specific-embodiment data. In this stage, we consider two post-training settings, namely a non-autoregressive (Non-AR) setting and an autoregressive (AR) setting, and train separately under each setting. The attention masks used in stage 1, stage 2, Non-AR, and AR are summarized in Figure 2. The two post-training settings mainly differ in sequence factorization and attention masking.

Under the Non-AR setting, the model denoises all video and action tokens within the full observation window in a single forward pass, where o_t denotes the observation at step t , $z_t = \mathcal{E}(o_t)$ is its VAE latent, the VAE temporally compresses every τ consecutive video frames into a single latent frame, and the action stream operates at f_{va} actions per raw video frame. Each latent frame therefore corresponds to $S_a = f_{va}\tau$ action tokens. Consequently, predicting K future video latents amounts to generating $K S_a$ actions, enabling high-frequency control while keeping the visual stream compact. As noted in Section 2.2, we continue to apply the same LingBot-VA-style noisy-conditioning augmentation to the conditioned-frame latent in Non-AR post-training. This improves robustness to noisy conditioning and helps the model recover from partially corrupted observations.

For long horizon tasks, we adopt an AR formulation based on chunk-level factorization, as illustrated in Figure 2(d). We partition each episode into non-overlapping chunks and process all chunks in parallel during training with a block-causal attention mask. The input sequence contains language tokens, conditioned-image

tokens, clean video tokens from teacher-forced past observations, and the noisy video and noisy action tokens of the current target chunk. Unlike Non-AR post-training, the AR setting does not apply the noisy-conditioning augmentation to the conditioned frame. For chunk k , tokens can attend only to clean visual context from preceding chunks, but never to future chunks. Importantly, our AR model do not involve clean action tokens, because doing so would break the unified relative-EEF action representation. At deployment, the model rolls out sequentially, using the newly observed frame as clean context for the next chunk.

To better align post-training with deployment, we adopt V2A-style attention in both operating modes. Under this attention pattern, action tokens attend to video and language tokens, while video tokens never attend to action tokens. Combined with the UniDiffuser formulation, where video and action are sampled with independent timesteps, this asymmetric dependency makes it possible to use an action-only suffix during inference: after a short joint denoising prefix, the video stream can be frozen and only the action stream continues to update while attending to the cached visual-language context. This design is motivated by efficiency rather than a change in training objective, and in practice it enables substantial inference acceleration without degrading task success rate.

2.4 Inference

2.4.1 Inference Optimization

A practical concern with world-action models is inference latency: jointly denoising a high-dimensional video latent and an action latent over many diffusion steps couples two of the slowest computations in modern generative modeling, and naively deployed WAMs typically run far below the control rates required for high frequency manipulation. We adopt and extend a stack of optimizations to tackle this problem. Crucially, we verify on the RoboTwin2.0 that this speedup is essentially lossless: average task success rates fluctuate within sub-percent margins across the optimized and unoptimized configurations, indicating that the gains come from removing genuinely redundant computation rather than from sacrificing model fidelity. We detail each component below.

Noise sampling. We use separate SNR-based timestep sampling for the two modalities during training. Specifically, we set $\text{timeshift} = 6$ for video and $\text{timeshift} = 1$ for action. This makes video timesteps more likely to be sampled from noisier regions, while action timesteps are sampled more uniformly. As a result, the model becomes more robust at predicting accurate actions under noisy visual conditions. This improves both robustness and convergence, allowing the number of inference steps to be reduced from 50 to 30 without performance degradation.

Compiling. We optimize the inference graph and fuse the operators with `torch.compile` to reduce the overhead of repeated denoising. Since the model is rewritten as a single-GPU, pure-PyTorch inference model, the main Transformer computation can be compiled directly at inference time. This optimization primarily improves execution efficiency for the repeated DiT forwards used during sampling.

DiT Cache. We adopt a DreamZero-style DiT caching strategy [36] to exploit temporal redundancy across denoising steps. Let v_t denote the velocity prediction at denoising step t . We measure the similarity between two consecutive predictions by

$$s_t = \frac{\langle v_t, v_{t-1} \rangle}{\|v_t\|_2 \|v_{t-1}\|_2}. \quad (7)$$

When $s_t > \gamma$, where γ is a predefined threshold, we skip a small number of subsequent DiT evaluations and approximate the skipped predictions from recent history, i.e.,

$$\hat{v}_{t+j} \approx v_t, \quad j = 1, \dots, k, \quad (8)$$

where k is the cache length. The cache can be applied to either the video velocity or, in action-only mode, the action velocity, and is reset for each inference call or chunk.

FP8 Quantization. We further reduce DiT inference cost with FP8 quantization. The implementation replaces eligible `nn.Linear` layers with FP8 linear layers. We store the weights in `float8_e4m3fn` with a per-tensor scale, dynamically quantize activations to FP8 at runtime, and compute matrix multiplications through `torch._scaled_mm`, returning outputs in the original compute dtype. Layers whose input or output dimensions are not divisible by 16 are skipped to satisfy the kernel alignment requirement. This optimization mainly targets the large linear projections in attention and MLP blocks, reducing memory bandwidth and GEMM cost on FP8-capable GPUs. Since quantization is applied after checkpoint loading and before compilation, the compiled graph traces the quantized linear operators directly.

V2A-Style Inference. For models trained with V2A attention, the video stream is prevented from attending to action tokens, while the action stream can still attend to visual and language context. This asymmetric dependency enables an efficient inference schedule. Instead of denoising video and action jointly for all sampling steps, we use a short joint denoising prefix followed by an action-only suffix:

$$(z_v^{(t+1)}, z_a^{(t+1)}) = \begin{cases} \Phi_{\text{joint}}(z_v^{(t)}, z_a^{(t)}), & t < N, \\ (z_v^{(N)}, \Phi_{\text{act}}(z_a^{(t)}; z_v^{(N)})), & t \geq N, \end{cases} \quad (9)$$

where $z_v^{(t)}$ and $z_a^{(t)}$ denote the video and action latents at denoising step t , respectively. After step N , the video latent is frozen, and the video-text branch is executed once to build per-layer cached keys and values for the fixed visual-language context. Subsequent denoising steps update only the action tokens: action queries attend to the cached video/text keys and values together with their own action keys and values. This removes the repeated video-stream computation from the latter part of sampling while preserving the same attention semantics as the V2A-style model. Combined with the other inference optimizations above, this design enables ours to reach a 11 Hz inference frequency, exceeding typical human reaction speed.

Action Smoothing. Following DreamZero [36], MotuBrain applies action chunk smoothing to improve execution smoothness. Each chunk is upsampled to twice its original temporal resolution, smoothed with a Savitzky-Golay filter, and then downsampled back to the original resolution.

Frequency-aware interpolation. After smoothing, the action sequence is interpolated according to the ratio between the model action frequency and the low-level control frequency. This frequency-aware interpolation preserves the temporal scale of the model prediction: rapid transitions remain fast during execution, while slower motion segments are expanded over a longer control duration. For manipulation tasks with different motion phases, this helps preserve the velocity profile predicted by the model.

Table 2 Cumulative speedup from inference optimizations. Each row applies the listed technique on top of all preceding ones. Latency is measured end-to-end on the non-autoregressive model. Per-step latency is reported where applicable.

Technique	Steps	Per-step (ms)	Latency (s)	Frequency (Hz)	Speedup
Baseline	50	95.0	4.90	0.20	1.00×
+ Noise sampling	30	95.0	2.90	0.34	1.69×
+ <code>torch.compile</code>	30	32.7	0.98	1.02	5.00×
+ FP8 quantization	30	29.3	0.88	1.14	5.57×
+ DiT cache	30	–	0.20	5.00	24.5×
+ V2A-style	30(action-only)	–	0.09	11.11	54.4×

2.4.2 Real-time Inference and Execution

Real-time control is essential for robotic deployment, while the inference latency of world action models is usually non-negligible. MotuBrain therefore decouples the model inference loop from the robot action

execution loop: the controller executes the current action chunk at the target control frequency, while the world action model asynchronously generates the next chunk from the latest observation.

However, directly switching to the newly generated chunk may introduce chunk-boundary discontinuities, such as action regression, velocity jumps, and high-frequency jitter, since adjacent chunks may be generated from different observations and action modes. To reduce such boundary mismatch, we adopt an RTC-inspired strategy [7]. The unexecuted part of the current chunk is used as a constraint for the next generation and fused before denoising, where the latency-affected prefix is treated as a frozen region and the remaining overlapping actions are used as soft constraints.

Let the current action chunk be

$$A^{old} = \{a_0^{old}, a_1^{old}, \dots, a_{H-1}^{old}\}, \quad (10)$$

where H denotes the prediction horizon. If a new inference request is launched after s actions have been executed, the remaining actions in the current chunk are

$$A^{remain} = \{a_s^{old}, a_{s+1}^{old}, \dots, a_{H-1}^{old}\}. \quad (11)$$

Given the end-to-end inference latency δ and the controller period Δt , the inference delay measured in action steps is defined as

$$d = \left\lceil \frac{\delta}{\Delta t} \right\rceil. \quad (12)$$

The first d steps of the next chunk correspond to the period in which the previous chunk is still being executed while the new inference is running. These actions are therefore fully constrained by the remaining actions from the previous chunk. For the following overlapping region, we apply a smooth decay weight to gradually reduce the influence of the previous chunk and allow the new prediction to take over.

Let ρ_i denote the normalized progress within the fusion window, where i is the action index in the new chunk, and L denotes the end of the fusion window. The exponential decay function used in our implementation is

$$g(\rho_i) = \frac{\rho_i (e^{\rho_i} - 1)}{e - 1}. \quad (13)$$

Since $g(\rho_i)$ increases from 0 to 1, the fusion weight is defined as

$$w_i = \begin{cases} 1, & 0 \leq i < d, \\ 1 - g(\rho_i), & d \leq i < L, \\ 0, & i \geq L. \end{cases} \quad (14)$$

The fused action is computed as

$$\tilde{a}_i^{new} = w_i a_i^{remain} + (1 - w_i) a_i^{new}. \quad (15)$$

To improve robustness under variable inference and communication latency, MotuBrain maintains a delay queue Q that stores recent inference delays, and the system uses $\hat{d}_{t+1} = \max(Q)$ as a conservative estimate for the next inference request. This estimated delay determines the length of the frozen prefix. The fusion window is also adjusted accordingly: when the estimated delay increases, more steps are treated as fully constrained; when the estimated delay decreases, more future actions are left to be updated by the new prediction. This makes the asynchronous execution more stable under fluctuating network and model latency.

3 Evaluations

In this section, we evaluate MotuBrain in both simulation and real world environment.

3.1 Evaluation on Simulation Environment

Following the protocol of RoboTwin 2.0 [11], we adopt a multi-task training setup where all models are trained with 2,500 demonstrations collected from clean scenes (50 per task) together with 25,000 demonstrations from heavily randomized scenes (500 per task). For training efficiency and temporal consistency across tasks, we downsample videos to 5 Hz and action sequences to 10 Hz. We further conduct ablation studies over multiple policy architectures to analyze how different design choices affect the performance of MotuBrain under the same data and evaluation setting.

Table 3 Robotwin2.0 Results. Following previous works, MotuBrain is fine-tuned from pre-trained weights on the official RoboTwin2.0 dataset (clean + randomized), yielding the evaluation results presented in the table. MotuBrain-Non-AR represents non-autoregressive mode.

Model	Clean	Randomized
# VLA Based		
π_0	65.9	58.4
X-VLA	72.9	72.8
$\pi_{0.5}$	82.7	76.8
starVLA	88.2	88.3
ABot-M0	81.2	80.4
LingBot-VLA	86.5	85.3
# World Model Based		
JEPA-VLA	73.5	-
Motus	88.7	87.0
LingBot-VA	92.9	91.6
Fast-WAM	91.9	91.8
Being-H0.7	90.2	89.6
MotuBrain-Non-AR w/o Pretrain, HBridge	89.1	88.8
MotuBrain-Non-AR w/o Pretrain	89.6	89.5
MotuBrain w/o Pretrain	91.5	91.3
MotuBrain-Non-AR	91.9	92.3
MotuBrain	95.8	96.1

As shown in Table 3, MotuBrain achieves the best average success rate on RoboTwin, reaching 95.8% in clean scenes and 96.1% in randomized scenes. It ranks first in both settings and is the only model on the leaderboard whose average score exceeds 95% under randomized evaluation. At the per-task level, as summarized in Table 4, MotuBrain attains a perfect score on 24 tasks under the clean setting and 25 tasks under the randomized setting, with 19 tasks achieving 100% success in both settings. Moreover, it surpasses 90% success on 42 clean tasks and 44 randomized tasks, demonstrating consistently strong generalization across the full 50-task benchmark and strong ability to act robustly in the world.

From the detailed breakdown in Table 4, the largest gains are concentrated on tasks that require robust multi-stage manipulation and precise interaction under visual variation. In particular, we observe clear improvements on handover and coordination-heavy tasks such as *Handover Block*, articulated-object interaction tasks such as *Open Microwave*, *Press Stapler*, and *Turn Switch*, as well as fine-grained spatial arrangement tasks including *Blocks Ranking Size*, *Move Can Pot*, *Place A2B Left*, and *Place Can Basket*. These results suggest that MotuBrain is especially effective on tasks demanding accurate temporal modeling, stable contact-rich control, and robustness to scene randomization.

We further observe a favorable multi-task scaling trend: as the number of training tasks increases, the average success rate of MotuBrain continues to improve, suggesting that more tasks expose the model to more reusable world knowledge about objects, contacts, and temporal transitions, as illustrated in Figure 3. In contrast, conventional VLA baselines exhibit noticeably weaker scaling in the same regime, indicating stronger task interference and less effective knowledge sharing across tasks. We also find that MotuBrain is substantially

Table 4 Per-task success rates on RoboTwin under clean and randomized evaluation settings.

Task	$\pi_{0.5}$		X-VLA		Motus		LingBot-VA		Fast-WAM		MotuBrain	
	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.
Adjust Bottle	100	99	100	99	89	93	90	94	100	100	99	100
Beat Block Hammer	96	93	92	88	95	88	96	98	99	97	100	100
Blocks Ranking RGB	92	85	83	83	99	97	99	98	100	100	100	100
Blocks Ranking Size	49	26	67	74	75	63	94	96	94	98	100	100
Click Alarmclock	98	89	99	99	100	100	99	100	100	100	100	100
Click Bell	99	66	100	100	100	100	100	100	100	100	100	100
Dump Bin Bigbin	92	97	79	77	95	91	89	96	97	96	99	100
Grab Roller	100	100	100	100	100	100	100	100	100	100	100	100
Handover Block	66	57	73	37	86	73	99	78	95	81	100	95
Handover Mic	98	97	0	0	78	63	94	96	99	100	100	100
Hanging Mug	18	17	23	27	38	38	40	28	58	62	55	43
Lift Pot	96	85	99	100	96	99	100	99	100	100	100	100
Move Can Pot	51	55	89	86	34	74	94	97	90	88	99	100
Move Pillbottle Pad	84	61	73	71	93	96	99	99	100	99	100	100
Move Playingcard Away	96	84	93	98	100	96	100	99	100	100	100	100
Move Stapler Pad	56	42	78	73	83	85	91	79	77	64	85	93
Open Laptop	90	96	93	100	95	91	92	94	98	100	100	99
Open Microwave	34	77	79	71	95	91	82	86	62	45	100	100
Pick Diverse Bottles	81	71	58	36	90	91	89	82	80	85	95	89
Pick Dual Bottles	93	63	47	36	96	90	100	99	100	96	100	100
Place A2B Left	87	82	48	49	88	79	97	93	95	93	100	99
Place A2B Right	87	84	36	36	91	87	97	95	93	99	95	99
Place Bread Basket	77	64	81	71	91	94	97	95	91	93	97	99
Place Bread Skillet	85	66	77	67	86	83	95	90	90	93	95	93
Place Burger Fries	94	87	94	94	98	98	97	95	96	99	99	100
Place Can Basket	62	62	49	52	81	76	81	84	71	69	85	93
Place Cans Plasticbox	94	84	97	98	98	94	100	99	99	96	100	100
Place Container Plate	99	95	97	95	98	99	99	97	96	100	97	100
Place Dual Shoes	75	75	79	88	93	87	94	89	94	88	93	95
Place Empty Cup	100	99	100	98	99	98	100	100	100	100	100	100
Place Fan	87	85	80	75	91	87	99	93	96	96	97	97
Place Mouse Pad	60	39	70	70	66	68	93	96	83	89	97	97
Place Object Basket	80	76	44	39	81	87	91	88	89	88	89	93
Place Object Scale	86	80	52	74	88	85	96	95	90	97	97	100
Place Object Stand	91	85	86	88	98	97	99	96	90	94	97	97
Place Phone Stand	81	81	88	87	87	86	97	97	97	99	97	93
Place Shoe	92	93	96	95	99	97	98	98	96	99	100	97
Press Stapler	87	83	92	98	93	98	85	82	90	97	100	100
Put Bottles Dustbin	84	79	74	77	81	79	87	91	95	90	87	87
Put Object Cabinet	80	79	46	48	88	71	85	87	94	89	92	93
Rotate QRcode	89	87	34	33	89	73	96	91	93	89	90	93
Scan Object	72	65	14	36	67	66	96	91	89	92	94	90
Shake Bottle	99	97	100	100	100	97	100	97	100	100	100	100
Shake Bottle Horizontally	99	99	99	100	100	98	100	99	100	100	100	100
Stack Blocks Three	91	76	6	10	91	95	99	98	95	97	100	100
Stack Blocks Two	97	100	92	87	100	98	100	98	100	100	100	100
Stack Bowls Three	77	71	76	86	79	87	86	83	80	81	86	90
Stack Bowls Two	95	96	96	93	98	98	94	98	92	98	92	98
Stamp Seal	79	55	76	82	93	92	96	97	90	94	100	98
Turn Switch	62	54	40	61	84	78	44	45	61	59	82	84
Average	82.74	76.76	72.80	72.84	88.66	87.02	92.90	91.50	91.88	91.78	95.80	96.08

more data-efficient than both conventional VLA baselines and the previous Motus model: under the same data scaling protocol, MotuBrain reaches stronger performance with fewer training trajectories and continues to improve more effectively as data grows, as shown in Figure 4.

Taken together, Figures 3 and 4 further suggest that, in our setting, increasing task diversity is more effective than merely scaling up the amount of data collected for a fixed task set, as evidenced by the steeper improvement trend in the task-scaling curve. This observation is consistent with the hypothesis that broader task coverage exposes the model to a richer set of interaction patterns, object affordances, and temporal transitions, thereby improving knowledge reuse and cross-task generalization more efficiently than data duplication alone.




Figure 3 Task scaling results. For each point on the curve, we train the model using data from only the specified number of tasks and keep optimization running until convergence, i.e., until the average success rate becomes approximately stable. As the number of training tasks increases, MotuBrain shows a clear upward trend in average success rate and consistently stronger scaling behavior than conventional VLA baselines.




Figure 4 Data scaling results. For each data budget, we uniformly subsample demonstrations from every task to keep the per-task data distribution balanced. The number of training steps is set proportional to the total amount of training data; in particular, training on the full 27,500-trajectory dataset uses 50,000 optimization steps. MotuBrain continues to benefit from additional training data, indicating that larger-scale supervision improves policy performance and robustness.

3.2 Evaluation on World Models

Beyond action execution, we also evaluate whether MotuBrain can understand and predict changes in the world itself. We adopt WorldArena [31], a unified benchmark that scores embodied world models through 16 numerical metrics across six perceptual sub-dimensions—visual quality, motion quality, content consistency, physics adherence, 3D accuracy, and controllability. We use the entire 2,500-trajectory release for high-resolution training and report benchmark submission results. Because MotuBrain is a unified world-action model, the same backbone simultaneously serves as the manipulation policy used in Section 3.1 and as the world predictor evaluated here. For WorldArena, we run the model in its forward-dynamics mode (FDM), which serves as an action-conditioned world model. We train and run inference at high resolution with a 5 Hz video stream paired with 10 Hz actions (1:2 ratio), using classifier-free guidance to improve instruction-following capability.

As shown in Figure 5, MotuBrain attains the highest EWMScore (63.77) among embodied world models and video baselines on the WorldArena leaderboard. Its margin over the strongest video-generation baseline shown here (Wan2.6, 59.80) is approximately four points. Notably, MotuBrain ranks first on motion quality and remains competitive on several other dimensions, indicating that strong dynamics modeling can be achieved without sacrificing overall perceptual quality.

The detailed per-metric breakdown in Table 5 shows that MotuBrain’s lead is driven primarily by its strong

performance in the **Motion Quality** dimension. The Flow Score, which measures the overall magnitude of motion in the predicted video, is the highest among all baselines, indicating that MotuBrain generates sequences with sustained inter-frame displacement rather than near-static clips that collapse to a stable but trivial solution. The leading Motion Smoothness indicates that the produced motion unfolds with realistic inertia rather than frame-to-frame jitter. The high Dynamic Degree, which measures whether motion concentrates in the regions that actually matter, further confirms that MotuBrain’s predicted activity is localized on the embodiment and the manipulated objects rather than diffused into background flicker. Together, these three metrics show that MotuBrain predicts motion that is substantial, smooth, and concentrated in embodiment-relevant regions.

This distinction matters because the original WorldArena study reports that EWMScore correlates only weakly with downstream action-planning success ($r=0.36$), reflecting the well-known perception–functionality gap: visually impressive world models often fail when used for control, whereas functionally useful ones often appear visually unpolished. MotuBrain helps bridge this gap. On the perceptual side, it ranks first on the EWMScore leaderboard among embodied world models and video baselines; on the functional side, it achieves an average success rate of 95.8% across the 50 RoboTwin manipulation tasks (Section 3.1), the highest among all evaluated VLA and world-model baselines. This cross-benchmark consistency suggests that the dynamics representation MotuBrain learns is simultaneously perceptually faithful and functionally actionable when consumed by a policy.

Table 5 Per-metric WorldArena results and overall EWMScore comparison. We compare MotuBrain with several publicly listed and representative entries on the WorldArena leaderboard. The best score in each row is shown in bold. EWMScore is the arithmetic mean of all 16 normalized metrics, scaled to [0, 100].

		MotuBrain	Veo3.1	Wan2.6	Ctrl-World	ABot-PW	GigaWorld-1
Visual Quality	Image Quality	0.4459	0.6557	0.6736	0.4244	0.6103	0.5118
	Aesthetic Quality	0.3977	0.4879	0.4440	0.3705	0.4183	0.4117
	JEPA Similarity	0.9667	0.5797	0.7257	0.9277	0.9036	0.9677
Motion Quality	Dynamic Degree	0.5148	0.1466	0.3363	0.4182	0.3933	0.3052
	Flow Score	0.4911	0.0826	0.2201	0.3357	0.2922	0.1864
	Motion Smoothness	0.8566	0.6785	0.8212	0.7734	0.7648	0.6833
Content Consistency	Subject Consistency	0.8240	0.7582	0.7315	0.8356	0.8056	0.8097
	Background Consistency	0.9021	0.9167	0.8429	0.9030	0.8902	0.8643
	Photometric Consistency	0.0574	0.3752	0.3457	0.1288	0.2052	0.2811
Physics Adherence	Interaction Quality	0.7174	0.8150	0.7316	0.6262	0.8196	0.7510
	Trajectory Accuracy	0.4793	0.1136	0.1218	0.4820	0.3150	0.5427
3D Accuracy	Depth Accuracy	0.8992	0.7427	0.7543	0.9325	0.7199	0.9844
	Perspectivity	0.9290	0.9964	0.9394	0.8366	0.9894	0.9560
Controlability	Instruction Following	0.8072	0.9714	0.8996	0.6768	0.9210	0.8214
	Semantic Alignment	0.8941	0.8379	0.8809	0.8868	0.8958	0.8942
	Action Following	0.0203	0.0852	0.0992	0.0390	0.0765	0.0028
EWMScore (↑)		63.77	57.77	59.80	59.98	62.63	62.34

Model	EWMScore	Rank	Visual Quality	Motion Quality	Content Consistency	Physics Adherence	3D Accuracy	Controllability	Action Following	Aesthetic Quality	Background Consist
MotuBrain	63.77 🏆	1	60.34	62.08 🏆	59.45	59.83	91.41	57.39	2.03	39.77	90.21
FlowWAM-FiveAges	63.71	2	62.28	40.84	67.40	69.04 🏆	96.76	58.73	3.50	40.83	89.97
ABot-PhysWorld(text)	62.63	3	64.41 🏆	48.34	63.37	56.73	85.46	63.11	7.65	41.83	89.02
GigaWorld-1	62.34	4	63.04	39.16	65.17	64.68	97.02 🏆	57.28	0.28	41.17	86.43
Ctrl-World	59.98	5	57.42	50.91	62.25	55.41	88.46	53.42	3.90	37.05	90.30
Wan2.6	59.80	6	61.44	45.92	64.00	42.67	84.68	62.66	9.92	44.40	84.29
CogvideoX	58.79	7	55.79	42.18	67.71	50.88	88.28	55.09	0.00	36.30	88.38
Veo3.1	57.77	8	57.44	30.26	68.34 🏆	46.43	86.96	63.15 🏆	8.52	48.79 🏆	91.67
IRASim	56.14	9	56.14	32.52	65.74	49.34	88.83	52.92	3.94	32.86	78.69
Tesseract	54.62	10	38.90	54.86	65.90	39.54	79.06	52.61	2.36	41.30	92.82
WoW	54.42	11	53.16	43.67	64.39	40.62	78.28	49.75	4.51	38.55	90.53
Cosmos-Predict 2.5(action)	54.29	12	58.09	31.52	61.72	43.98	87.52	50.57	1.31	34.14	72.72
Cosmos-Predict 2.5(text)	53.06	13	52.17	36.70	58.52	38.37	84.46	53.68	10.01 🏆	44.33	73.98
Vidar	51.92	14	47.57	39.06	56.39	41.05	80.82	52.62	7.86	39.87	78.24
Wan2.2	51.71	15	49.96	33.15	66.62	35.74	81.28	48.07	5.39	38.32	79.37
GigaWorld-0	50.96	16	42.54	49.23	58.20	36.64	70.92	50.11	7.55	40.15	88.15
RoboMaster	50.35	17	34.58	42.91	67.33	33.57	80.85	47.45	3.43	37.87	90.80
Genie Envisioner	41.37	18	29.46	40.50	59.37	12.68	70.58	35.80	1.07	26.39	87.54

Figure 5 WorldArena public leaderboard.

3.3 Real-World Experiments

We further investigate rapid deployment to new embodiments in real-world settings. Starting from a pretrained model, MotuBrain can be adapted to a new embodiment using only 50–100 same-embodiment trajectories. We validate this capability across multiple humanoid platforms. Notably, these results are achieved without relying on auxiliary components such as VLM-based planners, dual-system decompositions, external memory modules, or additional reinforcement/retry data. In other words, the native world action model alone provides a sufficiently strong control prior to enable practical transfer.

3.3.1 Real-World Performance

Quantitative Evaluation. We conduct a quantitative evaluation of MotuBrain on a suite of representative real-world household tasks. Performance is measured using a normalized scoring metric with a maximum value of 100, where each sub-task step contributes equally to the overall score. A step receives full credit if it is successfully completed on the first attempt; if it succeeds after one retry, two retries, or three or more retries, it receives 80%, 50%, or 0 credit, respectively.

For each task, the model is trained on 100 task-specific trajectories consisting of single, coherent, and successfully executed compound actions, without atomic-level annotations. We further assess performance across tasks with varying execution horizons and differing numbers of atomic actions. The quantitative results are summarized in Table 6, Table 7, and Table 8.

In the Making Oden task, which requires simultaneous bimanual manipulation with one arm pouring a drink and the other scooping dumplings, MotuBrain is evaluated over five consecutive trials, with an average

execution time of 33 seconds across seven atomic actions. The robot executes the subtasks concurrently with its left and right hands, achieving an overall score of 98.54, with only a single atomic action failure observed.

Table 6 Per-step performance evaluation of the bimanual Making Oden task. The overall score reaches 98.54, with only a slight deficiency in the pouring step.

Task Name	IDs	Sub-task Names	Total Scores	Our Scores
Making Oden	1	Grab the drink with your right hand	14.29	14.29
	2	Pours drink successfully with your right hand	14.29	12.83
	3	Put the beverage bottle down with your right hand	14.29	14.29
	4	Grab the spoon with your left hand	14.29	14.29
	5	Pick up the balls with your left hand	14.29	14.29
	6	Place the balls in the bowl with your left hand	14.29	14.29
	7	Put the spoon back with your left hand	14.26	14.26

In the Mixing Cocktails task, which involves an extended long-horizon sequence with complex multi-step dependencies across 15 atomic actions, the model is evaluated over an average execution time of 124 seconds. Across seven consecutive trials, MotuBrain attains an overall score of 97.34, demonstrating stable performance under substantially increased temporal complexity.

Table 7 Per-step real-world evaluation results for the long-horizon Mixing Cocktailtask. Each sub-task contributes equally to the final score, and MotuBrain achieves strong overall performance with only minor deficiencies in a small number of steps.

Task Name	IDs	Sub-task Names	Total Scores	Our Scores
Mixing Cocktails	1	Pick up the bottle	6.67	6.67
	2	Fill the mixing glass	6.67	5.72
	3	Return the bottle to its original position	6.67	6.67
	4	Grasp the milk bottle	6.67	6.67
	5	Pour milk into the glass	6.67	6.29
	6	Return the milk bottle to its original position	6.67	6.67
	7	Grasp the mixing glass	6.67	6.67
	8	Pour the mixed liquid into the serving glass	6.67	6.67
	9	Place the mixing glass back	6.67	6.67
	10	Grasp a peppermint leaf	6.67	6.67
	11	Place the peppermint leaf into the glass	6.67	6.67
	12	Grasp the prepared drink	6.67	6.67
	13	Place the prepared drink on the serving tray	6.67	5.53
	14	Squeeze the syrup bottle to complete the mixing	6.67	6.48
	15	Return to the home position	6.62	6.62

In the Flower Arrangement task, which requires precise and fine-grained manipulation for accurate flower placement, MotuBrain is evaluated over 10 consecutive trials, achieving an average execution time of 138 seconds across 10 atomic actions and an overall score of 83.30. Notably, for repetitive failure cases, MotuBrain demonstrates an inherent retry capability despite the absence of explicit recovery supervision during training, indicating that it captures the underlying long-horizon task objective rather than relying on fixed motion replay.

Quantitative Evaluation. We evaluate MotuBrain in real-world home environments on several everyday tasks, including flower arrangement, meal preparation, cocktail mixing, sofa tidying, and sink organization. These tasks encompass representative activities across multiple household environments. Across all settings, MotuBrain consistently accomplishes them on diverse humanoid robot embodiments, demonstrating robust generalization and the ability to reliably compose and execute fundamental multi-step household skills.

Few-shot Evaluation and Self-correction. Following the training-stage setup, we utilize only 100 trajectories of continuous atomic actions, without any additional sub-task annotations. In the flower-arrangement task, where the placement height varies across trials, MotuBrain is able to execute repeated long-horizon task sequences in a novel home environment, with each rollout lasting approximately 2-3 minutes. In these

Table 8 Per-step real-world evaluation results for the *Flower Arrangement* task. Our model demonstrates stable execution across repeated long-horizon trials.

Task Name	IDs	Sub-task Names	Total Scores	Our Scores
Flower Arrangement	1	Left hand picks up the first flower branch	10	8.8
	2	Insert flowers into vase	10	8.8
	3	Pick up the second branch with your left hand	10	7.3
	4	Insert flowers into vase	10	8.0
	5	Pick up the third flower branch with your left hand	10	5.6
	6	Insert flowers into vase	10	5.8
	7	Pick up the kettle with your right hand and aim it at the vase	10	10
	8	Hold water spray with left hand	10	9.5
	9	Release the kettle with your left hand and lower the kettle with your right hand	10	10
	10	Left hand delivering the vase to the table	10	9.5

experiments, the model receives the goal instruction only once at the beginning of each episode. During execution, it predicts future action states and continuously refines its policy through closed-loop integration of visual observations. When execution errors occur, such as failed insertion of a flower into the vase, the model leverages updated perceptual feedback to adapt its behavior and perform online correction. This long-horizon, few-shot evaluation demonstrates that MotuBrain possesses strong physical world modeling capabilities and achieves a significantly higher success rate than Vision-Language-Action models.

Bimanual Action Generalization. Our demonstrations include asymmetric bimanual behaviors, where the two arms execute distinct subtasks concurrently, such as pouring and scooping. Our model is robust to asynchronous subtask completion: when one subtask is externally fulfilled ahead of time, the other arm continues execution without interruption. For example, if a person helps fill the drink before the robot finishes pouring, the other arm can still continue the scooping behavior and complete the remaining sub-task. Such behavior suggests that MotuBrain does not rely on rigid inter-arm dependencies, but instead learns coordinated yet decoupled control under a shared temporal context.

Scene Generalization. We evaluate cross-scene generalization under a limited-data setting. In the flower-arrangement task, In the flower-arrangement task, our model trained on a single scene is deployed in a novel environment with previously unseen flowers and vases. Despite being trained on only one type of flower and vase, the model generalizes to four unseen flower-vase combinations and achieves a success rate above 80%. In contrast, the VLA-based baseline typically requires training on at least three object categories with diverse shapes and sizes before achieving reasonable generalization to a fourth unseen instance. These results indicate that MotuBrain captures object geometry of novel household objects and leverages it for downstream action reasoning.

3.3.2 Real-World Demonstrations

We visualize the real-world deployment of MotuBrain across a range of manipulation tasks. As shown in Figure 6, the model executes these tasks reliably and effectively in diverse household scenarios.



Place the toothbrush in the cup and put the soap back in its place.



Mix a cocktail using the milk and beverage on the table, place it on the tray, and serve it to the customer.



Pour a glass of juice while scooping a serving of dumplings from the pot into the bowl.



Put the clothes on the sofa into the laundry basket and put the pillows back in place.



Insert the flowers into the vase, then spray water with the watering can.

Figure 6 Real-world demonstrations on multiple long-horizon and dexterous manipulation tasks.

4 Conclusion and Future Work

We present MotuBrain, a unified world action model that learns robot control by jointly modeling future visual dynamics and action generation. By combining large-scale pretraining with lightweight robot adaptation, MotuBrain inherits broad semantic and physical priors while remaining practical for downstream deployment. Empirically, the model achieves strong and consistent performance across three complementary settings: large-scale simulation benchmarks, open-ended world-model evaluation, and real-world humanoid transfer. Together, these results suggest that a single pretrained world-action model can serve as a scalable foundation for both understanding how the world evolves and deciding how a robot should act within it.

Our study also highlights a broader design principle for embodied intelligence: action learning benefits substantially from being trained together with predictive world modeling rather than being optimized as an isolated imitation problem. The resulting representation is not only more robust to compounding execution errors, but also more transferable across tasks, embodiments, and environments.

There remain several important directions for future work. First, current adaptation still relies on a modest amount of same-embodiment robot data; reducing this requirement further would improve accessibility and deployment speed. Second, extending the framework to longer-horizon mobile manipulation, richer tactile interaction, and more dynamic human-centered environments would better test the limits of the learned world prior. Finally, integrating stronger uncertainty estimation, explicit memory, and online test-time adaptation may further improve robustness in open-world settings where disturbances and task variations cannot be fully covered during training.

5 Contributors

Data: Chendong Xiang*, Louis Liu*, Jiabao Liu*, James Li*, Zeyuan Wang, Hongzhe Bi, Hengkai Tan

Base Model: Zeyuan Wang*, Chendong Xiang*, Hengkai Tan*, Haitian Liu, Yao Feng, Ruowen Zhao, Shuhe Huang, Hongzhe Bi

Post-Training: Zeyuan Wang*, Chendong Xiang*, Haitian Liu*, Rongxu Cui*, Ruowen Zhao, Hengkai Tan, Jingrui Pang, Yao Feng, Shuhe Huang, Mengchen Cai, Yinze Rong

Evaluation: Rongxu Cui*, Zeyuan Wang*, Haitian Liu*, Chendong Xiang*, Hengkai Tan*, Ruowen Zhao*, Shuhe Huang*, Mengchen Cai, Runqing Wang, Kiro Jing, James Li, Yao Feng

Project Lead: Hengkai Tan

Advisor: Fan Bao, Jun Zhu

** denotes the core-contributors or leaders of each sub-module.*

References

- [1] Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. [arXiv preprint arXiv:2511.21631](#), 2025.
- [2] Fan Bao, Shen Nie, Kaiwen Xue, Chongxuan Li, Shi Pu, Yaole Wang, Gang Yue, Yue Cao, Hang Su, and Jun Zhu. One transformer fits all distributions in multi-modal diffusion at scale, 2023.
- [3] Fan Bao, Chendong Xiang, Gang Yue, Guande He, Hongzhou Zhu, Kaiwen Zheng, Min Zhao, Shilong Liu, Yaole Wang, and Jun Zhu. Vidu: a highly consistent, dynamic and skilled text-to-video generator with diffusion models, 2024.
- [4] Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. [arXiv preprint arXiv:2407.07726](#), 2024.
- [5] 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.
- [6] 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 Mothukuri, 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, 2026.
- [7] Kevin Black, Manuel Y Galliker, and Sergey Levine. Real-time execution of action chunking flow policies. [arXiv preprint arXiv:2506.07339](#), 2025.
- [8] 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.
- [9] Jake Bruce, Michael Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Bechtile, Feryal Behbahani, Stephanie Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando de Freitas, Satinder Singh, and Tim Rocktäschel. Genie: Generative interactive environments, 2024.
- [10] Qingwen Bu, Yanting Yang, Jisong Cai, Shenyuan Gao, Guanghui Ren, Maoqing Yao, Ping Luo, and Hongyang Li. Univla: Learning to act anywhere with task-centric latent actions, 2025.
- [11] Tianxing Chen, Zanzin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. [arXiv preprint arXiv:2506.18088](#), 2025.
- [12] Open X-Embodiment Collaboration. Open x-embodiment: Robotic learning datasets and rt-x models, 2023.
- [13] 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, 2023.
- [14] 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. [arXiv preprint arXiv:2507.12898](#), 2025.
- [15] Shenyuan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, and Chuang Gan. Adaworld: Learning adaptable world models with latent actions, 2025.
- [16] Yanjiang Guo, Lucy Xiaoyang Shi, Jianyu Chen, and Chelsea Finn. Ctrl-world: A controllable generative world model for robot manipulation. [arXiv preprint arXiv:2510.10125](#), 2025.
- [17] Haoran He, Yang Zhang, Liang Lin, Zhongwen Xu, and Ling Pan. Pre-trained video generative models as world simulators, 2025.

- [18] Xianglong He, Chunli Peng, Zexiang Liu, Boyang Wang, Yifan Zhang, Qi Cui, Fei Kang, Biao Jiang, Mengyin An, Yangyang Ren, Baixin Xu, Hao-Xiang Guo, Kaixiong Gong, Size Wu, Wei Li, Xuchen Song, Yang Liu, Yangguang Li, and Yahui Zhou. Matrix-game 2.0: An open-source real-time and streaming interactive world model, 2026.
- [19] 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.
- [20] 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.
- [21] Moo Jin Kim, Yihuai Gao, Tsung-Yi Lin, Yen-Chen Lin, Yunhao Ge, Grace Lam, Percy Liang, Shuran Song, Ming-Yu Liu, Chelsea Finn, et al. Cosmos policy: Fine-tuning video models for visuomotor control and planning. [arXiv preprint arXiv:2601.16163](#), 2026.
- [22] 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.
- [23] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, DuoJun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xincheng Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models, 2025.
- [24] Lin Li, Qihang Zhang, Yiming Luo, Shuai Yang, Ruilin Wang, Fei Han, Mingrui Yu, Zelin Gao, Nan Xue, Xing Zhu, et al. Causal world modeling for robot control. [arXiv preprint arXiv:2601.21998](#), 2026.
- [25] 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. [arXiv preprint arXiv:2508.05635](#), 2025.
- [26] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation, 2025.
- [27] 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.
- [28] NVIDIA, :, Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, Daniel Dworakowski, Jiaojiao Fan, Michele Fenzi, Francesco Ferroni, Sanja Fidler, Dieter Fox, Songwei Ge, Yunhao Ge, Jinwei Gu, Siddharth Gururani, Ethan He, Jiahui Huang, Jacob Huffman, Pooya Jannaty, Jingyi Jin, Seung Wook Kim, Gergely Klár, Grace Lam, Shiyi Lan, Laura Leal-Taixe, Anqi Li, Zhaoshuo Li, Chen-Hsuan Lin, Tsung-Yi Lin, Huan Ling, Ming-Yu Liu, Xian Liu, Alice Luo, Qianli Ma, Hanzhi Mao, Kaichun Mo, Arsalan Mousavian, Seungjun Nah, Sriharsha Niverty, David Page, Despoina Paschalidou, Zeeshan Patel, Lindsey Pavao, Morteza Ramezanali, Fitsum Reda, Xiaowei Ren, Vasanth Rao Naik Sabavat, Ed Schmerling, Stella Shi, Bartosz Stefaniak, Shitao Tang, Lyne Tchampi, Przemek Tredak, Wei-Cheng Tseng, Jibin Varghese, Hao Wang, Haoxiang Wang, Heng Wang, Ting-Chun Wang, Fangyin Wei, Xinyue Wei, Jay Zhangjie Wu, Jiashu Xu, Wei Yang, Lin Yen-Chen, Xiaohui Zeng, Yu Zeng, Jing Zhang, Qingsheng Zhang, Yuxuan Zhang, Qingqing Zhao, and Artur Zolkowski. Cosmos world foundation model platform for physical ai, 2025.
- [29] Jonas Pai, Liam Achenbach, Victoriano Montesinos, Benedek Forrai, Oier Mees, and Elvis Nava. mimic-video: Video-action models for generalizable robot control beyond vlas, 2025.
- [30] Team Seedance, De Chen, Liyang Chen, Xin Chen, Ying Chen, Zhuo Chen, Zhuowei Chen, Feng Cheng, Tianheng Cheng, Yufeng Cheng, Mojie Chi, Xuyan Chi, Jian Cong, Qinpeng Cui, Fei Ding, Qide Dong, Yujiao Du, Haojie Duanmu, Junliang Fan, Jiarui Fang, Jing Fang, Zetao Fang, Chengjian Feng, Yu Gao, Diandian Gu, Dong Guo,

- Hanzhong Guo, Qiushan Guo, Boyang Hao, Hongxiang Hao, Haoxun He, Jiaao He, Qian He, Tuyen Hoang, Heng Hu, Ruoqing Hu, Yuxiang Hu, Jiancheng Huang, Weilin Huang, Zhaoyang Huang, Zhongyi Huang, Jishuo Jin, Ming Jing, Ashley Kim, Shanshan Lao, Yichong Leng, Bingchuan Li, Gen Li, Haifeng Li, Huixia Li, Jiashi Li, Ming Li, Xiaojie Li, Xingxing Li, Yameng Li, Yiyang Li, Yu Li, Yueyan Li, Chao Liang, Han Liang, Jianzhong Liang, Ying Liang, Wang Liao, J. H. Lien, Shanchuan Lin, Xi Lin, Feng Ling, Yue Ling, Fangfang Liu, Jiawei Liu, Jihao Liu, Jingtuo Liu, Shu Liu, Sichao Liu, Wei Liu, Xue Liu, Zuxi Liu, Ruijie Lu, Lecheng Lyu, Jingting Ma, Tianxiang Ma, Xiaonan Nie, Jingzhe Ning, Junjie Pan, Xitong Pan, Ronggui Peng, Xueqiong Qu, Yuxi Ren, Yuchen Shen, Guang Shi, Lei Shi, Yinglong Song, Fan Sun, Li Sun, Renfei Sun, Wenjing Tang, Boyang Tao, Zirui Tao, Dongliang Wang, Feng Wang, Hulin Wang, Ke Wang, Qingyi Wang, Rui Wang, Shuai Wang, Shulei Wang, Weichen Wang, Xuanda Wang, Yanhui Wang, Yue Wang, Yuping Wang, Yuxuan Wang, Zijie Wang, Ziyu Wang, Guoqiang Wei, Meng Wei, Di Wu, Guohong Wu, Hanjie Wu, Huachao Wu, Jian Wu, Jie Wu, Ruolan Wu, Shaojin Wu, Xiaohu Wu, Xinglong Wu, Yonghui Wu, Ruiqi Xia, Xin Xia, Xuefeng Xiao, Shuang Xu, Bangbang Yang, Jiaqi Yang, Runkai Yang, Tao Yang, Yihang Yang, Zhixian Yang, Ziyang Yang, Fulong Ye, Bingqian Yi, Xing Yin, Yongbin You, Linxiao Yuan, Weihong Zeng, Xuejiao Zeng, Yan Zeng, Siyu Zhai, Zhonghua Zhai, Bowen Zhang, Chenlin Zhang, Heng Zhang, Jun Zhang, Manlin Zhang, Peiyuan Zhang, Shuo Zhang, Xiaohe Zhang, Xiaoying Zhang, Xinyan Zhang, Xinyi Zhang, Yichi Zhang, Zixiang Zhang, Haiyu Zhao, Huating Zhao, Liming Zhao, Yian Zhao, Guangcong Zheng, Jianbin Zheng, Xiaozheng Zheng, Zerong Zheng, Kuan Zhu, and Feilong Zuo. Seedance 2.0: Advancing video generation for world complexity, 2026.
- [31] Yu Shang, Zhuohang Li, Yiding Ma, Weikang Su, Xin Jin, Ziyu Wang, Lei Jin, Xin Zhang, Yinzhou Tang, Haisheng Su, et al. Worldarena: A unified benchmark for evaluating perception and functional utility of embodied world models, 2026.
- [32] Hengkai Tan, Yao Feng, Xinyi Mao, Shuhe Huang, Guodong Liu, Zhongkai Hao, Hang Su, and Jun Zhu. Anypos: Automated task-agnostic actions for bimanual manipulation, 2025.
- [33] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy, 2024.
- [34] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwei 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, Wenyan 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. [arXiv preprint arXiv:2503.20314](#), 2025.
- [35] Xiang Wang, Zhifei Zhang, He Zhang, Zhe Lin, Yuqian Zhou, Qing Liu, Shiwei Zhang, Yijun Li, Shaoteng Liu, Haitian Zheng, Jason Kuen, Yuehuan Wang, Changxin Gao, and Nong Sang. Hbridge: H-shape bridging of heterogeneous experts for unified multimodal understanding and generation, 2025.
- [36] Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chunling Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. [arXiv preprint arXiv:2602.15922](#), 2026.
- [37] Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-wam: Do world action models need test-time future imagination? [arXiv preprint arXiv:2603.16666](#), 2026.
- [38] Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, Yinan Zheng, Jiayin Zou, Yilun Chen, Jia Zeng, Ya-Qin Zhang, Jiangmiao Pang, Jingjing Liu, Tai Wang, and Xianyu Zhan. X-vla: Soft-prompted transformer as scalable cross-embodiment vision-language-action model, 2025.
- [39] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all, 2024.
- [40] Siyuan Zhou, Yilun Du, Jiaben Chen, Yandong Li, Dit-Yan Yeung, and Chuang Gan. Robodreamer: Learning compositional world models for robot imagination, 2024.
- [41] Chunling 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. In [Proceedings of Robotics: Science and Systems \(RSS\)](#), 2025.