UTILITY-WAM: CORRECTING FUTURE-INDUCED ACTION CHANGES

Anonymous authors
Paper under double-blind review

ABSTRACT

World Action Models (WAMs) can condition robot actions on generated visual futures, but this additional context is neither cost-free nor uniformly beneficial. We examine what a predicted future actually changes: the action delta Δa between the future-conditioned and future-free execution paths of the same model. Under a shared WAM checkpoint, this delta barely shifts average benchmark success ($+0.8$ points on LIBERO, -0.8 on RoboTwin 2.0), while generating visual futures increases inference latency by 1.7–1.9$\times$. These aggregate metrics conceal a concentrated failure mode that we term *future-induced negative transfer*: while predictive deltas benefit long-horizon tasks ($+2.4$ points on LIBERO-Long) and remain negligible on most routine manipulations, they induce severe degradation on precision tasks—with *Press Stapler* and *Place Can Basket* losing 54.0 and 34.0 points, recurring as top failure modes across independently trained WAM variants. We introduce **Utility-WAM** for in-place correction of future-conditioned control. Its primary learning component is a residual adapter that repairs future-induced velocity deltas under a frozen shared WAM by aligning them with ground-truth dynamic guidance and penalizing excess error beyond the future-free baseline. The adapter selectively eliminates the majority of severe task drops (reducing losses to 6.0 and 14.0 points) while preserving long-horizon foresight. To eliminate superfluous video computation, a lightweight pre-generation gate reuses counterfactual action outcomes to selectively bypass video generation on 75.5% of RoboTwin queries, cutting latency by 38% while improving average success to 93.3% on RoboTwin 2.0 and 98.9% on LIBERO. On physical bimanual manipulation across three real-world tasks, Utility-WAM succeeds in 53 of 60 trials, compared to 43 for FAST-WAM.

1 INTRODUCTION

World Action Models (WAMs) couple robot control with future visual prediction, but their deployment choices differ. DreamZero (Ye et al., 2026b), Vidar (Feng et al., 2025), and LingBot-VA (Li et al., 2026b) generate visual futures for action prediction; FAST-WAM (Yuan et al., 2026) omits future video at deployment; and AdaWAM (Tang et al., 2026) routes visual reasoning by execution phase. These approaches mainly ask whether or when to spend computation on visual prediction. We instead ask what a predicted future changes: the action delta Δa it induces relative to a matched future-free path. This delta may help, hurt, or be negligible, while its video-generation cost is always paid.

FAST-WAM already suggests this tension: its separately trained future-generating variants improve LIBERO but degrade RoboTwin 2.0 relative to future-free inference (Yuan et al., 2026). Our controlled comparison under one shared Stage-1 checkpoint reproduces this small sign reversal (Table 2), but the averages hide where the harm lies. Following the standard definition of negative transfer as deterioration caused by an additional information source (Wang et al., 2019; Zhang et al., 2023), we call deterioration relative to the future-free path *future-induced negative transfer* (Section 2) and find it to be a concentrated failure mode rather than a uniform effect.

We therefore treat the action delta Δa as the central object of study. Empirically, this delta provides meaningful guidance on long-horizon tasks, remains negligible on most standard manipulations (26 of 50 RoboTwin tasks change by at most one point), but induces severe failure on precision tasks

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

such as *Press Stapler* and *Place Can Basket* (hereafter the *two severe tasks*), which also recur as the most degraded tasks across FAST-WAM variants (Appendix C.2). Because pure routing cannot rectify a corrupted action vector once future conditioning is invoked, the foundational challenge is to *repair future-induced action changes while preserving the predictive foresight where it genuinely helps.*

To address this, we introduce **Utility-WAM**, which decouples test-time imagination into targeted action correction and selective compute allocation. Rather than discarding future generation or distilling it into a current-only policy—which compromises long-horizon performance—Utility-WAM retains the future-conditioned path while attaching a lightweight residual adapter to a frozen shared WAM. The adapter directly repairs harmful action shifts by aligning predicted velocity deltas with beneficial corrections supported by ground-truth futures. Furthermore, because physical imagination remains computationally expensive and unneeded on routine decisions, a pre-generation gate reuses the adapter's counterfactual action outcomes to selectively bypass video generation entirely when no action benefit is expected.

Our contributions are as follows:

- **Characterizing Future-Induced Negative Transfer:** We formalize test-time imagination via the action delta Δa relative to a matched future-free reference. We demonstrate that small benchmark-average effects mask severe, reproducible performance collapse on precision tasks across diverse WAM configurations.
- **Targeted Action Correction:** We propose a residual adapter that repairs future-induced velocity deltas via ground-truth future alignment and an excess-error penalty. On RoboTwin 2.0, the adapter recovers 89% and 59% of the losses on the two severe tasks (accounting for 80% of its overall 1.7-point gain) while preserving gains on the remaining 48 tasks and long-horizon suites.
- **Adaptive Deployment and Physical Validation:** By deriving a pre-generation utility gate from paired adapter outcomes, Utility-WAM invokes video prediction on only 24.5% of RoboTwin queries (cutting latency by 38%) and achieves reported averages of 98.9% on LIBERO and 93.3% on RoboTwin 2.0. On real-world bimanual hardware, it improves success from 71.7% to 88.3% (+16.6 pp).

2 PRELIMINARIES

Notation and Flow-Matching WAM. Let $x = (o, c, s)$ denote the current observation, instruction, and proprioceptive state, $a_{1:H}$ an action chunk, and V a visual future. The shared WAM jointly models video and action tokens with flow matching: for a clean target z, noise ϵ, and level τ, $z_\tau = (1-\tau)z + \tau\epsilon$ with target velocity $v^* = \epsilon - z$, trained with the standard flow-matching MSE. Its two action paths share parameters: π_{wo} blocks action access to future video, whereas π_{w} conditions on its generated future.

Future-Induced Negative Transfer. Treating the predicted future as the additional source of negative transfer (Wang et al., 2019), we measure any configuration π on evaluation episodes $\mathcal{D}$ by its success change relative to the future-free path,

$$\Delta J_{\mathcal{D}}(\pi) = J_{\mathcal{D}}(\pi) - J_{\mathcal{D}}(\pi_{\text{wo}}), \tag{1}$$

where $J_{\mathcal{D}}$ is closed-loop success. Future-induced negative transfer is $\Delta J_{\mathcal{D}}(\pi_{\text{w}}) < 0$; unlike classical transfer, both policies share parameters and data. ΔJ is the closed-loop outcome of the per-query action delta Δa (Eq. (3)). All tables report ΔJ in percentage points, so negative entries indicate harm.

3 METHOD

3.1 OVERVIEW: CORE CORRECTION AND SELECTIVE DEPLOYMENT

Utility-WAM's main learning component is a residual adapter that repairs the action changes induced by generated futures. A lightweight gate trained on the adapter's paired outcomes then selects

between this corrected path and the future-free path. We parameterize the shared WAM, adapter, and gate by θ, ϕ, and ω, respectively. Building on the WAM formulation in Section 2, let $\hat{v}_{\mathrm{w}}$ denote the future-conditioned velocity after adapter correction (Eq. (6)). At inference, the gate is evaluated once from the current context and selects one complete action flow:

$$\frac{\mathrm{d}\hat{z}_\tau}{\mathrm{d}\tau} = \begin{cases} \hat{v}_{\mathrm{w}}, & b_\omega(x) = 1, \\ v_{\mathrm{wo}}, & b_\omega(x) = 0, \end{cases} \qquad b_\omega(x) = \mathbf{1}[g_\omega(x) > \eta], \tag{2}$$

where $g_\omega(x)$ is the gate score. The future-free flow uses the future-blocked mask; the corrected future-conditioned flow uses the future-visible mask and applies the adapter to action velocities from evolving video features.

For paired branches sharing initial action noise, the N Euler updates collapse into the action delta

$$\Delta a = \hat{a}_{\mathrm{w}} - \hat{a}_{\mathrm{wo}} = \sum_{i=0}^{N-1} \Delta_i (\hat{v}_{\mathrm{w},i} - v_{\mathrm{wo},i}) \,, \tag{3}$$

where Δ_i is the i-th step and each velocity follows its own branch trajectory. Both branches are run only during training to form Δa; at inference only the selected branch runs. Stage 2 shapes the local terms, and a binary label derived from the adapter-corrected endpoint Δa supervises Stage 3.

Training proceeds in three stages. Stage 1 jointly trains the shared video and action DiTs under two attention masks: M_{wo} blocks action access to future-video tokens, whereas M_{w} permits it. Both paths retain current observations and share all parameters. We minimize

$$\mathcal{L}_1 = \mathcal{L}_{\mathrm{video}} + \lambda_{\mathrm{wo}} \mathcal{L}_{\mathrm{act}}^{\mathrm{wo}} + \lambda_{\mathrm{w}} \mathcal{L}_{\mathrm{act}}^{\mathrm{w}}, \tag{4}$$

where the video loss and two masked action losses are balanced by λ_{wo} and λ_{w}. Stage 2 freezes θ and trains only the adapter ϕ on the future-conditioned path, using the future-free velocity as its reference. Stage 3 freezes both components and trains the gate ω from paired final-action errors. Appendix A.1 specifies the masks and stage-wise sampling, and Figure 1 summarizes the three training stages and the resulting inference-time routing.

3.2 Correcting Future-Induced Action Changes

Utility-WAM first repairs future-induced action changes. With the expert residual $r_a^* = a_{1:H} - \hat{a}_{\mathrm{wo}}$ and D valid action entries, we measure the *utility* of an action delta Δa (Eq. (3)) by its paired action-error improvement:

$$u_{\mathrm{act}} = \ell(\hat{a}_{\mathrm{wo}}, a_{1:H}) - \ell(\hat{a}_{\mathrm{w}}, a_{1:H}) = \frac{2 \underbrace{\langle r_a^*, \Delta a \rangle}_{\text{alignment}} - \underbrace{\|\Delta a\|^2}_{\text{magnitude}}}{D}, \tag{5}$$

where ℓ is MSE. Negligible deltas give $u_{\mathrm{act}} \approx 0$ and misdirected ones give $u_{\mathrm{act}} < 0$. The adapter shapes both the direction and the magnitude of Δa during generation. For the same current inputs x, noisy action chunk z_τ, and noise level τ, let v_{wo} and v_{w} denote the action velocity predictions without and with predicted future video, respectively. Their difference $v_{\mathrm{w}} - v_{\mathrm{wo}}$ is the future-induced velocity delta; we reserve *correction* for the adapter residual introduced below. Because v_{w} need not have a lower mean squared error relative to the flow-matching target v^* than v_{wo}, this velocity delta is not intrinsically beneficial.

We therefore introduce a residual adapter to correct this velocity delta. Let q_τ^a denote the action queries at action noise level τ, and let F_k^+ denote the spatially pooled future-video features after the k-th video update. The adapter predicts a residual from these post-update features, yielding the refined velocity

$$\hat{v}_{\mathrm{w}} = v_{\mathrm{w}} + \Delta v_\phi(q_\tau^a, F_k^+). \tag{6}$$

To supervise this residual, the frozen WAM receives the same noisy expert action under no-future, predicted-future, and noised-ground-truth-future conditioning, producing v_{wo}, v_{w}, and v_{GT}. The adapter is trained on self-generated features F_k^+, matching its inference-time input distribution, while v_{GT} identifies whether the future supports a beneficial velocity delta. The matched-noise rollout and update sampling are detailed in Appendix A.1.

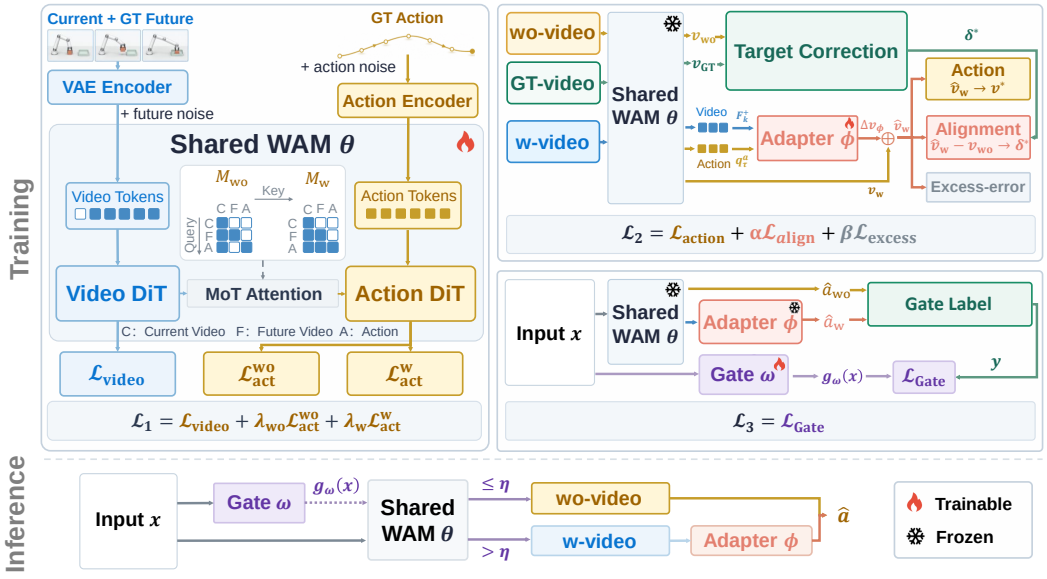

Figure 1: **Utility-WAM repairs future-induced action changes before allocating future-generation compute.** Stage 1 trains a shared WAM with matched future-free and future-visible action paths. Stage 2 freezes the WAM and trains a residual adapter on self-generated future features, using beneficial ground-truth-future velocity deltas and the future-free path as supervision. These two stages establish the core correction method. Stage 3 reuses the adapter's paired final-action outcomes to train a gate. At inference, the gate runs before video generation: it either executes the future-free path and skips that compute, or invokes future generation and the corrected future-conditioned path. Solid arrows denote inference data flow; dashed orange arrows denote training-only supervision.

Let $d(u, v)$ denote the mean squared error over valid action entries. With a relative improvement margin $\rho = 0.05$, we construct

$$\delta^* = \begin{cases} v_{\mathrm{GT}} - v_{\mathrm{wo}}, & d(v_{\mathrm{GT}}, v^*) < (1-\rho)\, d(v_{\mathrm{wo}}, v^*), \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

Using the refined velocity $\hat{v}_{\mathrm{w}}$ from Eq. (6), we optimize

$$\mathcal{L}_2 = \mathbb{E}\left[\underbrace{w(\tau)\, d(\hat{v}_{\mathrm{w}}, v^*)}_{\text{Action}} + \alpha \underbrace{d(\hat{v}_{\mathrm{w}} - v_{\mathrm{wo}}, \delta^*)}_{\text{Alignment}} + \beta \underbrace{\max\{d(\hat{v}_{\mathrm{w}}, v^*) - d(v_{\mathrm{wo}}, v^*),\, 0\}}_{\text{Excess-error}} \right]. \tag{8}$$

Here $w(\tau)$ is the FAST-WAM scheduler weight, and α, β weight alignment and excess error. The three objective terms jointly regulate the direction and magnitude of the action correction: the action term fits expert demonstration velocities, the alignment term steers $\hat{v}_{\mathrm{w}} - v_{\mathrm{wo}}$ toward the velocity shift supported by ground-truth futures (collapsing to zero when privileged futures offer no advantage), and the excess-error regularizer penalizes local flow-matching error exceeding the future-free baseline.

Why correction is applied to the future-conditioned path. Both execution paths share identical current observations, but the future-conditioned branch incorporates synthetic representations generated through noisy sampling. While test-time visual imagination provides dynamic foresight on complex horizons, imperfect video generation introduces distribution shifts that degrade action precision. Equation (8) leverages v_{wo} as an anchor: when ground-truth future conditioning yields negligible benefit, the adapter suppresses deviations from the reliable future-free reference; when future conditioning is demonstrably beneficial, it mirrors this correction in the spirit of professor forcing (Lamb et al., 2016). Unlike PFD, which distills future-induced velocity residuals into a current-only policy (Fang et al., 2026), this preserves visual foresight while constraining its negative side effects.

3.3 Skipping Futures that Correction Does Not Make Useful

Correction is soft: the excess-error term constrains local velocity error but does not make every endpoint correction useful, and the corrected branch still generates video on every query. The adapter's paired outcomes identify where this happens. Let $y = \mathbf{1}[u_{\text{act}} > 0]$, with u_{act} (Eq. (5)) evaluated on the adapter-corrected Δa. Because Δa is unavailable before video generation, a lightweight gate learns $g_\omega(x) \approx P(y = 1 \mid x)$ from current inputs with $\mathcal{L}_3 = \mathcal{L}_{\text{Gate}} = \text{BCE}(g_\omega(x), y)$ (Appendix A.1), so queries on which correction yields no improvement run the future-free path and skip video generation. Sweeping its threshold traces the accuracy–compute frontier in Figure 2; Appendix B gives the budgeted routing objective, default calibration, and hard-budget rule.

3.4 Inference

We evaluate the gate before video generation and execute only the selected path. The future-free path performs ten action updates; the future-conditioned path pairs each video update with an adapter-refined action update. We fix η from held-out validation data before test evaluation. Appendices A.1 and B give the solver, update order, and threshold-selection rule.

4 Experiments

Experiments first test the core prediction: correction should shrink the severe losses while retaining performance on the remaining tasks and the long-horizon gain. We then test whether the adapter-derived gate improves on always-on correction at lower compute and against compute-matched baselines.

4.1 Simulation Experiments

Setup. We build on FAST-WAM (Yuan et al., 2026) with a Wan2.2-TI2V-5B backbone and train one model per benchmark. All controlled variants share the same Stage-1 checkpoint, isolating the effects of adaptation and routing, and are evaluated on identical initial states and random seeds. LIBERO evaluation covers 40 tasks and 2,000 trials, executing 10 of 32 predicted actions; RoboTwin 2.0 uses 100 trials per task under clean and randomized conditions, executing 24 of 32 actions. Appendix A.2 and Appendix A.3 provide architecture, optimization, data, and rollout details.

4.2 Diagnosing Future-Induced Negative Transfer

Benchmark Level: A Small Average Effect. Relative to Shared WAM (wo), future conditioning changes average success by $+0.8$ points on LIBERO and -0.8 points on RoboTwin 2.0 (Table 2). The LIBERO gain is concentrated on Long ($+2.4$ points versus at most $+0.4$ elsewhere), and the RoboTwin decrease is not larger under randomization (-1.1 Clean, -0.5 Random). FAST-WAM's separately trained Joint and IDM variants show the same trend (Appendix C.2).

Task Level: Concentrated Severe Failures. The averages hide where the harm lies. Table 1 summarizes the task-level changes $\Delta J_t(\pi_{\text{w}})$. On RoboTwin 2.0, 21 of 50 tasks have $\Delta J_t < 0$, yet the median is 0.0: the aggregate loss is driven by *Press Stapler* (-54.0) and *Place Can Basket* (-34.0), while the other 19 harmed tasks lose at most 2.5 points each. These two tasks are also among the three most harmed for both FAST-WAM future variants (Appendix Table 9), so the severe failures recur across three future-usage mechanisms. We therefore treat them as the primary targets for correction; Appendix C.3 relates each task's loss to its adapter gain.

4.3 Core Correction and Selective Deployment

We evaluate mitigation in training order (Tables 2 and 3).

Core Result: Gains Concentrate Where Futures Hurt, Without Sacrificing Other Tasks. The adapter-loss rows in Table 3 isolate the objective with the shared WAM frozen, future generation

Table 1: **Task-level prevalence and severity of future-induced negative transfer.** Entries are ΔJ of Shared WAM (w) in percentage points (negative values indicate harm). "Harmed" counts tasks with $\Delta J_t < 0$. Entries are computed from the complete per-task simulation evaluation logs; Appendix Table 9 lists the ten most harmed RoboTwin 2.0 tasks.

Evaluation	Aggregate ΔJ	Harmed tasks	Median ΔJ_t	ΔJ_t range
LIBERO	$+0.8$	$7/40$	0.0	$[-6.0, +6.0]$
RoboTwin 2.0	-0.8	$21/50$	0.0	$[-54.0, +5.5]$

Table 2: **Controlled diagnosis and mitigation of future-induced negative transfer.** All entries are success rates (%) under a shared Stage-1 checkpoint. The fixed w-versus-wo comparison measures the aggregate effect of future conditioning. Utility-WAM (Correction) isolates the core correction method with future generation always enabled; Utility-WAM (Selective) adds the adapter-derived gate. Bold and underlined entries denote the best and second-best results in each column, respectively (including ties).

Method	LIBERO					RoboTwin 2.0		
	Spatial	Object	Goal	Long	Avg.	Clean	Random	Avg.
Shared WAM (wo)	97.2	99.0	<u>98.0</u>	95.4	97.4	92.4	91.3	91.9
Shared WAM (w)	97.6	<u>99.2</u>	<u>98.0</u>	97.8	98.2	91.3	90.8	91.1
Utility-WAM (Correction)	**98.6**	**100.0**	97.8	<u>98.4</u>	<u>98.7</u>	<u>93.3</u>	<u>92.3</u>	<u>92.8</u>
Utility-WAM (Selective)	<u>98.4</u>	**100.0**	**98.2**	**98.8**	**98.9**	**93.8**	**92.7**	**93.3**

enabled, and the gate disabled. Training the adapter on self-generated futures with action loss alone ($\mathcal{L}_{\text{action}}$), akin to scheduled sampling (Bengio et al., 2015), changes success by ≤ 0.3 points; this capacity-matched control confirms that added parameters alone cannot resolve future-induced degradation.

Instead, the gains stem from the complementary pairing of alignment and error regularization in Eq. (8). The alignment term ($\mathcal{L}_{\text{align}}$) provides constructive steering toward privileged corrections δ^*, recovering +0.9 points on RoboTwin 2.0 and +0.3 on LIBERO. Meanwhile, the excess-error penalty ($\mathcal{L}_{\text{excess}}$) serves as an asymmetric safety anchor: it remains largely inactive on routine tasks, but penalizes velocity drift exceeding the future-free baseline. Together, they curb residual degradation during noisy rollouts, contributing an essential +0.5 points on RoboTwin 2.0 (reaching 92.8%) and +0.2 on LIBERO (98.7%).

Under always-video routing, this combined objective produces the intended targeted-recovery pattern: gains concentrate where futures cause severe harm. The two severe tasks supply 1.36 of its 1.7-point RoboTwin gain (80%): their losses shrink from 54.0 to 6.0 and from 34.0 to 14.0 points, where the excess-error penalty crucially arrests catastrophic drift. The other 48 tasks still improve by $+0.35$ points on average, and LIBERO-Long keeps its gain (98.4% versus 97.8%). Thus, correction removes concentrated harm without sacrificing the benefit of futures (Appendix Table 9); LIBERO improves by 0.5 points.

Gate: Skipping Futures that Correction Does Not Make Useful. We fix the gate thresholds on the validation sets ($\eta = 0.439$ for LIBERO and 0.541 for RoboTwin 2.0; Appendix B). Figure 2 reports success against N_{eff} as we sweep the threshold, tracing the accuracy–compute frontier of Eq. (11).

At the calibrated operating points, the gate generates video on 87% of LIBERO queries and 24.5% of RoboTwin 2.0 queries ($N_{\text{eff}} = 8.70$ and 2.45) and raises success from 98.7% and 92.8% under always-on correction to 98.9% and 93.3%. The gain comes from queries that correction does not make useful: on the 48 RoboTwin 2.0 tasks outside the two severe ones, the gain over the future-free path rises from $+1.35$ to about $+1.9$ points, whereas the two severe tasks keep the recovery obtained by correction (Appendix C.3). On RoboTwin 2.0 the gate cuts latency by 38% relative to

Table 3: **Adapter-objective ablations and selective routing.** Success entries are benchmark averages (%) under a shared Stage-1 checkpoint. All rows share identical evaluation episodes and random seeds.

Configuration	Routing	Adapter objective	LIBERO	RoboTwin 2.0	N_{eff} LIB./RT
Shared WAM (wo)	Fixed: no video	–	97.4	91.9	0/0
Shared WAM (w)	Fixed: video	–	98.2	91.1	10/10
Action only	Fixed: video	Action	98.2	91.4	10/10
+ Alignment	Fixed: video	Action + Align.	98.5	92.3	10/10
+ Excess-error	Fixed: video	Action + Excess	98.3	91.8	10/10
Utility-WAM (Correction)	Fixed: video	Full	98.7	92.8	10/10
Utility-WAM (Selective)	Learned	Full	**98.9**	**93.3**	8.70/2.45

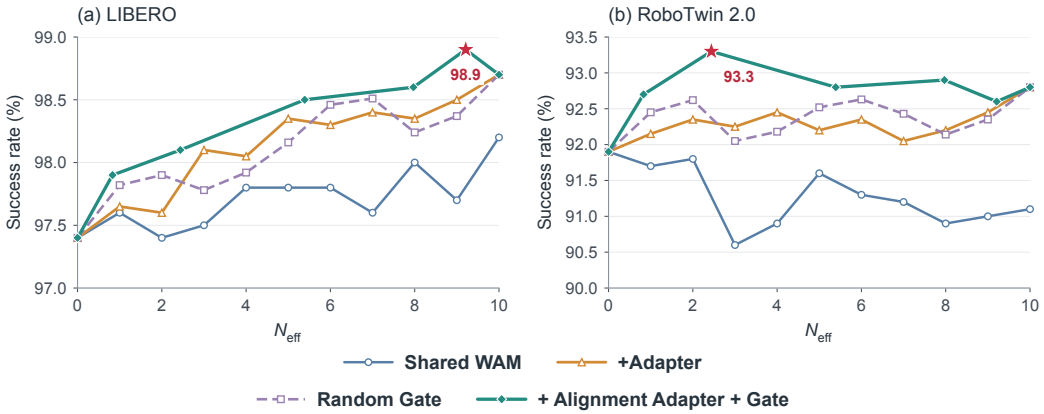

Figure 2: **Accuracy–compute frontier of selective routing.** Success rate versus effective video generation steps on LIBERO and RoboTwin 2.0; moving left saves video generation. N_{eff} denotes the average number of video generation steps per action query. Red stars mark the validation-calibrated operating points; the remaining points form a test-time threshold sweep. The selective configuration is compared with two compute-matched baselines. The correction-only curve is a static schedule that updates the future video only during the first n of the ten denoising steps of every query, so that $N_{\text{eff}} = n$. The dashed Random Gate curve routes each query to the corrected future-conditioned path with probability $N_{\text{eff}}/10$.

always-video inference, adding 17% over the future-free path for $+1.4$ points; on LIBERO, where futures help on average (Section 4.2), it keeps most queries on the future-conditioned path and saves 6.7% (Appendix Table 11).

On both benchmarks, the gate curve lies above adapter-only static truncation and random routing at every evaluated intermediate budget ($0 < N_{\text{eff}} < 10$). Random routing reaches at most 98.4% on LIBERO ($N_{\text{eff}} \in \{8, 9\}$) and 92.6% on RoboTwin 2.0 ($N_{\text{eff}} \in \{2, 3\}$), below the gate at its operating points, and truncating video generation on every query also remains below the gate. The learned gate therefore improves the measured accuracy–compute frontier over these two compute-matched controls. Unlike phase-supervised routing (Tang et al., 2026), the gate needs no phase labels and applies unchanged to variable-length real-world episodes.

4.4 Comparison with Reported Results

Table 4 compares Utility-WAM with representative VLAs and WAMs, with FAST-WAM as the primary reported comparison because we follow its data and evaluation pipeline. Its future-generating Joint and IDM variants raise LIBERO-Long from 95.2% to 96.8% and 97.6%, respectively, but lower RoboTwin 2.0 from 91.8% to 90.6% and 91.3%. The adapter alone already exceeds every re-



(a) AgileX PiPER

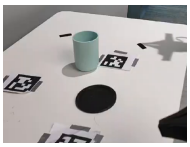

(b) Place Empty Cup



(c) Insert Can

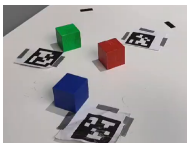

(d) Stack Three Blocks

Figure 3: **Real-world platform and tasks.** (a) Bimanual platform (image: AgileX Robotics). (b–d) Initial task configurations.

Table 5: **Real-world success counts** (20 trials per task) on the tasks in Figure 3b–d. Shared WAM and Utility-WAM (Correction) follow the configurations of Table 2 and are evaluated under the same protocol. Bold and underlined entries denote the best and second-best results in each column, respectively (including ties).

Method	Cup	Can	Blocks	Total
FAST-WAM	**18/20**	14/20	11/20	43/60 (71.7%)
Shared WAM (wo)	17/20	16/20	14/20	47/60 (78.3%)
Shared WAM (w)	16/20	16/20	13/20	45/60 (75.0%)
Utility-WAM (Correction)	**18/20**	17/20	15/20	50/60 (83.3%)
Utility-WAM (Selective)	**18/20**	**18/20**	**17/20**	**53/60 (88.3%)**

ported LIBERO average (98.7%) and is within 0.1 points of the best reported RoboTwin 2.0 average (92.8% versus 92.9%); with the gate, Utility-WAM reaches 98.9% and 93.3%, retaining the long-horizon benefit of generated futures while removing their concentrated harm. While cross-paper comparisons involve independently trained models, the controlled ablations in Table 2 isolate the exact contribution of each component under a shared backbone.

Table 4: **Comparison with reported results** on LIBERO and RoboTwin 2.0. † marks values not reported in the cited papers. Bold/underline: best/second-best (including ties). Sources are in Appendix A.4.

	LIBERO					RoboTwin 2.0		
Method	Spatial	Object	Goal	Long	Avg.	Clean	Random	Avg.
$\pi_{0.5}$† (Physical Intelligence et al., 2025)	98.8	98.2	98.0	92.4	96.9	82.7	76.8	79.8
LingBot-VA (Li et al., 2026b)	98.5	99.6	97.2	98.5	98.5	92.9	91.6	92.2
Cosmos-Policy (Kim et al., 2026)	98.1	**100.0**	98.2	97.6	98.5	–	–	–
Efficient-WAM (Li et al., 2026a)	–	–	–	–	–	86.7	85.7	86.2
AHA-WAM (Cai et al., 2026)	–	–	–	–	–	93.4	92.2	92.8
AdaWAM (Tang et al., 2026)	98.0	99.6	97.1	**99.1**	98.5	93.1	91.4	92.2
FAST-WAM (Yuan et al., 2026)	98.2	**100.0**	97.0	95.2	97.6	91.9	91.8	91.8
FAST-WAM-Joint (Yuan et al., 2026)	**99.6**	99.4	98.2	96.8	98.5	90.8	90.3	90.6
FAST-WAM-IDM (Yuan et al., 2026)	98.8	97.8	97.8	97.6	98.0	91.2	91.3	91.3
PFD (Fang et al., 2026)	98.6	99.2	**99.2**	95.4	98.1	93.1	**92.7**	92.9
Utility-WAM (Correction)	98.6	**100.0**	97.8	98.4	98.7	93.3	92.3	92.8
Utility-WAM (Selective)	98.4	**100.0**	98.2	98.8	**98.9**	**93.8**	**92.7**	**93.3**

4.5 Real-World Evaluation

Setup. Figure 3 shows the bimanual platform and task configurations. We use two AgileX PiPER arms with three RGB views and evaluate three tasks: *Place Empty Cup*, *Insert Can*, and *Stack Three Blocks*. Each task has 200 teleoperated demonstrations. FAST-WAM and Utility-WAM (Selective) are trained on the same demonstrations, with Utility-WAM following the three-stage procedure of Section 3, and each method is evaluated on 20 trials per task with varied initial layouts. Hardware, task definitions, and data collection details are in Appendices D.1 and D.2.

Results. Utility-WAM (Selective) succeeds in 53 of 60 trials versus 43 for FAST-WAM (Table 5), a 16.7-point improvement with an unadjusted two-sided 95% Newcombe interval of $[2.3, 30.4]$ points. The two methods tie on *Place Empty Cup*; the gains come from *Insert Can* ($+4$ successes), which requires aligning the can with a recess, and *Stack Three Blocks* ($+6$), the longest multi-stage task, consistent with the simulation finding that generated futures help most on long-horizon tasks. This

evaluation demonstrates that selective visual reasoning transfers effectively to physical hardware under matching demonstration budgets.

5 Related Work

World Action Models and Test-Time Futures. WAMs learn actions together with future visual prediction, either jointly (Wu et al., 2024; Cheang et al., 2024; Cen et al., 2025; Guo et al., 2024; Zhu et al., 2025; Bi et al., 2026; Kim et al., 2026; Ye et al., 2026b) or through predictive representations and inverse dynamics (Tian et al., 2025; Hu et al., 2024; Li et al., 2026b). Their deployment choices differ: FAST-WAM omits future generation, GigaWorld-Policy makes it optional, AHA-WAM reuses video context, Efficient-WAM reduces its cost, and AdaWAM routes it using phase supervision (Yuan et al., 2026; Ye et al., 2026a; Cai et al., 2026; Li et al., 2026a; Tang et al., 2026). We instead study the induced action change: we first repair it, then reuse the adapter's paired outcomes, rather than phase supervision, to decide where futures are generated.

Privileged Futures and Distillation. PFD (Fang et al., 2026) distills ground-truth-future corrections into a current-only policy, sacrificing predictive context at inference. In contrast, Utility-WAM preserves test-time visual imagination on long-horizon horizons while using adaptive gating to avert negative transfer on precision subtasks.

Imperfect Predictions and Negative Transfer. Learned rollouts are useful but fallible: I2A learns to interpret them, while STEVE and MBPO limit their use (Racanière et al., 2017; Buckman et al., 2018; Janner et al., 2019). Training on ground-truth inputs but deploying on model predictions creates exposure bias or covariate shift (Bengio et al., 2015; Ross et al., 2011; Ning et al., 2023); generated intermediate reasoning can likewise hurt in robotics and language models (Zhao et al., 2025; Liu et al., 2025). Following negative-transfer definitions (Wang et al., 2019; Zhang et al., 2023), we treat the predicted future as the additional source while holding parameters and data fixed. Feature contamination (Zhang et al., 2024) offers a related representation-level account.

6 Conclusion and Discussion

Visual futures in World Action Models induce an action delta Δa relative to future-free execution. While this predictive delta provides essential guidance on long-horizon tasks, it remains largely redundant on standard manipulations and induces severe failures on precision tasks—all while imposing significant inference latency on every decision.

Utility-WAM resolves this tension by decoupling test-time visual imagination into targeted action correction and selective deployment. A residual adapter attached to a frozen shared WAM directly aligns predicted velocity shifts with ground-truth guidance, eliminating the majority of severe task drops while preserving long-horizon gains. A lightweight utility gate, trained on paired counterfactual outcomes, selectively bypasses future generation when no tangible action benefit is anticipated, reducing video generation by up to 75.5% and latency by 38%.

Generality. The phenomenon of future-induced negative transfer recurs across diverse WAM paradigms, including joint denoising and inverse dynamics. Our correction framework applies broadly to any architecture with dual action paths and a flow-matching action head: joint-denoising WAMs (Ye et al., 2026b; Kim et al., 2026) expose these paths naturally through attention masks, while inverse-dynamics models (Li et al., 2026b; Feng et al., 2025) can attach the adapter directly to the action head.

Limitations. Our evaluations use a common Stage-1 checkpoint due to the computational cost of pretraining 5B diffusion backbones; while cross-task consistency and threshold robustness are verified, variance across independently retrained backbones remains uncharacterized. Furthermore, the utility gate relies on an offline action-error proxy rather than closed-loop task returns, leaving room for future work to explore online reinforcement-learned deployment policies.

AI USE STATEMENT

In this work, we used generative AI tools for language polishing and improving the readability of the manuscript. We have reviewed all AI-assisted work to ensure that the revisions preserve the intended meaning and technical content. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015.

Hongzhe Bi, Hengkai Tan, Shenghao Xie, Zeyuan Wang, Shuhe Huang, Haitian Liu, Ruowen Zhao, Yao Feng, Chendong Xiang, Yinze Rong, et al. Motus: A unified latent action world model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 35101–35113, 2026.

Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Jisong Cai, Long Ling, Shiwei Chu, Zhongshan Liu, Jiayue Kang, Zhixuan Liang, Wenjie Xu, Yinan Mao, Weinan Zhang, Xiaokang Yang, et al. AHA-WAM: Asynchronous horizon-adaptive world-action modeling with observation-guided context routing. *arXiv preprint arXiv:2606.09811*, 2026.

Jun Cen, Chaohui Yu, Hangjie Yuan, Yuming Jiang, Siteng Huang, Jiayan Guo, Xin Li, Yibing Song, Hao Luo, Fan Wang, et al. WorldVLA: Towards autoregressive action world model. *arXiv preprint arXiv:2506.21539*, 2025.

Chi-Lam Cheang, Guangzeng Chen, Ya Jing, Tao Kong, Hang Li, Yifeng Li, Yuxiao Liu, Hongtao Wu, Jiafeng Xu, Yichu Yang, et al. GR-2: A generative video-language-action model with web-scale knowledge for robot manipulation. *arXiv preprint arXiv:2410.06158*, 2024.

Tianxing Chen, Zanxin 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.

Pengcheng Fang, Hongli Chen, and Xiaohao Cai. Privileged foresight distillation: Zero-cost future correction for world action models. *arXiv preprint arXiv:2604.25859*, 2026.

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. URL https://arxiv.org/abs/2507.12898.

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. *Advances in Neural Information Processing Systems*, 37:112386–112410, 2024.

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.

Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

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.

Alex M. Lamb, Anirudh Goyal, Ying Zhang, Saizheng Zhang, Aaron Courville, and Yoshua Bengio. Professor forcing: A new algorithm for training recurrent networks. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Jiajun Li, Tiecheng Guo, Yifan Ye, Rongyu Zhang, Xiaowei Chi, Qianpu Sun, Ying Li, Yunfan Lou, Yan Huang, Zhihe Lu, et al. Efficient-WAM: A 1b-parameter world-action model with low-cost future imagination. *arXiv preprint arXiv:2606.10040*, 2026a.

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*, 2026b.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.

Ryan Liu, Jiayi Geng, Addison J. Wu, Ilia Sucholutsky, Tania Lombrozo, and Thomas L. Griffiths. Mind your step (by step): Chain-of-thought can reduce performance on tasks where thinking makes humans worse. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, 2025.

Mang Ning, Enver Sangineto, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Input perturbation reduces exposure bias in diffusion models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 2023.

Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: A vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.

Sébastien Racanière, Théophane Weber, David P. Reichert, Lars Buesing, Arthur Guez, Danilo Rezende, Adrià Puigdomènech Badia, Oriol Vinyals, Nicolas Heess, Yujia Li, et al. Imagination-augmented agents for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.

Yinzhou Tang, Jingbo Xu, Yu Shang, Zihao Song, Chen Gao, Wei Wu, and Yong Li. Dreaming when necessary: Advancing world action models with adaptive multi-modal reasoning. *arXiv preprint arXiv:2606.07089*, 2026.

Yang Tian, Sizhe Yang, Jia Zeng, Ping Wang, Dahua Lin, Hao Dong, and Jiangmiao Pang. Predictive inverse dynamics models are scalable learners for robotic manipulation. In *International Conference on Learning Representations*, volume 2025, pp. 92033–92052, 2025.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. Characterizing and avoiding negative transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11293–11302, 2019.

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 *International Conference on Learning Representations*, volume 2024, pp. 10641–10662, 2024.

Angen Ye, Boyuan Wang, Chaojun Ni, Guan Huang, Guosheng Zhao, Hao Li, Hengtao Li, Jie Li, Jindi Lv, Jingyu Liu, et al. GigaWorld-Policy: An efficient action-centered world–action model. *arXiv preprint arXiv:2603.17240*, 2026a.

Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026b.

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.

Tianren Zhang, Chujie Zhao, Guanyu Chen, Yizhou Jiang, and Feng Chen. Feature contamination: Neural networks learn uncorrelated features and fail to generalize. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 60446–60495. PMLR, 2024. URL https://proceedings.mlr.press/v235/zhang24cj.html.

Wen Zhang, Lingfei Deng, Lei Zhang, and Dongrui Wu. A survey on negative transfer. *IEEE/CAA Journal of Automatica Sinica*, 10(2):305–329, 2023.

Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. CoT-VLA: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1702–1713, 2025.

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. *arXiv preprint arXiv:2504.02792*, 2025.

A Implementation Details

A.1 Detailed Three-Stage Training and Inference

Shared-WAM masks. Let C, F, and A denote current-frame, future-video, and action tokens, with each relation interpreted as query-to-key access. Following FAST-WAM, M_{wo} permits $C \to C$, $F \to \{C, F\}$, and $A \to \{C, A\}$, while M_{w} additionally permits $A \to F$. At future-free inference, F is omitted and the remaining tokens use the corresponding submask. Both training paths retain the video objective.

Adapter sampling. For each adapter update, we sample one of the ten video denoising updates and roll out the frozen video model through it to obtain post-update features F_k^+. The ground-truth future is noised to the matched video level with the same initial video noise, while the action noise level follows the action flow-matching scheduler. The no-future, predicted-future, and ground-truth-future paths receive the same noisy expert action. The adapter uses q_τ^a as queries and F_k^+ as keys and values, followed by a feed-forward network and a zero-initialized output projection.

Gate labels. Each gate-training input contributes one paired action sample from the two frozen paths. They share Gaussian action noise, solver, and action schedule; the future-conditioned path additionally draws video noise. The gate is trained after adapter correction. Subtracting and summing the paired Euler updates cancels the common initial noise and gives Eq. (3); the resulting Δa supervises Stage 3.

Inference update. Starting from Gaussian noise, we solve Eq. (2) with the ten-step explicit-Euler solver. In the future-conditioned path, each iteration updates the video latent, extracts F_k^+, refines the action velocity, and then updates the action latent. The future-free path skips video updates. The exact schedule is given below, and Appendix B specifies gate calibration.

A.2 Base Model and Training Process

Base Model. The Wan backbone (Wan et al., 2025) models videos in a spatiotemporal latent space. A pretrained VAE encodes video frames into compact representations, and the video DiT performs conditional generation in this space. We retain the pretrained video weights and initialize the action DiT backbone by interpolating compatible video DiT parameters to its target dimensions. During training, the VAE remains frozen, and language embeddings are precomputed and cached.

Training Hyperparameters. Table 6 summarizes the model, inference, and three-stage training settings. For gate training, we retain the checkpoint with the lowest validation loss and apply early stopping with a patience of three epochs and a minimum improvement of 10^{-4}.

Adapter Architecture. The residual adapter uses a single cross-attention block, where projected action features attend to spatially pooled future-video features, followed by a residual feed-forward network and a zero-initialized output projection that predicts action-velocity corrections. For LIBERO, we use a hidden dimension of 256 and a $256 \to 512 \to 256$ feed-forward network, resulting in approximately 1.58M trainable parameters. For RoboTwin 2.0, we increase the hidden dimension to 512 and use a $512 \to 1024 \to 512$ feed-forward network, resulting in approximately 4.21M trainable parameters. The adapter parameters are shared across denoising steps and are used only in the future-conditioned branch.

Gate Architecture. The routing gate encodes each current camera view using a shared lightweight residual CNN. Each view is pooled into a 4×4 grid of 128-dimensional visual tokens, augmented with spatial position and camera embeddings. Language and proprioceptive features are fused into a 128-dimensional query, which attends to the visual tokens through a single four-head cross-attention block. A residual feed-forward network and an MLP prediction head produce a scalar logit, converted by a sigmoid into $g_\omega(x)$. We use the same architecture on both benchmarks, with approximately 1.3M trainable parameters. The gate operates exclusively on current inputs and makes its routing decision before future generation.

Table 6: Model and training hyperparameters. Paired values denote LIBERO / RoboTwin 2.0. Global batch sizes include gradient accumulation.

Model and inference settings			
Video backbone	Wan2.2-TI2V-5B		
Action DiT depth L_a	30		
Action DiT hidden dimension d_a	1024		
Action horizon H	32		
Video frames per sample T	9		
Multi-view fusion	Concatenation before VAE encoding		
Adapter fusion	Cross-attention (action queries, video keys/values)		
Adapter feature dimension d_{adapt}	256 / 512		
Adapter output projection initialization	Zero		
Gate feature fusion	MLP		
Action denoising steps N_a	10		
Future-video updates N_v	$\{0, 10\}$, selected by the gate		
Training hyperparameter	Stage 1 Shared WAM	Stage 2 Adapter	Stage 3 Gate
Optimizer	AdamW	AdamW	AdamW
Learning rate	10^{-4}	10^{-4}	10^{-4}
Weight decay	10^{-2}	10^{-4}	10^{-4}
Adam (β_1, β_2)	$(0.9, 0.95)$	$(0.9, 0.95)$	$(0.9, 0.999)$
Global batch size	128 / 128	48 / 96	64 / 64
Training budget (LIBERO)	10 epochs	30,000 steps	Up to 20 epochs
Training budget (RoboTwin)	5 epochs	20,000 steps	Up to 20 epochs
Learning-rate schedule	Cosine	Cosine	Constant
Warmup ratio	5%	5%	–
Gradient clipping norm	1.0	1.0	1.0
Training precision	BF16	BF16	FP32

Flow-matching schedule and sampling order. We follow the released FAST-WAM continuous flow scheduler rather than sampling a logit-normal inference grid. Let $S = 1000$ denote the number of training timesteps. We use the original paper and released-checkpoint setting $s = 5.0$ for both the video and action modalities; the later repository revision that changes the default action shift to 1.0 is not used in our experiments. The scheduler maps $u \in [0, 1]$ to the noise level

$$\phi_s(u) = \frac{s\,u}{1 + (s-1)u}, \qquad \tau = \phi_s(u), \qquad t = S\tau. \tag{9}$$

During training, the video and action branches draw their own $u \sim \mathcal{U}(0, 1)$ independently. During inference with $N = 10$ steps, we set $u_i = 1 - i/N$, $\tau_i = \phi_s(u_i)$, and $\Delta_i = \tau_{i+1} - \tau_i$ for $i = 0, \ldots, N-1$. FAST-WAM uses the explicit-Euler update

$$z_{i+1} = z_i + \Delta_i f_\theta(z_i, S\tau_i; \mathcal{C}), \tag{10}$$

where f_θ is the velocity predictor and $\mathcal{C}$ its conditioning. The no-video branch applies these ten updates only to the action latent. The future-conditioned branch pairs video and action updates by the common index i: it updates the video latent first, extracts F^+_{i+1}, applies the adapter, and then updates the action latent. During adapter training, we instead sample a video update k uniformly from $\{1, \ldots, 10\}$ and sample the action noise level independently. Classifier-free guidance uses scale 1.0. Each sample contains 32 actions and nine video frames, including the current frame, so adjacent video frames correspond to four action steps.

Loss weights. For shared-WAM training, we set $\lambda_{\text{wo}} = \lambda_{\text{w}} = 0.5$ in Eq. (4). For adapter training, we set the alignment weight to $\alpha = 1.0$ and the excess-error weight to $\beta = 0.5$ in Eq. (8). The relative improvement margin in Eq. (7) is $\rho = 0.05$. The scheduler weight in Eq. (8) is the FAST-WAM training weight $w(\tau) \propto r(\tau) - \min_i r(\tau_i)$ with $r(\tau) = \exp[-2(\tau - \frac{1}{2})^2]$, normalized to mean one over the discrete training grid $\tau_i = \phi_s(1 - i/S)$, $i = 0, \ldots, S-1$.

A.3 Simulation Benchmarks

LIBERO. We follow the FAST-WAM data and evaluation pipeline on the Spatial, Object, Goal, and Long suites of LIBERO (Liu et al., 2023). Each suite contains 500 demonstrations spanning ten tasks. The model receives a fixed external view and a wrist view, each resized to 224×224 and concatenated horizontally before VAE encoding. Actions are seven-dimensional and proprioceptive states are eight-dimensional; action and state normalization follows the FAST-WAM min–max preprocessing. A sample contains 32 actions and nine video frames. At evaluation, the policy samples one chunk, executes its first ten actions, and replans, for a total of 2,000 trials across the 40 tasks.

RoboTwin 2.0. We follow the FAST-WAM multi-task setup on RoboTwin 2.0 (Chen et al., 2025), using 2,500 demonstrations from clean scenes and 25,000 demonstrations with heavy scene randomization over more than 50 bimanual tasks. Inputs comprise a high camera and two wrist cameras, each resized to 240×320 and combined using the FAST-WAM three-camera layout. Actions and proprioceptive states are both 14-dimensional and use the released z-score statistics. Each sample again contains 32 actions and nine video frames. At evaluation, the policy samples one chunk, executes its first 24 actions, and replans. We report 100 trials per task under both clean and randomized conditions.

A.4 Reported Result Sources

LIBERO provenance. LingBot-VA (Li et al., 2026b), Cosmos Policy (Kim et al., 2026), AdaWAM (Tang et al., 2026), the three FAST-WAM variants (Yuan et al., 2026), and PFD (Fang et al., 2026) are copied from their respective original papers. The cited $\pi_{0.5}$ paper does not report LIBERO; its values are taken from the authors' official OpenPI LIBERO release[1].

RoboTwin 2.0 provenance. LingBot-VA (Li et al., 2026b), Efficient-WAM (Li et al., 2026a), AHA-WAM (Cai et al., 2026), AdaWAM (Tang et al., 2026), the three FAST-WAM variants (Yuan et al., 2026), and PFD (Fang et al., 2026) are taken from their latest audited original-paper tables;

[1] Official $\pi_{0.5}$ LIBERO results.

PFD's Clean and Random scores (93.11 and 92.69) are rounded to one decimal place. The cited $\pi_{0.5}$ paper does not contain RoboTwin experiments; its values are traceable to LingBot-VA's Appendix comparison table (Li et al., 2026b) and are marked with a dagger. Utility-WAM values are produced by our evaluation protocol. Published averages are retained when available; otherwise, we compute the arithmetic mean of the displayed Clean and Random scores before rounding to one decimal place.

Tables 2 and 3 provide the controlled comparisons for method attribution.

B Gate Threshold Selection

Budgeted routing objective. Under a budget B on the fraction of queries that generate video (equivalently, $N_{\text{eff}} = 10\,\mathbb{E}_x[b(x)]$ video steps per query), the gate of Section 3.3 routes queries toward corrections that are more likely to reduce action error:

$$\max_{b:\mathcal{X}\to\{0,1\}} \mathbb{E}_x\big[b(x)\big(2y(x)-1\big)\big] \quad \text{s.t.} \quad \mathbb{E}_x[b(x)] \le B. \tag{11}$$

With $p(x) = P(y = 1 \mid x)$, Eq. (11) is solved by routing the queries with the largest $p(x)$ among those with $p(x) > 1/2$, up to budget B, so a score that ranks queries by $p(x)$ suffices.

Validation labels. We determine the routing threshold η exclusively from the gate validation set, which is disjoint from the gate-training inputs and from all test episodes, with the WAM, adapter, and gate frozen. For each validation input x_i, we draw one paired future-conditioned/no-video action sample with the shared action noise of Section 3.3, compute $u_{\text{act},i}$ by Eq. (5), set $y_i = \mathbf{1}[u_{\text{act},i} > 0]$, and record the gate logit $h_i = \log[g_\omega(x_i)/(1 - g_\omega(x_i))]$. Each validation input thus contributes exactly one label–logit pair; no test data are used.

Hard budget. Because binary cross-entropy is a proper scoring rule for y, ranking $g_\omega(x)$ approximates ranking the conditional probability of a beneficial branch correction. For a prescribed budget B, we use the validation quantile

$$\eta_B = \max\left\{\frac{1}{2}, Q_{1-B}(\{g_\omega(x_i)\}_{x_i\in\mathcal{D}_{\text{val}}})\right\}, \qquad b_{\omega,B}(x) = \mathbf{1}[g_\omega(x) > \eta_B], \tag{12}$$

which bounds the empirical validation routing rate by B (up to ties) while excluding scores that do not predict a beneficial correction.

Default operating point. Without an external budget, we fit Gaussians with a shared variance to the logits of the two classes, $h \mid y = j \sim \mathcal{N}(\mu_j, \sigma^2)$ with μ_j the class-mean logit, and apply the Bayes rule under equal priors and misclassification costs. Assuming $\mu_1 > \mu_0$, the decision boundary is the midpoint of the class means, independent of σ^2, so

$$\eta = \left[1 + \exp\left(-\frac{\mu_0 + \mu_1}{2}\right)\right]^{-1}. \tag{13}$$

This rule does not require calibrated gate probabilities. Calibration uses 5,408 LIBERO and 10,892 RoboTwin 2.0 validation inputs. The fitted logit boundaries are $h^* = -0.2452$ and $h^* = 0.1644$, which yield $\eta = 0.439$ and $\eta = 0.541$, respectively, and each threshold is kept fixed throughout test evaluation. A deployment with a hard latency budget instead uses Eq. (12), moving along the frontier in Figure 2.

C Simulation Setup and Additional Results

C.1 Datasets and Evaluation Protocols

Validation Set and Threshold Selection. We reserve a task-stratified validation set whose trajectories are disjoint from those used to train the gate. For LIBERO, this set contains 169 trajectories, with 32 temporally distributed observation points per trajectory, yielding 5,408 samples. For RoboTwin 2.0, it contains 2,723 trajectories, with four observation points per trajectory, yielding 10,892 samples. Appendix B describes how the routing threshold is determined from this set; no test rollouts are used.

Closed-Loop Evaluation. Each policy query predicts a chunk of 32 actions using 10 action denoising steps. The policy executes the first 10 actions on LIBERO and the first 24 actions on RoboTwin 2.0 before obtaining a new observation and replanning. At each replanning step, the gate applies the fixed benchmark-specific threshold to determine whether to generate future observations. Task success is determined by the benchmark-provided success checks. Each rollout terminates upon success or upon reaching the task-dependent execution limit in the FAST-WAM evaluation implementation. All configurations are evaluated on the same episodes: for each task and condition, the i-th trial uses the same initial state (object layout and scene randomization) and the same random seeds under every configuration, so all between-configuration comparisons are paired at the episode level. We report success rates measured over the closed-loop evaluation rollouts.

C.2 Replication Across Future-Usage Mechanisms

Table 7 compares our shared-checkpoint sign reversal with the two future-generating variants of FAST-WAM (Yuan et al., 2026): Joint, which jointly denoises future video and actions, and IDM, which generates a future before an inverse-dynamics model. In all three comparisons, future conditioning raises the LIBERO average, with the largest gain on Long, and lowers RoboTwin 2.0 success under both clean and randomized conditions. The FAST-WAM comparisons are not controlled: each future-generating variant is trained separately, so they confound future conditioning with training differences. We use them only to show that the sign reversal is not specific to our checkpoint and rely on the shared-checkpoint comparison for attribution.

Table 7: **Replication of the benchmark-level sign reversal.** Entries are ΔJ, future-conditioned minus future-free success in percentage points; negative values indicate harm. FAST-WAM rows use the published results of Yuan et al. (2026), comparing separately trained models; the last row is the shared-checkpoint comparison of Table 2.

		LIBERO		RoboTwin 2.0		
Comparison	Stage 1	Avg.	Long	Clean	Random	Avg.
FAST-WAM-Joint − FAST-WAM	Separate	$+0.9$	$+1.6$	-1.0	-1.5	-1.3
FAST-WAM-IDM − FAST-WAM	Separate	$+0.4$	$+2.4$	-0.7	-0.4	-0.6
Shared WAM (w) − Shared WAM (wo)	Shared	$+0.8$	$+2.4$	-1.1	-0.5	-0.8

Using the published per-task RoboTwin 2.0 results of FAST-WAM, we compute ΔJ_t for each of the 50 tasks, averaged over clean and randomized conditions (Table 8). For both variants, future generation helps on more tasks than it hurts, and the median task change is non-negative; the aggregate is nonetheless negative because a few tasks fail severely. The task-level changes of the two variants are strongly rank-correlated (Spearman 0.84), and *Press Stapler* and *Place Can Basket* are among the three most harmed tasks for both. Changes estimated separately under clean and randomized conditions, which share no evaluation episodes, are also positively correlated (Spearman 0.42 for Joint and 0.40 for IDM). These statistics are descriptive, but they indicate that harm recurs on a subset of tasks rather than spreading uniformly, so future use should be corrected where harm occurs rather than enabled or disabled globally.

Table 8: **Task-level negative transfer in FAST-WAM on RoboTwin 2.0.** ΔJ_t is the future-generating variant minus FAST-WAM, averaged over clean and randomized success, computed from the published per-task results of Yuan et al. (2026). Negative values indicate harm.

Variant	Harmed	Helped	Tied	Median	≥ 5 harmed / helped	Most harmed task
FAST-WAM-Joint	12/50	26/50	12/50	$+0.5$	6 / 5	Open Microwave (-45.0)
FAST-WAM-IDM	17/50	20/50	13/50	0.0	5 / 8	Press Stapler (-40.0)

C.3 Task-Level Negative-Transfer Analysis

Metrics. For task t, we compute $\Delta J_t(\pi) = J_t(\pi) - J_t(\pi_{\text{wo}})$ against the future-free branch of the same Stage-1 checkpoint. We report the number of tasks with $\Delta J_t(\pi_{\text{w}}) < 0$, the median across

tasks, and the range. The sign count is a descriptive prevalence statistic: with finite rollouts, a negative point estimate is not by itself evidence that an individual task has nonzero population ΔJ_t.

Per-task changes on RoboTwin 2.0. Table 9 lists ΔJ_t for the ten most harmed RoboTwin 2.0 tasks, alongside FAST-WAM's separately trained Joint and IDM variants, and summarizes all 50 tasks. Our three columns share the future-free reference, so the change from the always-video column to the correction and selective columns directly measures how much of the loss each component removes.

Why the per-task columns are highly consistent. All configurations are evaluated on identical initial states and random seeds (Appendix C.1), so each entry is a paired difference over the same 200 episodes and a step of 0.5 points is a single episode. Because the adapter and the gate leave the Shared WAM frozen, most episodes end with the same outcome under every configuration: across all 50 tasks, adjacent columns differ by at most four episodes (2.0 points) except on the two severe tasks. The paired design removes variation in initial states and sampling noise, but not training-seed variation, so small per-task differences should be read as descriptive trends.

Table 9: **Per-task success changes on RoboTwin 2.0: the ten most harmed tasks.** All entries are ΔJ_t in percentage points relative to the future-free reference, averaged over clean and randomized conditions; negative values indicate harm. Tasks are ranked by the mean ΔJ_t over the three future-conditioned comparisons (Joint, IDM, and Shared WAM (w)), most negative first, with ties broken by Shared WAM (w). FAST-WAM columns use the published per-task results of Yuan et al. (2026) with FAST-WAM as reference. Our columns use Shared WAM (wo) as the common reference and are paired over the same episodes. The bottom rows summarize all 50 tasks.

	FAST-WAM		Ours		
Task	Joint	IDM	Shared WAM (w)	Correction	Selective
Press Stapler	−42.5	−40.0	−54.0	−6.0	−6.5
Place Can Basket	−33.5	−37.5	−34.0	−14.0	−17.0
Open Microwave	−45.0	0.0	−1.5	−1.0	−0.5
Blocks Ranking Size	−9.0	−11.5	−2.5	−2.0	−0.5
Open Laptop	−8.5	−7.0	−2.5	−2.0	−0.5
Place Object Basket	−5.0	−4.0	0.0	0.0	0.0
Place Dual Shoes	0.0	−5.0	−0.5	−0.5	−0.5
Adjust Bottle	−1.5	−3.5	0.0	0.0	0.0
Place Shoe	−1.5	−1.0	−1.0	−0.5	0.0
Rotate QRcode	+0.5	−2.5	−1.0	−0.5	0.0
Harmed tasks (all 50)	12/50	17/50	21/50	18/50	10/50
Median (all 50)	+0.5	0.0	0.0	0.0	0.0

Does correction target the harmed tasks? For the always-future adapter, define the task-level recovery $R_t = J_t(\pi_{\mathrm{w+adapter}}) - J_t(\pi_{\mathrm{w}})$. We compare mean recovery for harmed ($\Delta J_t(\pi_{\mathrm{w}}) < 0$) and non-harmed tasks, that is, whether correction preferentially repairs tasks harmed by predicted futures or shifts all tasks uniformly.

Table 10: **Adapter recovery on harmed and non-harmed tasks.** Harmed tasks have $\Delta J_t < 0$ under Shared WAM (w). Recovery values are percentage points; positive values indicate improvement.

Benchmark	Recovery, harmed	Recovery, non-harmed
LIBERO	+1.43	+0.36
RoboTwin 2.0	+3.55	+0.36

On both benchmarks, harmed tasks recover more than non-harmed tasks. On RoboTwin 2.0 the difference comes from the two severe tasks; the other 19 harmed tasks lose at most 2.5 points and recover $+0.34$ points on average, matching the $+0.36$ on non-harmed tasks. Correction therefore repairs severe harm where it occurs and otherwise acts as a small uniform gain, consistent with

Section 4.3. The harmed and non-harmed means are computed on all trials and may be inflated by regression to the mean; the externally selected tasks below avoid this selection.

Externally identified harmed tasks. Selecting tasks by their observed ΔJ_t and then measuring recovery on the same trials invites regression to the mean. As a complementary check, we examine the two severe tasks, which FAST-WAM's independently trained and evaluated per-task results already identify as severely harmed (Table 8), so their recovery is not inflated by selecting on our own evaluation noise. The adapter removes 89% and 59% of their losses (Table 9). The gate plays a complementary role: it improves benchmark-level success at lower cost (Section 4.3), while on these two tasks the residual losses change from 6.0 to 6.5 and from 14.0 to 17.0 points. Their recovery therefore comes from correction, consistent with the gate selecting between complete action paths rather than repairing the future-conditioned path. The gate's gain instead comes from the remaining 48 RoboTwin 2.0 tasks, where correction changes little: relative to Shared WAM (wo), they gain 1.0 point under uncorrected always-video inference, 1.35 under always-on correction, and about 1.9 under selective deployment, despite far less video generation (see Table 3 for the complete ablation breakdown).

C.4 Inference Efficiency

We measure end-to-end inference latency from CPU inputs to CPU action outputs on a single H100 80GB GPU, using BF16 and a batch size of 1. Table 11 reports latency and N_{eff}. Compared with the shared WAM that always generates future video, Utility-WAM reduces latency by 6.7% on LIBERO and 37.9% on RoboTwin 2.0. Future-free inference remains faster: relative to Shared WAM (wo-video), Utility-WAM adds 59% latency on LIBERO for $+1.5$ points and 17% on RoboTwin 2.0 for $+1.4$ points, compared with 70% and 89% for always generating futures, which gains $+0.8$ points on LIBERO and loses 0.8 points on RoboTwin 2.0.

Table 11: Mean video generation steps and inference latency per action query. Bold and underlined values denote the best and second-best results, respectively (lower is better).

	LIBERO		RoboTwin 2.0	
Method	N_{eff}	Latency (ms)↓	N_{eff}	Latency (ms)↓
FAST-WAM	**0.00**	<u>275.5</u>	**0.00**	**279.0**
Shared WAM (wo-video)	**0.00**	**274.4**	**0.00**	<u>280.0</u>
Shared WAM (w-video)	10.00	467.5	10.00	528.7
Utility-WAM	<u>8.70</u>	436.3	<u>2.45</u>	328.2

D Real-World Experimental Details and Qualitative Results

D.1 Real-World Hardware and Task Details

Hardware Platform. As shown in Figure 3, our real-world platform consists of two six-DoF AgileX PiPER manipulators, each equipped with a two-finger gripper. The observation system comprises one fixed front-view RGB camera and one wrist-mounted RGB camera on each arm. The robot state is represented by a 14-dimensional vector containing the six joint positions and gripper opening of each arm.

Visual Observations. The recorded images for *Place Empty Cup* and *Stack Three Blocks* have a resolution of 640×480 pixels for all three views. For *Insert Can*, the recorded front-view images have a resolution of 848×480 pixels, while both wrist views have a resolution of 640×480 pixels. All resolutions are reported as width $\times$ height. In the training pipeline, each view is resized to 320×240 pixels, and the three views are combined into a multi-view image of 320×384 pixels.

Action Representation and Execution. The policy uses a 14-dimensional absolute-position action representation comprising the joint and gripper targets of both arms. The training data are sampled at 10 Hz, and each training window contains 32 action steps, corresponding to a 3.2 s horizon.

Workspace and Initialization. All tasks are performed on a tabletop within the reachable workspace of the bimanual platform. The manipulated objects and their targets are placed in task-relevant regions that permit the required grasping, transport, placement, and insertion motions. Initial object placements are varied within these regions across trials while preserving task feasibility and the intended object–target relationships.

Tasks and Success Criteria. The initial task configurations are illustrated in Figures 3b–3d. For *Place Empty Cup*, the robot must grasp an empty cup and place it upright on the designated coaster. Success requires the cup to remain stably supported by the coaster after gripper release. For *Insert Can*, the robot must grasp a can and insert it into the designated recess in a case. Success requires the can to be seated in the recess and remain upright after release; a can resting on the rim does not satisfy the task objective. For *Stack Three Blocks*, the robot must sequentially arrange three blocks into a single tower. Success requires the completed tower to remain stable without gripper support.

D.2 Demonstration Collection

Demonstration Collection. We collect 200 demonstrations per task using AgileX's official teleoperation setup[2]. During teleoperation, we record RGB observations from the fixed front-view camera and both wrist-mounted cameras, together with the measured joint positions and gripper openings of both arms. Each demonstration is associated with a task instruction. The recorded robot state is a 14-dimensional vector containing six joint positions and one gripper opening per arm.

D.3 Evaluation Protocol and Representative Rollouts

Evaluation protocol. Each configuration in Table 5 is evaluated on 20 trials per task with varied initial layouts. Success requires stable placement on the coaster, a seated upright can, or a stable three-block tower, respectively. Table 5 reports the per-task success counts.

Figure 4 presents keyframe sequences extracted from successful real-world evaluation videos of Utility-WAM. For *Place Empty Cup*, the robot grasps the cup, transfers it to the designated coaster, and releases it in an upright position. For *Insert Can*, the robot picks up the can, moves it toward the case, aligns it with the target recess, and completes the insertion before withdrawing the gripper. For *Stack Three Blocks*, the robot first places the green block on the red block and then adds the blue block on top, forming a stable three-block tower. The final frame of each sequence shows the completed task configuration after gripper withdrawal. These sequences illustrate successful executions but are not comparative evidence.

[2] https://github.com/agilexrobotics/openpi-agilex

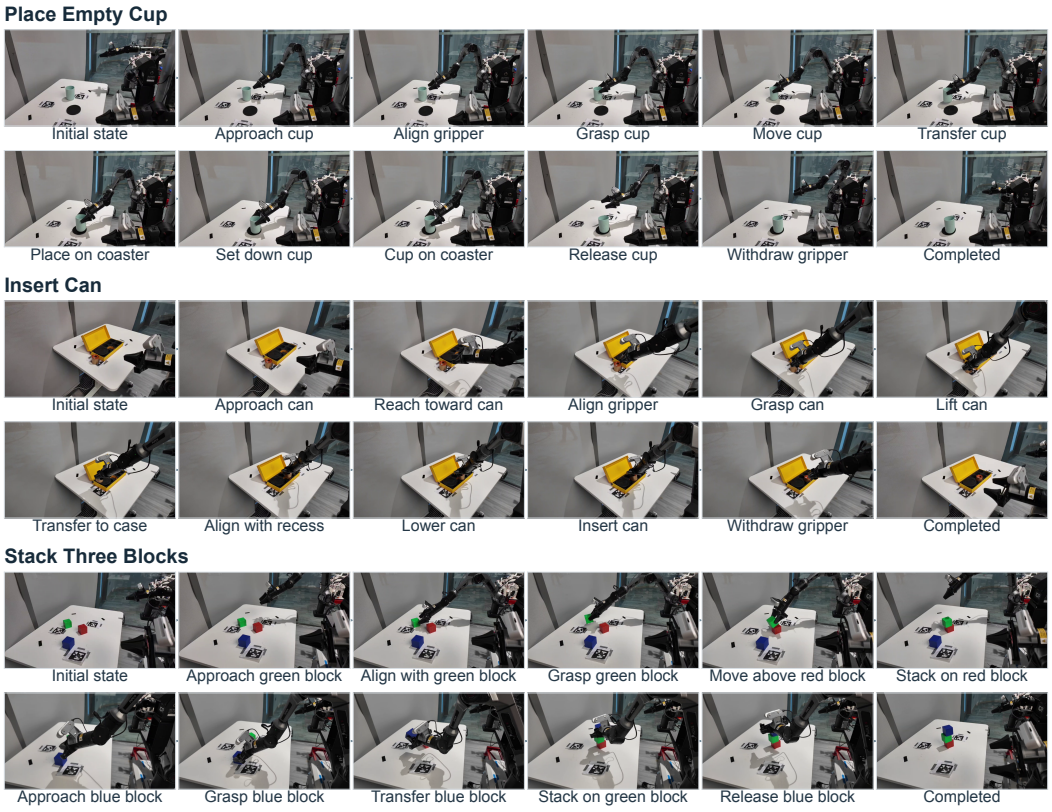

Figure 4: **Representative successful real-world evaluation sequences.** Keyframes for *Place Empty Cup*, *Insert Can*, and *Stack Three Blocks* are shown from top to bottom and progress from left to right, illustrating the main manipulation stages and the final task configurations after gripper withdrawal. These examples are qualitative and are not used for method attribution.